Oracle® Enterprise Manager Database Lifecycle Management Administrator's Guide





Oracle Enterprise Manager Database Lifecycle Management Administrator's Guide, 24ai Release 1 (24.1)

F97207-03

Copyright © 2015, 2025, Oracle and/or its affiliates.

Primary Author: Oracle Corporation

Contributors: Enterprise Manager 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, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

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

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

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

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

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

Contents

	Preface	
	Audience	XXV
	Documentation Accessibility	XXV
	Related Resources	XXV
	Conventions	XXV
Par	t I Introduction to Enterprise Manager Database Lifecycle Manag	ement
1	Introduction to Database Lifecycle Management	
	Information Map for Lifecycle Management Solutions	1-1
	Overview of the Lifecycle Management Solutions	1-2
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	2-5
	Overview of User Accounts	2-6
	Creating Designer User Account	2-7
	Creating Operator User Account	2-7
	(Optional) Setting Up My Oracle Support	2-8
	(Optional) Configuring Self-Update	2-8
	(Optional) Setting Restricted Accesses for the Root Components	2-9
	Security Models for running Root scripts	2-9
	Manually Staging the Root Components For Fleet Maintenance (Patching)	2-9
	Restricting the Root User Access For Fleet Maintenance (Patching)	2-11
	Root Components for Provisioning	2-11
	Manually Staging the Root Components for Provisioning	2-12
	Restricting the Root User Access	2-13
	(Ontional) Deploying Agents on OCI Passuress	2-1/



Part II Database Provisioning

Overview of Database Provisioning	
Introduction to Database Provisioning	3-1
Supported Use Cases and Targets Provisioned Using Database Provisioning Proced	
Setting Up Database Provisioning	3-5
Meeting Basic Infrastructure and Host Requirements	3-6
Understanding Administrator Privileges for Provisioning Database	3-6
Prerequisites for Designers	3-7
Prerequisites for Operators	3-9
Creating Database Provisioning Profiles	3-9
Viewing, Creating, and Deleting Database Provisioning Profiles Using EMCLI	3-11
Viewing Database Provisioning Profiles Using EMCLI	3-12
Creating Database Provisioning Profiles Using EMCLI	3-12
Deleting Database Provisioning Profiles Using EMCLI	3-12
Creating Installation Media	3-13
Creating Database Templates	3-14
Uploading Database Templates to Software Library	3-15
Pre/Post Scripts for Database Provisioning	3-16
Pre/Post Script Examples for Database Provisioning	3-31
Creating Database Provisioning Entities	3-34
Creating an Entity for Oracle Database Software from a Shared NFS location	n 3-34
Creating an Oracle Database Clone from an External Storage	3-35
Creating an Oracle Clusterware Clone from a Reference Home	3-36
Creating an Oracle Clusterware Clone from an External Storage	3-37
Downloading Cluster Verification Utility	3-38
Creating and Provisioning Oracle Databases	
Oracle Database Topology	4-1
Creating Databases	4-3
Getting Started with Creating Databases	4-3
Creating an Oracle Database	4-4
Prerequisites for Creating an Oracle Database	4-4
Procedure for Creating an Oracle Database	4-5
Creating Oracle Real Application Clusters Database	4-9
Prerequisites for Creating an Oracle Real Application Clusters Database	4-9
Procedure for Creating an Oracle Real Application Clusters Database	4-10
Creating Oracle Real Application Clusters One Node Database	4-12



Prerequisites for Creating an Oracle RAC One Node Database	4-13
Procedure for Creating an Oracle Real Application Clusters One Node Database	4-14
Provisioning Databases	4-16
Getting Started with Provisioning Oracle Databases	4-16
Provisioning and Creating Oracle Databases	4-17
Prerequisites for Provisioning Databases	4-18
Procedure for Provisioning Databases	4-18
Provisioning Oracle Databases with Oracle Automatic Storage Management	4-23
Prerequisites for Provisioning Oracle Databases with Oracle Automatic Storage Management	4-23
Procedure For Provisioning Databases with Oracle Automatic Storage Management	4-23
Provisioning Oracle Database Software Only	4-28
Prerequisites for Provisioning Oracle Database Software Only	4-29
Procedure for Provisioning Oracle Database Software Only	4-29
Provisioning Oracle Database Software from a Shared NFS location	4-31
Using Non-Root Credentials for Provisioning Oracle Databases	4-32
Getting Started with Provisioning Oracle Grid Infrastructure for Oracle Databases Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle Automatic	5-1
Storage Management	5-2
Prerequisites for Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle ASM	5-2
Procedure for Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle ASM	5-2
Provisioning Oracle Grid Infrastructure and Oracle Database Software Only	5-9
Prerequisites for Provisioning Oracle Grid Infrastructure and Oracle Database Software	
Only	5-9
Procedure for Provisioning Oracle Grid Infrastructure and Oracle Database Software	
Only	5-9
	5-9
Provisioning Oracle Grid Infrastructure for Oracle Real Application Clusters Databases Getting Started with Provisioning Grid Infrastructure for Oracle RAC Databases	6-1
Provisioning Oracle Grid Infrastructure for Oracle Real Application Clusters Databases Getting Started with Provisioning Grid Infrastructure for Oracle RAC Databases Oracle Real Application Clusters Database Topology	6-1 6-3
Provisioning Oracle Grid Infrastructure for Oracle Real Application Clusters Databases Getting Started with Provisioning Grid Infrastructure for Oracle RAC Databases Oracle Real Application Clusters Database Topology Oracle Flex Clusters and Flex ASM	6-1
Provisioning Oracle Grid Infrastructure for Oracle Real Application Clusters Databases Getting Started with Provisioning Grid Infrastructure for Oracle RAC Databases Oracle Real Application Clusters Database Topology Oracle Flex Clusters and Flex ASM Provisioning Grid Infrastructure with Oracle Real Application Clusters Database and Configuring Database with Oracle Automatic Storage Management	6-1 6-3 6-4 6-4
Provisioning Oracle Grid Infrastructure for Oracle Real Application Clusters Databases Getting Started with Provisioning Grid Infrastructure for Oracle RAC Databases Oracle Real Application Clusters Database Topology Oracle Flex Clusters and Flex ASM Provisioning Grid Infrastructure with Oracle Real Application Clusters Database and Configuring Database with Oracle Automatic Storage Management Prerequisites for Provisioning Grid Infrastructure with Oracle RAC Database	6-1 6-3 6-4 6-4 6-5
Provisioning Oracle Grid Infrastructure for Oracle Real Application Clusters Databases Getting Started with Provisioning Grid Infrastructure for Oracle RAC Databases Oracle Real Application Clusters Database Topology Oracle Flex Clusters and Flex ASM Provisioning Grid Infrastructure with Oracle Real Application Clusters Database and Configuring Database with Oracle Automatic Storage Management	6-1 6-3 6-4 6-4



Provisioning Oracle Real Application Clusters Database with File System on an Existing Cluster	6-12
Prerequisites for Provisioning Oracle RAC Database with File System on an Existing	6-13
Cluster Procedure for Provisioning Oracle RAC with ASM or File System on an Existing Cluster	6-13
Provisioning Oracle Real Application Clusters Database with File System on a New Cluster	6-18
Prerequisites for Provisioning Oracle RAC Database with File System on a New Cluster	6-18
Procedure for Provisioning Oracle RAC Database with File System on a New Cluster	6-18
Using No Root Credentials for Provisioning Oracle Real Application Clusters (Oracle RAC)	0-10
Databases	6-25
Provisioning Oracle Real Application Clusters One (Oracle RAC On Node Databases	e)
Getting Started with Provisioning Oracle RAC One Node Databases	7-1
Deployment Procedures for Provisioning Oracle RAC One Node Databases	7-2
Provisioning Oracle RAC One Node Databases	7-2
Prerequisites for Provisioning Oracle RAC One Node Databases	7-2
Procedure for Provisioning Oracle RAC One Node Databases	7-4
Provisioning Oracle Real Application Clusters	
Getting Started with Provisioning Oracle Real Application Clusters for 12c and 18c	8-1
Core Components Deployed When Provisioning Oracle RAC	8-2
Cloning a Running Oracle Real Application Clusters	8-2
Prerequisites for Cloning a Running Oracle Real Application Clusters	8-3
Procedure for Cloning a Running Oracle Real Application Clusters	8-3
Provisioning Oracle Real Application Clusters Using Gold Image	8-9
Prerequisites for Provisioning Oracle Real Application Clusters Using Gold Image	8-9
Procedure for Provisioning Oracle Real Application Clusters Using Gold Image	8-10
Provisioning Oracle Real Application Clusters Using Archived Software Binaries	8-16
Prerequisites for Provisioning Oracle Real Application Clusters Using Archived Software Binaries	8-16
Procedure for Provisioning Oracle Real Application Clusters Using Archived Software	
Binaries	8-17
Sample Cluster Configuration File	8-23
Provisioning Oracle Real Application Clusters (Oracle RAC) Databases Using No Root Credentials	8-23
Extending Oracle Real Application Clusters	
Getting Started with Extending Oracle Real Application Clusters	9-1
Extending Oracle Real Application Clusters	9-1
Prerequisites for Extending Oracle Real Application Clusters	9-1



10	Deleting or Scaling Down Oracle Real Application Clusters				
	Getting Started with Deleting or Scaling Down Oracle Real Application Clusters	10-1			
	Deleting the Core Components of Oracle Real Application Clusters	10-2			
	Deleting the Entire Oracle RAC	10-2			
	Prerequisites for Deleting the Entire Oracle RAC	10-2			
	Procedure for Deleting the Entire Oracle RAC	10-3			
	Scaling Down Oracle RAC by Deleting Some of Its Nodes	10-5			
	Prerequisites for Scaling Down Oracle RAC by Deleting Some of Its Nodes	10-5			
	Procedure for Scaling Down Oracle RAC by Deleting Some of Its Nodes	10-5			
11	Provisioning Oracle Database Replay Client				
	Getting Started with Provisioning Oracle Database Replay Client	11-1			
	Cloning a Running Oracle Database Replay Client	11-2			
	Prerequisites for Cloning a Running Oracle Database Replay Client	11-2			
	Procedure for Cloning a Running Oracle Database Replay Client	11-3			
	Provisioning an Oracle Database Replay Client Using Gold Image	11-5			
	Prerequisites for Provisioning an Oracle Database Replay Client Using Gold Image	11-5			
	Procedure for Provisioning an Oracle Database Replay Client Using Gold Image	11-6			
	Provisioning an Oracle Database Replay Client Using Installation Binaries	11-8			
	Prerequisites for Provisioning an Oracle Database Replay Client Using Installation Binaries	11-8			
	Procedure for Provisioning an Oracle Database Replay Client Using Installation Binaries	11-9			
12	Provisioning Oracle Standby Databases				
	Overview of Creating a Standby Database	12-1			
	Creating a New Physical Standby Database (single-instance only)	12-1			
	Step 1: Determine the backup type	12-3			
	Step 2: Set up the backup options	12-3			
	Step 3: Select the Oracle home in which to create the standby database	12-3			
	Step 4: Set up the location for standby database files	12-4			
	Step 5: Provide standby database configuration parameters	12-4			
	Step 6: Review the information before clicking Finish	12-4			
	Creating a New Logical Standby Database (single-instance only)	12-5			
	Step 1: Determine the backup type	12-6			
	Step 2: Set up the backup options	12-6			
	Step 3: Select the Oracle home in which to create the standby database	12-6			
	Step 4: Set up the location for standby database files	12-7			
	Step 5: Provide standby database configuration parameters	12-7			



	Step 6. Review the information before clicking Fillish	12-7
	Managing an Existing Standby Database with Data Guard Broker	12-8
	Creating a Primary Database Backup Only	12-9
	Data Guard	12-9
	Fast-Start Failover REST API	12-13
	Fast-Start Failover Management REST API	12-13
	Observer Management REST API	12-22
13	Managing Pluggable Databases Using Enterprise Manager	
	Getting Started With Managing Pluggable Databases Using Enterprise Manager	13-1
	Overview of Managing Pluggable Databases Using Enterprise Manager	13-2
	Provisioning Pluggable Databases Using Enterprise Manager	13-3
	Creating a New Pluggable Database Using Enterprise Manager	13-4
	Prerequisites for Creating a New Pluggable Database	13-4
	Creating a New Pluggable Database	13-4
	Plugging In an Unplugged Pluggable Database Using Enterprise Manager	13-10
	Prerequisites for Plugging In an Unplugged Pluggable Database	13-10
	Plugging In an Unplugged Pluggable Database	13-11
	Cloning a Pluggable Database Using Enterprise Manager	13-19
	Prerequisites for Cloning a Pluggable Database	13-20
	Cloning a Pluggable Database	13-21
	Migrating a Non-CDB as a Pluggable Database Using Enterprise Manager	13-31
	Prerequisites for Migrating a Non-CDB as a Pluggable Database	13-31
	Migrating a Non-CDB as a Pluggable Database	13-32
	Relocating Pluggable Databases Using Enterprise Manager	13-35
	Prerequisites for Relocating Pluggable Databases	13-36
	Relocating Pluggable Databases	13-36
	Removing Pluggable Databases Using Enterprise Manager	13-37
	Unplugging and Dropping a Pluggable Database Using Enterprise Manager	13-37
	Prerequisites for Unplugging and Dropping a Pluggable Database	13-38
	Unplugging and Dropping a Pluggable Database	13-38
	Deleting Pluggable Databases Using Enterprise Manager	13-42
	Prerequisites for Deleting Pluggable Databases	13-43
	Deleting Pluggable Databases	13-43
	Viewing Pluggable Database Job Details Using Enterprise Manager	13-47
	Viewing Create Pluggable Database Job Details	13-47
	Viewing Unplug Pluggable Database Job Details	13-48
	Viewing Delete Pluggable Database Job Details	13-49
	Administering Pluggable Databases Using Enterprise Manager	13-50
	Switching Between Pluggable Databases Using Enterprise Manager	13-51



14	Cloning	Oracle	Databases	and	Pluggable	Databases

Creating a Full Clone Database	14-1
Creating a Full Clone Database Using the Clone Wizard	14-1
Creating a Full Clone Database Using EM CLI	14-4
Refresh a Cloned Database	14-7
Cloning an Oracle Database Using Recovery Manager (RMAN) Backup	14-7
Creating a Full Clone Pluggable Database	14-9
Creating a Full Clone Pluggable Database Using the Clone Wizard	14-9
Creating a Full Clone Pluggable Database Using EMCLI	14-12
About Deployment Procedures	
Overview of the Provisioning Page	15-1
Granting Roles and Privileges to Administrators	15-3
Granting Roles and Privileges to Administrators on the Deployment Procedure	15-4
Granting Roles and Privileges to Administrators on Software Library	15-5
Components of a Procedure	15-6
Target List	15-6
Procedure Variables	15-7
Phases and Steps	15-8
Types of Phases	15-8
Types of Procedure Steps	15-8
Performing Tasks on Procedure Steps	15-10
Creating a Procedure	15-11
Adding Rolling or Parallel Phase	15-11
Adding Steps	15-12
Managing Deployment Procedures	15-17
Viewing, Editing, Deleting a Procedure	15-17
Editing and Saving Permissions of a Procedure	15-17
Tracking the Procedure Execution and Status of Deployment Procedures	15-18
Rescheduling a Procedure	15-19
Reverting a Procedure	15-19
Setting Step Level Grace Period	15-19
Creating, Saving, and Launching User Defined Deployment Procedure (UDDP)	15-20
Step 1: Creating User Defined Deployment Procedure	15-21
Step 2: Saving and Launching User Defined Deployment Procedure with Default Inputs	15-21
Step 3: Launching and Running the Saved User Defined Deployment Procedure	15-29
Step 4: Tracking the Submitted User Defined Deployment Procedure	15-30



	Procedure Instance Execution Page	15-30
	Comparison Between the Existing Design and the New Design for Procedure Instance	
	Execution Page	15-31
	Overview of the Procedure Instance Execution Page	15-32
	Investigating a Failed Step for a Single or a Set of Targets	15-34
	Retrying a Failed Step	15-34
	Creating an Incident Dump	15-35
	Viewing the Execution Time of a Deployment Procedure	15-36
	Searching for a Step	15-36
	Downloading a Step Output	15-36
	Accessing the Job Summary Page	15-36
Part	Database Maintenance: Updates (Patching) and Upgrades	
16	Database Maintenance Concepts	
	Patch Management Overview	16-1
	Understanding the Patching Workflow	16-1
17	Getting Started with Updates (Patching) and Upgrades	
	Overview of Updates and Upgrades	17-1
	About Database Fleet Maintenance	17-2
	Before you begin: Meeting the Prerequisites	17-5
	Meeting Basic Infrastructure Requirements for Patching	17-5
	User Roles and Privileges for Fleet Maintenance	17-5
	Preferred Credential Support for Fleet Maintenance	17-6
	Searching for Patches on My Oracle Support	17-7
	Search and View Patches in Oracle Software Library	17-8
	Downloading Patch Recommendations and Patches	17-8
	Setting up MOS Connection for Patch Recommendations and Patch Downloads	17-11
	Setting Up the Infrastructure for Patching in Offline Mode (Not Connected to MOS)	17-12
	Security Models for running Root scripts	17-16
	Prerequisites for performing Fleet Maintenance on Windows	17-16
18	Database Updates and Upgrades Using Fleet Maintenance	
	Get Organized With Standardization Advisor	18-1
	Use Fleet Maintenance Hub to get to Patch Compliance	18-3
	Patch Recommendations	18-7
	Image Maintenance	18-8
	Image Management	18-11



	Update Operations	18-12
	Update Databases Using Fleet Maintenance	18-12
	Track, Fix and Resume Update Procedures	18-17
	Use of Oracle Home Path Variables	18-18
	Custom Pre and Post Scripts for Fleet Operations	18-21
	Custom Script Considerations	18-21
	Software Library Entity For Uploading The Script	18-23
	How to Get URN for SWLIB Entity	18-23
	Pre/Post Script Usage	18-23
	Fleet Maintenance Automation	18-24
	Database Fleet Maintenance Update REST APIs	18-24
	Fleet Maintenance EM CLI Operations	18-25
	Image Maintenance	18-28
	Subscription Tasks	18-41
	Deploy Operations	18-45
	Migrate the Listeners	18-53
	Rollback Operations	18-55
	Cleanup Operations	18-66
	Emergency Patching	18-71
	Update Operations	18-74
	Patch History	18-82
19	Upgrade databases and Grid Infrastructure Using Fleet Maintenance	
	Plan Your Upgrades	19-1
	Prepare Your Upgrade	19-2
	Upgrading Databases using AutoUpgrade	19-3
	Data Guard Rolling Upgrade	19-4
	About Rolling Upgrades	19-4
	Prerequisites to Rolling Upgrades	19-4
	Submitting a Rolling Upgrade Procedure for a Primary Database With One Physical Standby Database	19-5
	Viewing a Running or Completed Rolling Upgrade Procedure in Enterprise Manager	19-8
	Downgrade Databases	19-9
20	Oracle Fleet Patching and Provisioning and Enterprise Manager Integration	
	FPP for Enterprise Manager Prerequisites	20-2
	FPP for Enterprise Manager Prerequisites FPP Setup Automation	20-2 20-3
	FPP Setup Automation	20-3



	Deploy a Gold Image and Patch Using FPP	20-13
	FPP Patching Using Fleet Maintenance UI	20-13
	Single Instance Database Patching Using FPP	20-13
	Deploy and Patch Independently, Starting with a Standby Database	20-14
	Deploy and Patch Together, Starting with a Standby Database	20-17
	FPP Patching Using EM CLI	20-19
	FPP Patching Using REST API	20-20
	Opt Out of the FPP Series	20-20
21	Patch Exadata PDBs Using Fleet Maintenance	
22	Patching End-to-End Use Cases	
	End-to-End Use Case: Patching Your Data Center	22-1
	The Challenge of Patching Your Data Center	22-1
	The Enterprise Manager Solution	22-1
	Executing the Example Scenario	22-2
	Create Administrators with the Required Roles	22-2
	Set Up the Infrastructure	22-2
	Analyze the Environment and Identify Patches to be Applied	22-3
	Identify the Relevant Patches	22-3
	Execute Database Patching with Fleet Maintenance	22-3
	Check the Status of the Patching Activities	22-3
	Verify If the Targets Have Been Patched	22-3
	Fleet Maintenance End to End Use Cases	22-3
Part	IV Manage Configuration, Compliance and Change Management	
23	Manage Configuration Information	
	Configuration Management	23-1
	Comparisons	23-3
	About Comparisons	23-3
	About Comparison Templates	23-4
	About One-Time Comparisons	23-4
	About Configuration Drift	23-5
	About Configuration Consistency	23-7
	About the Definition Library	23-8
	Get Started Using Comparisons	23-9
	Creating or Editing a Comparison Template	23-9



Set Up a Comparison Template	23-11
Managing Comparison Templates	23-12
Considerations Before Creating a Comparison	23-13
Set Up a Drift or Consistency Comparison	23-13
Create Notifications for Comparisons	23-14
Exclude Target To Be Compared In a Template	23-14
Exclude Target To Be Compared From Reference Target	23-15
Oracle Home Consistency With Target Properties	23-16
Get Started Using Rules	23-17
About Rules Expression and Syntax	23-18
Create a Value Constraint Rule	23-19
Create a Matching Rule	23-20
Create a Rule for Including and Excluding Configuration Items	23-20
Rule Examples	23-21
Work with Comparison Results	23-23
About Consistency Management (System) Comparison Results	23-23
About Drift (Target) Comparison Results	23-24
Synchronizing Configuration Extension Files	23-24
Comparisons and Drift Management with Oracle Enterprise Manager	23-26
About the Dashboard	23-26
Create a Drift Definition	23-28
Create a Consistency Definition	23-29
Create a One-Time Comparison	23-31
Manage the Definition Library	23-32
Manage Comparison Templates	23-33
Compare Configurations with Scripts	23-33
Configuration Searches	23-34
Managing Configuration Searches	23-34
Searching for a Configuration Search	23-35
Running a Configuration Search	23-35
Editing a Configuration Search	23-36
Deleting a Configuration Search	23-36
Importing or Exporting a Configuration Search	23-36
Creating a Configuration Search	23-36
Creating a New Configuration Search	23-37
Creating a Configuration Search from an Existing Configuration Search	23-38
Creating a Configuration Search Using SQL	23-38
Configuration Browser	23-38
Viewing Configuration Data	23-39
Working with Saved Configurations	23-41
Working with Inventory and Usage Details	23-42
Configuration History	23-43



Accessing Configuration History	23-44
Working with Configuration History	23-44
Searching History	23-44
Annotating Configuration Changes	23-46
Scheduling a History Search and Creating a Notification List	23-46
Saving History to a File	23-47
Saving Configuration History	23-47
Creating a Search Using SQL	23-47
Viewing History Job Activity	23-47
Configuration Extensions and Collections	23-48
Working with Configuration Extensions	23-48
Creating a Custom Target Type	23-48
Creating or Editing a Configuration Extension	23-49
Using the Files & Commands Tab	23-50
Using the SQL Tab	23-51
Setting Up Credentials When Creating a Configuration Extension	23-51
Setting Up Rules	23-52
Managing Configuration Extensions	23-52
About Configuration Extensions and Versioning	23-53
About Configuration Extensions and Privileges	23-54
About Configuration Extensions and Deployment	23-54
Deploying and Undeploying Configuration Extensions	23-55
Editing a Deployment of Configuration Extensions	23-56
Viewing a Configuration Collection	23-56
Extending Configuration Data Collections	23-56
Extending Existing Target Collections	23-57
Adding New Target Data Collections	23-57
Configuration Topology Viewer	23-58
About Configuration Topology Viewer	23-59
Examples of Using Topology	23-59
Viewing a Configuration Topology	23-60
Determining System Component Structure	23-60
Determining General Status of Target's Configuration Health	23-61
Getting Configuration Health/Compliance Score of a Target	23-61
Analyzing a Problem and Viewing a Specific Issue in Detail	23-61
About Dependency Analysis	23-62
About Impact Analysis	23-62
Creating a Custom Topology View	23-63
Deleting a Custom Topology View	23-63
Excluding Relationships from a Custom Topology View	23-64
Including Relationships to a Target in a Custom Topology View	23-64
Creating a Relationship to a Target	23-64



Deleting a Relationship from a Target	23-65
Controlling the Appearance of Information on a Configuration Topology Graph	23-65
Parsers	23-66
Managing Parsers	23-67
About XML Parsers	23-68
About the Default XML Parser	23-68
About the Generic XML Parser	23-69
XML Parser Examples	23-69
About Format-Specific Parsers	23-71
Database Query Parser Parameters	23-72
Database Query Paired Column Parser Parameters	23-72
Directory Parser Parameters	23-73
E-Business Suite Parser Parameters	23-73
Galaxy CFG Parser Parameters	23-73
MQ-Series Parser Parameters	23-74
Siebel Parser Parameters	23-74
Unix Installed Patches Parser Parameters	23-74
Unix Recursive Directory List Parser Parameters	23-74
About Columnar Parsers	23-75
Columnar Parser Parameters	23-76
About Properties Parsers	23-77
Basic Properties Parser Parameters	23-78
Advanced Properties Parser Parameters	23-79
Advanced Properties Parser Constructs	23-80
Using Parsed Files and Rules	23-84
Sample XML File Parsing and Rule Application	23-84
Sample Non-XML File Parsing and Rule Application	23-85
Sample SQL Query Parsing and Rule Application	23-87
Relationships	23-89
Manage Compliance	
Overview of Compliance Management	24-1
Terminology Used in Compliance	24-2
Access the Compliance Features	
Roles and Privileges Needed for Compliance Features	24-4
Configure Compliance Management	
About Compliance Frameworks	24-7
Operations on Compliance Frameworks	24-8
Creating a Compliance Framework	24-9
Creating Like a Compliance Framework	24-10
Editing a Compliance Framework	24-10



Deleting a Compliance Framework	24-11
Exporting a Compliance Framework	24-12
Importing a Compliance Framework	24-12
Browsing Compliance Frameworks	24-13
Searching Compliance Frameworks	24-13
Browsing Compliance Framework Evaluation Results	24-13
Searching Compliance Framework Evaluation Results	24-13
Browsing Compliance Framework Errors	24-14
Searching Compliance Framework Errors	24-14
Verifying Database Targets Are Compliant with Compliance Frameworks	24-14
About Compliance Standards	24-15
Operations on Compliance Standards	24-17
Creating a Compliance Standard	24-18
Associating a Compliance Standard	24-19
Self Update for Compliance Standards	24-20
Create Like for Compliance Standards	24-21
Editing a Compliance Standard	24-21
Deleting a Compliance Standard	24-21
Exporting a Compliance Standard	24-22
Importing a Compliance Standard	24-22
Browsing Compliance Standards	24-23
Searching Compliance Standards	24-23
Browsing Compliance Standard Evaluation Results	24-23
Searching Compliance Standard Evaluation Results	24-23
Browsing Compliance Standard Errors	24-23
Searching Compliance Standard Errors	24-24
Associating a Compliance Standard with Targets	24-24
Associating a Compliance Standard with a Group Target	24-25
Enabling Security Metrics	24-26
Considerations When Creating Compliance Standards	24-26
About Compliance Standard Rule Folders	24-27
Creating Rule Folders	24-27
Managing Rule Folders in a Compliance Standard	24-27
About Compliance Standard Rules	24-28
Operations on Compliance Standards Rules	24-29
Creating a Repository Compliance Standard Rule	24-30
Creating an Agent-side Rule	24-32
Creating a Manual Rule	24-33
Creating a Missing Patches Compliance Standard Rule	24-34
Creating a Configuration Consistency Rule	24-37
Creating Configuration Drift Rule	24-38
Creating Like a Compliance Standard Rule	24-39



Editing a Compliance Standard Rule	24-39
Deleting a Compliance Standard Rule	24-40
Exporting a Compliance Standard Rule	24-40
Importing a Compliance Standard Rule	24-41
Browsing Compliance Standard Rules	24-41
Searching Compliance Standard Rules	24-41
Using Corrective Actions	24-42
Evaluate Compliance	24-43
Access Compliance Statistics	24-44
Use the Compliance Dashboard Effectively	24-45
View Compliance Summary Information	24-46
View Target Compliance Evaluation Results	24-46
View Compliance Framework Evaluation Results	24-47
Manage Violations	24-47
Investigate Compliance Violations and Evaluation Results	24-49
Investigating Violations of Repository Compliance Standard Rules and Targe Causing Violations	ts 24-49
Viewing All the Violations Reported for Your Enterprise	24-50
Examples of Viewing Violations	24-50
Investigate Evaluation Errors	24-54
Analyze Compliance Reports	24-54
Overview of Compliance Score and Importance	24-55
Compliance Score of a Compliance Standard Rule -Target	24-56
Compliance Score of a Compliance Standard for a Target	24-56
Compliance Framework Compliance Score	24-57
Parent Node Compliance Score	24-57
About the Compliance Dashboard	24-58
Examples	24-63
Create Repository Rule Based on Custom Configuration Collections	24-63
Create Compliance Standard Agent-side and Manual Rules	24-66
Suppress Violations	24-74
Clear Violations	24-75
Manage Database Schema Changes	
Overview of Change Management for Databases	25-1
Using Schema Baselines	25-2
Overview of Scope Specification	25-3
About Capturing a Schema Baseline Version	25-3
About Working With A Schema Baseline Version	25-4
About Working With Multiple Schema Baseline Versions	25-4
Exporting and Importing Schema Baselines	25-4



	Creating Directory Objects for Export and Import	25-5
	Using Schema Comparisons	25-5
	Defining Schema Comparisons	25-6
	About Working with Schema Comparison Versions	25-8
	Using Schema Synchronizations	25-8
	About Defining Schema Synchronizations	25-8
	Creating a Synchronization Definition from a Comparison	25-11
	Working with Schema Synchronization Versions	25-11
	About the Schema Synchronization Cycle	25-11
	Creating Additional Synchronization Versions	25-14
	Using Change Plans	25-14
	About Working with Change Plans	25-15
	Creating a Change Plan	25-15
	Creating and Applying a Change Plan From a Schema Comparison	25-15
	Using External Clients to Create and Access Change Plans in Enterprise Manager	25-18
	Submitting Schema Change Plans From SQL Developer Interface	25-19
	Using Database Data Comparison	25-19
	Requirements for Database Data Comparisons	25-20
	Comparing Database Data and Viewing Results	25-22
	t V Database Security	
26	Managing Oracle Audit Vault and Database Firewall	
26	Managing Oracle Audit Vault and Database Firewall	
26 27		
26 27	Managing Oracle Audit Vault and Database Firewall	
26 27	Managing Oracle Audit Vault and Database Firewall	
26 27 28	Managing Oracle Audit Vault and Database Firewall	
27	Managing Oracle Audit Vault and Database Firewall Using Oracle Data Redaction	
27	Managing Oracle Audit Vault and Database Firewall Using Oracle Data Redaction	
27	Managing Oracle Audit Vault and Database Firewall Using Oracle Data Redaction Managing Oracle Database Vault and Privilege Analysis	29-2
27	Managing Oracle Audit Vault and Database Firewall Using Oracle Data Redaction Managing Oracle Database Vault and Privilege Analysis Managing Oracle Transparent Data Encryption	29-2
27	Managing Oracle Audit Vault and Database Firewall Using Oracle Data Redaction Managing Oracle Database Vault and Privilege Analysis Managing Oracle Transparent Data Encryption Prerequisites for Transparent Data Encryption Operations	
27	Managing Oracle Audit Vault and Database Firewall Using Oracle Data Redaction Managing Oracle Database Vault and Privilege Analysis Managing Oracle Transparent Data Encryption Prerequisites for Transparent Data Encryption Operations Encrypting a Tablespace in Online Mode	29-2
27	Managing Oracle Audit Vault and Database Firewall Using Oracle Data Redaction Managing Oracle Database Vault and Privilege Analysis Managing Oracle Transparent Data Encryption Prerequisites for Transparent Data Encryption Operations Encrypting a Tablespace in Online Mode Decrypting a Tablespace in Online Mode	29-2 29-3



30 Managing Enterprise Data Governance

Overview of Enterprise Data Governance	30-1
About Enterprise Data Governance	30-1
What Are Protection Policies?	30-1
What Are Application Signatures?	30-2
Using Enterprise Data Governance	30-2
The Enterprise Data Governance Dashboard	30-3
Working with Sensitive Database Discovery Results	30-3
Working with Metadata Discovery Jobs	30-3
Creating a Metadata Discovery Job	30-3
Managing Automatic Metadata Discovery	30-4
Managing Metadata Discovery Results	30-4
Working with Data Discovery Jobs	30-5
Creating a Data Discovery Job	30-5
Managing Data Discovery Results	30-5
Creating Custom Application Signatures	30-6
Appendixes	
Checking Host Readiness Before Provisioning or Patching	A-1
Setting Up User Accounts Before Provisioning	A-1
Configuring SSH	A-2
Shell Limits	A-2
Root Setup (Privilege Delegation)	A-2
Environment Settings	A-2
Kernel Requirements	A-3
Node Time Requirements	A-4
Package Requirements	A-4
Memory and Disk Space Requirements	A-4
Network & IP Address Requirements	A-5
Storage Requirements	A-6
Installation Directories and Oracle Inventory	A-6
Using Enterprise Manager Command Line Interface	A-7
Overview	A-7
Provisioning Using EM CLI	A-7
Creating the Properties File to Submit a Deployment Procedure	A-8
Using Properties File from an Existing Execution of a Deployment Procedure	A-11
Launching a Procedure using an Existing Saved Procedure	A-14
Provisioning Pluggable Databases	A-15



WorkFlow Examples Using EM CLI Commands	A-20
Provisioning Oracle Database Software	A-20
Provisioning Oracle WebLogic Server	A-22
Prerequisites for Provisioning Oracle WebLogic Server	A-22
Provisioning Oracle WebLogic Server Using the Provisioning Profile	A-22
Scaling Up or Scaling Out Middleware Deployment Procedure	A-24
Provisioning User Defined Deployment Procedure	A-29
Prerequisites for Provisioning User Defined Deployment Procedure	A-29
Adding Steps and Phases to User Defined Deployment Procedure Using GUI	A-29
Using EM CLI commands to Run an Instance of the Procedure	A-30
Creating a New Generic Component by Associating a Zip File	A-30
Step 1: Identifying the Parent Folder in Software Library	A-30
Step 2: Creating a Generic Component Entity	A-32
Step 3: Associating a Zip File to the Generic Component	A-34
Step 4: Verifying the Newly Created Entity	A-34
Migrate and Remove a Software Library Storage Location	A-35
Step 1: Adding a Destination Storage Location for Migrating Files	A-35
Step 2: Migrate and Remove an existing storage location	A-35
Adding ATS Service Test from Using EM CLI	A-36
Using emctl partool Utility	A-38
Overview of Provisioning Archive Files	A-38
Overview of emctl partool Utility	A-38
Checking Oracle Software Library	A-40
Exporting Deployment Procedures	A-40
Obtaining Deployment Procedure's GUID	A-40
Creating PAR File	A-41
Importing PAR Files	A-41
Importing Using Command Line Interface	A-42
Importing Using the Enterprise Manager Console	A-43
Troubleshooting Issues	A-43
Troubleshooting Patching Issues	A-43
Oracle Software Library Configuration Issues	A-44
My Oracle Support Connectivity Issues	A-45
Host and Oracle Home Credential Issues	A-46
User Account and Role Issues	A-47
Workaround for Errors	A-47
Oracle Fleet Patching and Provisioning: Diagnostic and Troubleshooting	A-49
Troubleshooting Database Provisioning Issues	A-51
Grid Infrastructure Root Script Failure	A-51
SUDO Error During Deployment Procedure Execution	A-51
Prerequisites Checks Failure	A-51
Oracle Automatic Storage Management (Oracle ASM) Disk Creation Failure	A-52



A-52
A-52
A-53
A-53
A-53
A-54
A-54
A-54
A-55
A-56
A-56
A-56

Index



List of Figures

2-1	Setting Up Your Infrastructure WorkFlow	2-2
2-2	Software Library Console	2-3
2-3	Software Library Administration	2-4
2-4	Add Security Rule	2-14
2-5	Privilege Delegation Setting	2-15
2-6	Add Host Target Installation Details	2-16
3-1	Database Provisioning Solution in Enterprise Manager	3-1
4-1	Oracle Database Topology	4-2
6-1	Oracle RAC Database Topology	6-3
13-1	Managing Pluggable Databases	13-3
13-2	Provision Pluggable Database Console	13-5
13-3	Creating a New Pluggable Database: Creation Options Page	13-6
13-4	Creating a New Pluggable Database: Identification Page	13-7
13-5	TDE Options	13-7
13-6	Creating a New Pluggable Database: Storage Page	13-8
13-7	Creating a New Pluggable Database: Schedule Page	13-9
13-8	Creating a New Pluggable Database: Review Page	13-9
13-9	Creating a New Pluggable Database: Procedure Activity Page	13-10
13-10	Provision Pluggable Database Console	13-12
13-11	Plugging In an Unplugged Pluggable Database: Creation Options Page	13-13
13-12	Plugging In an Unplugged Pluggable Database: Identification Page	13-14
13-13	Plugging In an Unplugged Pluggable Database: PDB Template Location Section	13-15
13-14	Plug in a TDE Enabled PDB	13-16
13-15	Plugging In an Unplugged Pluggable Database: Storage Page	13-17
13-16	Plugging In an Unplugged Pluggable Database: Schedule Page	13-18
13-17	Plugging In an Unplugged Pluggable Database: Review Page	13-19
13-18	Plugging In an Unplugged Pluggable Database: Procedure Activity Page	13-19
13-19	Provision Pluggable Database Console	13-22
13-20	Cloning a Pluggable Database: Specifying Creation Options	13-23
13-21	Cloning a Pluggable Database: Specifying CDB Host Credentials	13-23
13-22	Cloning a Pluggable Database: Identification Page	13-24
13-23	Cloning a Pluggable Database (Full Clone): Storage Page	13-25
13-24	Cloning a Pluggable Database (Snap Clone): Storage Page	13-26
13-25	Cloning a Pluggable Database: Schedule Page	13-27
13-26	Cloning a Pluggable Database: Review Page	13-27
13-27	Cloning a Pluggable Database: Procedure Activity Page	13-28



13-28	Provision Pluggable Database Console	13-32
13-29	Migrating a Non-CDB: Method Page	13-33
13-30	Migrating a Non-CDB: Selecting a Non-CDB	13-33
13-31	Migrating a Non-CDB: Specifying Credentials	13-34
13-32	Migrating a Non-CDB: Specifying Stage Location	13-34
13-33	Migrating a Non-CDB: Specifying PDB Properties and Datafile Location	13-34
13-34	Migrating a Non-CDB: Specifying Import Directory	13-34
13-35	Migrating a Non-CDB: Schedule Page	13-35
13-36	Migrating a Non-CDB: Review Page	13-35
13-37	Provision Pluggable Database Console	13-39
13-38	Unplugging a Pluggable Database: Select PDB Page	13-39
13-39	Unplugging a Pluggable Database: Destination Page	13-41
13-40	Unplugging a Pluggable Database: Schedule Page	13-42
13-41	Unplugging a Pluggable Database: Review Page	13-42
13-42	Provision Pluggable Database Console	13-44
13-43	Deleting Pluggable Databases: Specifying Credentials	13-44
13-44	Deleting Pluggable Databases: Schedule Page	13-45
13-45	Deleting Pluggable Databases: Review Page	13-46
13-46	Deleting Pluggable Databases: Procedure Activity Page	13-46
15-1	Accessing the Provisioning Page	15-3
15-2	Procedure Execution Page	15-32
17-1	Fleet Maintenance Operations Lifecycle Overview	17-2
17-2	Fleet Maintenance Roles	17-3
17-3	Rolling Mode of Patching	17-10
17-4	Parallel Mode of Patching	17-10
17-5	Searching for Patches	17-13
17-6	Downloading Patches from My Oracle Support	17-14
18-1	Fleet Maintenance Target Subscription	18-3
18-2	Subscribe Targets to Image	18-4
18-3	Patch Recomendations for Images	18-4
18-4	Target Patch Compliance with Images	18-6
18-5	Database Drift Use Case	18-73
20-1	FPP Workflow	20-1
22-1	Patching Your Datacenter Use Case	22-1
23-1	Member Settings	23-10
23-2	Template Settings	23-10
23-3	Comparison and Drift Management Dashboard	23-27



24-1	Compliance Standard Definition	24-15
24-2	Violations for a Compliance Standard	24-51
24-3	Violations Using the Target Compliance Tab	24-52
24-4	Show Details Page	24-52
24-5	Event Details and Guided Resolution	24-53
24-6	Compliance Summary Region on Enterprise Summary Page	24-53
24-7	How Compliance Score of a Compliance Standard-Target Is Calculated	24-57
24-8	Compliance Score of Parent Node	24-57
24-9	Compliance Dashboard with Engineered Systems Tab	24-58
24-10	Compliance Summary	24-60
24-11	Compliance Evaluation Report	24-61
24-12	Oracle Database Security Assessment Report	24-62
24-13	Oracle Database Sensitive Data Assessment	24-63
24-14	Completed Create Configuration Extension Page	24-68
24-15	Completed Compliance Standard Rule Details Page	24-69
24-16	Completed Compliance Standard Rule Check Definition Page	24-69
24-17	Completed Compliance Standard Rule Test Page	24-70
24-18	Completed Compliance Standard Rule Review Page	24-70
24-19	Completed Manual Rule Page	24-72
24-20	Completed Create Compliance Standard Pop-Up	24-72
24-21	Compliance Standard Rules	24-73
24-22	Compliance Standards Library Page	24-73
24-23	Completed Target Association Page	24-73
24-24	CS1 - DB Check Compliance Standard in Evaluation Results Tab	24-74
24-25	Manage Violations Page - Unsuppressed Violations	24-74
24-26	Evaluation Results Page After Violation Is Suppressed	24-75
24-27	Manage Violations Page Showing the Suppressed Violations Tab	24-75
24-28	Clearing Manual Rule Violations	24-76
25-1	Steps in a Change Plan	25-15
27-1	Data Redaction Page	27-1
29-1	Transparent Data Encryption Home Page	29-1



Preface

Enterprise Manager Database Lifecycle Management Administrator's Guide familiarizes you with Oracle Enterprise Manager database lifecycle management solutions. It enables you to employ its services for tasks like provisioning databases, applying patches for updates and upgrades to databases and grid infrastructure, monitoring configuration and compliance, and ensuring database security in both on-premises data centers and the cloud.

Audience

This guide is intended for administrators seeking to address the complexities of database lifecycle management, specifically in the realms of discovery, provisioning, patching (updates and upgrades), and monitoring. This guide caters to administrators assuming two distinct roles: designer and operator.

In the designer role, an administrator takes on the role of a system administrator, carrying out pivotal database operations such as Image Creation and Image Refresh. As an operator, an administrator executes update and upgrade of databases and grids, monitor configuration and compliance of targets within Enterprise Manager.

Documentation Accessibility

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

Access to Oracle Support

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

Related Resources

For more information, see the documents available in the Oracle Enterprise Manager documentation library: Enterprise Manager Documentation.

Conventions

The following conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.



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



Part I

Introduction to Enterprise Manager Database Lifecycle Management

This part contains the following chapters:

- Introduction to Database Lifecycle Management
- Setting Up Your Infrastructure



Introduction to Database Lifecycle Management

This chapter covers the following:

- Information Map for Lifecycle Management Solutions
- Overview of the Lifecycle Management Solutions

Information Map for Lifecycle Management Solutions

The Information Map for Lifecycle Management Solutions lists the chapters and sections relevant to the various lifecycle management solutions offered by Enterprise Manager. Consider this an information roadmap to learn about the solution and perform the required operations.

Table 1-1 Information Map for Lifecycle Management Solutions

Domain	Target Type	Solution Area	Reference Links
Database	Oracle Single- Instance Database Oracle Grid Infrastructure Oracle Real Application Server (Oracle RAC)	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
	Oracle RAC One Database	Update (patching)	Database Maintenance Concepts
		Upgrade	Upgrade databases and Grid Infrastructure Using Fleet Maintenance
		Change Management	Manage Database Schema Changes
		Configuration Management	Manage Configuration Information
		Compliance Management	Manage Compliance
		Enterprise Data Governance	Managing Enterprise Data Governance
	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

Managing Enterprise Data Governance

Domain **Target Type Solution Area** Reference Links Configuration Enterprise Provisioning Oracle Database Replay Provisioning and Compliance Manager Client managed targets Change Management Manage Database Schema Changes Configuration **Manage Configuration Information** Management Compliance Manage Compliance Management

Enterprise Data

Governance

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

Overview of the 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 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.

Table 1-2 Lifecycle Management Solutions

Solution Area	Coverage
Discovery	 Automatically discovers software deployments using IP scanning techniques (NMAP).
	 Converts unmanaged software deployments to managed targets, so that their health can be monitored.
	 Offers an integrated workflow for deploying Oracle Management Agents and discovering targets on selected auto-discovered hosts.
	 All discovery solutions are covered in the Oracle Enterprise Manager Host Lifecycle Management Guide under: Discovering Hosts and Software Deployments.



Table 1-2 (Cont.) Lifecycle Management Solutions

Solution Area	Coverage
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.
Patching	 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).
	 Perform migrate, deploy and update patching operations using Fleet Maintenance's user interface or through EM CLI or REST API.
	 Validates patches for applicability in your environment, validates patching operations, and automatically receives patches to resolve conflicts.
	 Fleet Maintenance integrates Oracle Fleet Patching and Provisioning (FPP) solution. Leverage native capabilities of Enterprise Manager automated risk assessment, security patch recommendations, and proactively patch databases and Grid Infrastructure.
	 Offers out-of-place patching (only for standalone databases), in-place patching, and rolling and parallel patching modes, both in offline and online 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.
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 Enterprise Manager 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.



Table 1-2 (Cont.) Lifecycle Management Solutions

Solution Area	Coverage	
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. 	

Note:

The provisioning and patch management solutions are based on Oracle-supplied predesigned deployment procedures, that help you accomplish the provisioning and patching tasks. The three Oracle-supplied deployment procedures used in provisioning and patching are: Deploy an Oracle Home, Migrate Listener, and Switch a Database to run from new Oracle Home.

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.



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.



For patching targets deployed on Microsoft Windows hosts make sure that Cygwin is installed on the target hosts before patching the targets.

For information on how to install Cygwin on a host, see Installing Cygwin and Starting the SSH Daemon in *Oracle Enterprise Manager Cloud Control Basic Installation Guide.*.

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) Deploying Agents on OCI Resources
- (Optional) Configuring Self-Update
- (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	Additional / Value Add setup (optional) Configuring Self-Update	(Optional) Configuring Self-Update



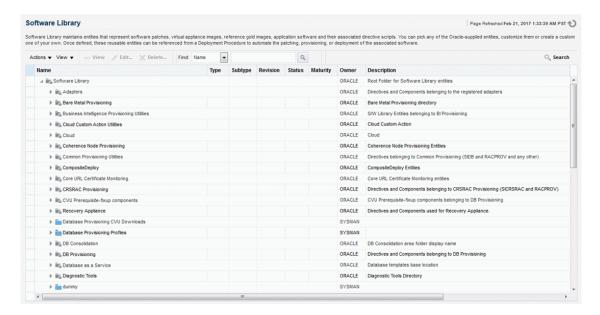
Ensure that the OMS is patched appropriately to the required level. For information about the patches that need to be applied on the Enterprise Manager Cloud Control Management Server (OMS) for using the Provisioning and Patching features, see My Oracle Support note 822485.1.

Setting Up Oracle Software Library

Oracle Software Library is one of the core features offered by Oracle Enterprise Manager, 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 Enterprise Manager, 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



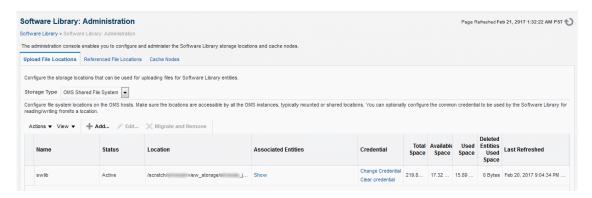
To start using the Software Library to create and manage entities, the Software Library Storage Locations must be configured. System Administrators are responsible for configuring the Software Library storage locations, following which the Software Library becomes usable. Enterprise Manager offers the 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
 - 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
 - NFS Locations
 - 3. Management Agent Locations

You can configure the storage locations from the Administration Console. To do so, in Enterprise Manager, 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



For information on configuring Software Library, see Configuring a Software Library in *Oracle Enterprise Manager Administrator's Guide*.



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, 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 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 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 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 Do the following:

If you do not have direct access or the required credentials for the normal operating system user account (Oracle) OR if you do not have direct access or the required credentials for the privileged account (root).

Set up the Privilege Delegation as follows:

- Create Privilege Delegation (PDP) Template either for SUDO. PowerBroker, or PAM. To do so, see Privilege Needed for Creating a Privilege Delegation
- Apply the created template on the Management Agents of the target hosts.
- Create Named Credentials for normal operating system user account Oracle with privileges to run as SUDO or PowerBroker, for more information see Creating Named Credentials.

Create Named Credentials for privileged users account root with privileges to run as SUDO or PowerBroker, for more information see Creating Privileged Credentials.

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 Oracle Enterprise Manager Security Guide.

If you have direct access or Do the following: the required credentials for the normal operating system user account (Oracle) OR if you have direct access or the required credentials for the privileged account (root).

- Create Named Credentials for normal operating system user account Oracle, for more information see Creating Named Credentials
 - Create Named Credentials for privileged user accounts (root) Credentials, for more information see Creating Privileged Credentials
- 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 Oracle Enterprise Manager Security Guide.

Privileged Access Management

For complete information on setting up PAM in your Enterprise Manager deployment see: (Optional) Setting a PAM User Account.

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 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. There are three main administrator accounts *Designer*, *Operator* and a *Super Administrator* that oversees the previous mentioned roles.

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.
- Publishing gold images for updates and upgrades
- Defining pre-post scripts that need to be invoked as part of the provisioning and patching flows. For advanced users, customizing the default provisioning procedures according to the needs of the organization within the limitations laid down for those deployment procedures.

The predefined Oracle role for a Designer is <code>EM_ALL_DESIGNER</code>, this role in turn includes fine grained roles where you can specifically set <code>EM_PROVISIONING_DESIGNER</code> for provisioning tasks, and <code>EM_PATCH_DESIGNER</code> 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.
- Deploying the gold images and updating/upgrading the targets by switching to the newly deployed Oracle homes.

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.



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:

- 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.



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.
- 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.



Enterprise Manager accesses, My Oracle Support 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 Enterprise Manager 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 Enterprise Manager 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:

- In Enterprise Manager, 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 Administrator's Guide*.
- From the Actions menu, select Subscribe to ensure that you receive notification whenever a provisioning bundle is available for download.
- In the Updates Home page, select update of Type Provisioning Bundle and from the Actions menu, select Open.
- 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 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.

Setting restricted accesses for root components can be useful for administrators that want to review and pre-stage to a read only location to avoid potentially malicious changes, or for administrators that don't want to provide root credentials to uses.

Primarily, this section covers:

- Security Models for running Root scripts
- Root Components for Provisioning

Security Models for running Root scripts

Root Components for Fleet Maintenance (Patching)

- Manually Staging the Root Components For Fleet Maintenance
- Restricting the Root User Access For Fleet Maintenance

Manually Staging the Root Components For Fleet Maintenance (Patching)

By default, the root component is automatically staged on the target host (at the location defined by <code>%emd_emstagedir%</code>). Before you manually stage a root component you need to consider the following critical points:

 After picking up a Release Updates (RU) or Plug-in release, you need to stage the root scripts component into a temporary location and do the same with the pre-staged files.



- Review the changes between versions and update the pre-stage location with the new set of files.
- Note that the RU or Plug-in will only work with the latest files shipped with that update.
 Validate and, either accept the changes or defer the uptake of the RU or Plug-in until the concern is resolved as per your internal process.
- Once the RU or Plug-in update is accepted, incorporate any custom changes made in the previous version to the new root scripts and pre-staged files.
- Pre-stage the root scripts for production use and continue with the regular manually staging steps.

If you want to manually stage (pre-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.
- 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.
- Select Root Scripts, and then click View.
- 7. Click Stage.
- 8. 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.



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.

- 9. Unzip the file staged in the previous step.
- 10. Select Common Root, and then click View.
- 11. Click Stage.

You now have configured the options for a Pre-Staged Root Dispatcher for fleet maintenance, for more information please see: Fleet Management Software.

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 For Fleet Maintenance (Patching)

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 <code>/etc/sudoers</code> policy file. Since provisioning involves invoking all the commands on the target host using Management Agent, you must include <code>nmosudo</code> in the <code>sudoers</code> file with the set of commands that is possible to run. To provide restricted **root** access to the Management Agent user, edit the <code>/etc/sudoers</code> file such that it has the following contents:

<command_alias> represents the alias that describes the entities that can be run by mngmtuser
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%.

You can also restrict root access by setting the **dispatchLoc** flag within the **input_file**. This parameter has been added, to aide operations that stage and use a dispatcher location, using an existing, read-only location to use pre-staged dispatcher as well as fleet scripts.

Fleet Maintenance operations allow you to create a read-only mount point location for root-owned types of scripts. You can restrict access to **root** only by setting the **dispatchLoc** flag within the **input_file**. You may choose, optionally, to pre-stage all the necessary files in this location. If you pre-stage, you must include in the **dispatchLoc**: the Root Dispatcher, unzipped Root Scripts and common root.

For more information on the input_file please review: Fleet Management Software

Root Components for Provisioning

This section covers the following:

- Manually Staging the Root Components for Provisioning
- Restricting the Root User Access



Manually Staging the Root Components for Provisioning

By default, the root component is automatically staged on the target host (at the location defined by <code>%emd_emstagedir%</code>). Before you manually stage a root component you need to consider the following critical points:

- After picking up a Release Updates (RU) or Plug-in release, you need to stage the root scripts component into a temporary location and do the same with the pre-staged files.
- Review the changes between versions and update the pre-stage location with the new set of files.
- Note that the RU or Plug-in will only work with the latest files shipped with that update.
 Validate and, either accept the changes or defer the uptake of the RU or Plug-in until the concern is resolved as per your internal process.
- Once the RU or Plug-in update is accepted, incorporate any custom changes made in the previous version to the new root scripts and pre-staged files.
- Pre-stage the root scripts for production use and continue with the regular manually staging steps.

If you want to manually stage (pre-stage) the root component at a custom location before initiating the patching process, follow these steps:

- From the Enterprise menu, select Provisioning and Patching, then select Software Library.
- 2. Navigate to Patching And Provisioning, and then Components.
- 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.



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 <code>/etc/sudoers</code> 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 <code>/etc/sudoers</code> file such that it has the following contents:

```
Cmnd_Alias SR_PATCH_DISPATCHER = \
/u01/app/oracle/agent/agent_13.5.0.0.0/sbin/nmosudo *, \
/u01/app/oracle/agent/agent_13.5.0.0.0/sbin/nmosudo DEFAULT_PLUGIN
DEFAULT_FUNCTIONALITY DEFAULT_SUBACTION DEFAULT_ACTION perl *, \
/u01/app/oracle/agent/agent_13.5.0.0.0/sbin/nmosudo DEFAULT_PLUGIN
DEFAULT_FUNCTIONALITY DEFAULT_SUBACTION DEFAULT_ACTION id, \
/u01/app/oracle/agent/agent_13.5.0.0.0/sbin/nmosudo DEFAULT_PLUGIN
DEFAULT_FUNCTIONALITY DEFAULT_SUBACTION DEFAULT_ACTION
/tmp/fleet/dispatchLoc/root_dispatcher.sh *
oracle ALL=(root)NOPASSWD:SR_PATCH_DISPATCHER
grid ALL=(root)NOPASSWD:SR_PATCH_DISPATCHER
```

<command_alias> represents the alias that describes the entities that can be run by oracle or
grid 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%.



(Optional) Deploying Agents on OCI Resources

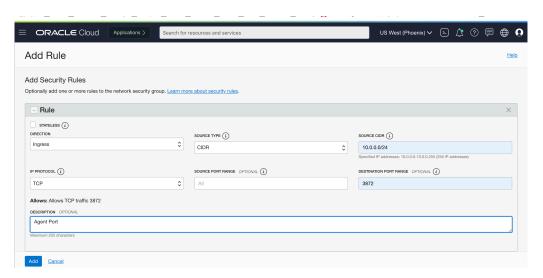
With Enterprise Manager you can manage database assets running in a combination of onpremises data centers or in the Oracle Cloud Infrastructure (OCI) from the same user interface.

Prerequisites To Deploying Agents on OCI resources

Before you begin please ensure the following prerequisites that need to be performed within OCI prior to discovery in Enterprise Manager

- Provision a Compute Virtual Machine or Database System and then deploy an agent on the same network VCN and Subnet for target discovery and monitoring. The agent needs to use the same SSH keys that you have used for provisioning Compute Virtual Machines or Database system when deploying an agent.
 - a. Open port 3872 on the VCN security list to permit connection.





- 2. Log in to the node where the agent will be deployed and open the firewall port 3872.
 - a. Use the following firewall rules for database systems:

```
sudo iptables -I INPUT -p tcp -m state --state NEW -m tcp -s 10.0.1.0/25 --dport 3872 -m comment --comment "Required foraccess to Agent Listener" -j ACCEPT sudo iptables -I INPUT -p tcp -m tcp --dport 3872 -j ACCEPT sudo service iptables save sudo service iptables reload
```

b. Use the following commands for RHEL 7.7 and above:

```
firewall-cmd --zone=public --add-port=3872/tcp --permanent
firewall-cmd --reload
iptables-save | grep 3872
```

Create named credentials on Enterprise Manager with the SSH key provided during Compute VM or Database system provisioning, to do this in Enterprise Manager navigate to Setup>Security and click on Named Credentials. a. Credential Name: OCI SSH Key

b. Credential Type: SSH Key Credentials

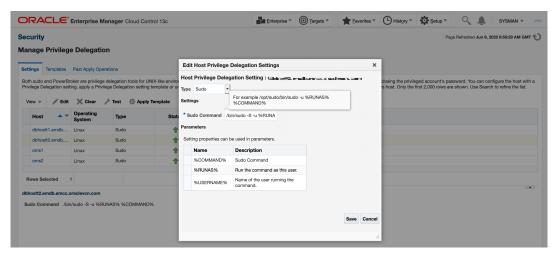
c. Scope: Globald. Username: opce. Run Privilege: sudo

Run as: oracle

f.

4. Set up Privilege Delegation. Navigate to **Setup>Security>Privilege Delegation**. Select the installed agent and make sure **Privilege Delegation** is set for the agent.

Figure 2-5 Privilege Delegation Setting



5. For single node deployments add the agent hosts IP and host name to the Enterprise Manager hosts file and add Enterprise Manager details on the Agent machines host file:



This step is not required for Multi Node environments as all communication goes through the load balancers.

a. Example of setup needed for an OMS host in a Single Node environment:

```
[opc@oms1 ~]$ cat/etc/hosts
10.0.1.3 <oms node 1>
10.0.1.4 <oms node 2>
10.0.1.6 <db host 1>
10.0.1.7 <db host 2>
```

b. Example of setup needed for an agent deployed on a Compute VM in a Single Node environment:

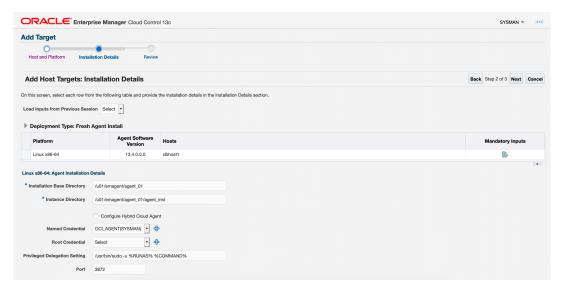
```
[opc@dbhostt1 ~]$ cat/etc/hosts
10.0.1.6 dbhost1.sample.com dbhost1
192.168.16.18 dbhost1-priv.sample.com dbhost1-priv
10.0.1.8 dbhost1-vip.sample.com dbhost1-vip
10.0.1.7 dbhost2.sample.com dbhost2
```

```
192.168.16.19 dbhost2-priv.sample.com dbhost2-priv
10.0.1.9 dbhost2-vip.sample.com dbhost2-vip
10.0.1.3 omsnode1
10.0.1.4 omsnode2
```

To deploy Enterprise Management agents on OCI resources follow these steps:

Add the agent using the graphical interface by navigating to: Setup>Add Target>Add
 Targets Manually>Add Hosts Targets. Click on Install Agent on Host and enter
 Installation directory, Instance directory and credentials.

Figure 2-6 Add Host Target Installation Details



- After the agent deployment, make sure the agent is shown on the Enterprise Manager All Targets page.
- 3. Perform a target Discovery after these steps are done, this will integrate an OCI resource into Enterprise Manager. For more information see: Discovering and Adding Database Targets in Oracle Enterprise Manager Cloud Control Administrator's Guide.



Remote agents are supported for Lifecycle Management activities, for more information on remote agents see: Installing Oracle Remote Management Agents.

Once the OCI database target has been discovered you can manage specific database actions. For more information, see Managing Pluggable Databases Using Enterprise Manager .

The following management actions can be performed on an OCI PDB:

- Creating a New Pluggable Database Using Enterprise Manager
- Deleting Pluggable Databases Using Enterprise Manager
- Cloning a Pluggable Database Using Enterprise Manager
- Relocating Pluggable Databases Using Enterprise Manager



- Refresh a Cloned Database
- Creating a Full Clone Pluggable Database

(Optional) Setting a PAM User Account

Oracle Enterprise Manager is compatible with PAM (Privileged Access Management), allowing you to eliminate the need for hard-coded application credentials embedded in applications, scripts, or configuration files. Allowing highly sensitive passwords to be centrally stored, logged, and managed within the Vault.

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 are SUDO, PowerBroker, and Privileged Access Manager (PAM). This support is offered using the Privilege Delegation mechanism available in Enterprise Manager.

To Configure PAM in your environment see the following resources:

- For a conceptual overview of Privilege Delegation and the authentication tools supported by it, see Oracle Enterprise Manager Enterprise Manager Security Guide.
- For complete information on setting up PAM in your Enterprise Manager deployment see: Privileged Access Management Integration with Enterprise Manager.
- For a in depth look at configuring PAM and sample scripts see Oracle MOS Note: 2948386.1.



Part II

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
- 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
- Fast-Start Failover REST API



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, supported targets and deployment procedures, 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 Enterprise Manager. As part of the database provisioning solution, Enterprise Manager 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 Enterprise Manager.

Provisioning Solution Provisioning Provisioning Provisioning Provisioning Provisioning Provisioning Oracle RAC Creating Upgrading Single Instance Oracle RAC Grid Pluggable SIHA Databases One Databases Databases Infrastructure Databases Databases Extending Deleting RAC RAC

Figure 3-1 Database Provisioning Solution in Enterprise Manager

For this release, database provisioning features are as follows:

Designer and Operator Roles

Enterprise Manager 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, preset 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 Enterprise Manager Console to be used in your provisioning activities.

Creating Databases Using Enterprise Manager

Enterprise Manager now enables you to create databases from the Enterprise Manager 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 Enterprise Manager.

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 3-1 lists the database deployment procedures offered by Cloud Control and the various targets that can be provisioned.



Table 3-1 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	Pluggable Databases (available only as a part of Oracle Database 12c Release 1 and above)



This deployment procedure cannot be manually invoked from the Database Provisionin g page. This deployment procedure is invoked internally when you use the Provision Pluggable Database console or EM CLI to provision pluggable databases.

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

Table 3-1 (Cont.) Database Deployment Procedures and Targets Provisioned

Deployment Procedure	Targets Provisioned
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 3-2 lists various use cases for database provisioning deployment procedures.

Table 3-2 use cases for Database Deployment Procedures

Deployment Procedure	use case	Link
Provision Oracle Database	Provisioning and Creating Single- Instance Databases	Provisioning and Creating Oracle Databases
	Provisioning Single-Instance Database with Oracle Automatic Storage Management Provisioning Single Instance Oracle Instance Or	 Provisioning Oracle Databases with Oracle Automatic Storage Management Provisioning Oracle Database Software Only
	 Provisioning Single-Instance Database Software Only Provisioning Oracle Grid Infrastructure with Single-Instance 	Provisioning Oracle Grid Infrastructure and Oracle Databases with Oracle Automatic Storage Management
	Database and Configuring Database with Oracle Automatic Storage Management	Provisioning Oracle Grid Infrastructure and Oracle Database Software Only
	 Provisioning Oracle Grid Infrastructure and Single-Instance 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 Grid Infrastructure with Oracle Real Application Clusters Database and Configuring Database with Oracle Automatic Storage Management Provisioning Oracle Real Application
	 Provisioning Oracle Real Application Clusters Database with File System 	Clusters Database with File System on an Existing Cluster
	 on an Existing Cluster Provisioning Oracle Real Application Clusters Database with File System on a New 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 an Oracle Database Creating Oracle Real Application Clusters Database
	 Creating Oracle Real Application Clusters One database 	Creating Oracle Real Application Clusters One Node Database



Table 3-2 (Cont.) use cases for Database Deployment Procedures

Deployment Procedure	use case	Link
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 Provision Oracle Clusterware / Oracle RAC for UNIX and RDBMS versions 10g/11g	 Cloning a Running Oracle Real Application Clusters Provisioning Oracle Real Application Clusters Using Gold Image Provisioning Oracle Real Application Clusters Using Archived Software Binaries 	 Cloning a Running Oracle Real Application Clusters Provisioning Oracle Real Application Clusters Using Gold Image Provisioning Oracle Real Application Clusters Using Archived Software Binaries
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 	Client Using Gold Image

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

Complete these requirements by doing 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 Installing Cygwin and Starting the SSH Daemon in *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.

Understanding Administrator Privileges for Provisioning Database

Table 3-3 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 3-3 Privileges for Using Deployment Procedures

Role	Target Privileges	Resource Privileges	Implementation Recommendation
EM_PROVISIONING_ DESIGNER	Instance Privileges: • Target Name: Database Target Note:	 Resource Type: Software Library Entity Privilege: Manage Any Software Library Entity Instance Privileges: View Credential privilege is required for the named credentials. 	 Used for profile creation of the Oracle Home, Grid Home and Database Can customize the deployment
	This privilege is required only if the target is	Note: This privilege is required only if the target is owned by a different user.	procedure and also saved with the input values • Required when you want to grant and restrict access to deployment
	owned by a different user. Privilege: Connect		 Can create or manage software library component used as an pre and post scripts
EM_PROVISIONING_	Target (Required for database profile creation.) N/A	Instance Privileges: View Credential privilege is	Provision Grid
OPERATOR		required for the named credentials.	Infrastructure, Real Application Clusters (RAC),
		Note:	and Single Instance Oracle Home.
		This privilege is required only if the target is owned by a different user.	 Enables the creation of SI and RAC databases, and also allows
			management of pluggable databases. • Launch privileges are required to initiate the Deployment Procedure.

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.



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:

Note:

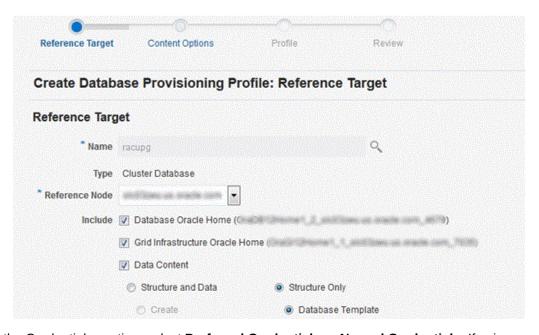
A Database provisioning profile created from a CDB reference target cannot be used to provision non CDB Databases. CDB reference targets can only be used to provision CDB Databases.



- 1. From the Targets menu, select Databases, then select Administration, then select Provisioning and Patching, and click on 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, 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.



- 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



 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**.



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 FMCLI

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:

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

```
emcli get instances
```

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
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=clustername
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=OmsShared
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:

- Create a temporary location mkdir /tmp/installmedia.
- Navigate to the following URL:

```
http://www.oracle.com/technetwork/database/enterprise-edition/downloads/
index.html
```

- 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 the 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**.
- In the Create Installation Media: Describe page, enter the Name and Description for the component, and click Next.
- 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.
 - 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.
- 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:

- 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.



For **Structure as well as data**, a PDB cannot be renamed. This template is restoring the physical backup of the source database. The destination database will be the same as 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 **Targets** menu, select **Databases**, then select **Administration**, then select **Provisioning and Patching**, and then select **Software Library**.
- 2. On the Software Library Home page, select the folder where you want to upload the database template.
- From the Actions menu, select Create Entity, then select Component. Alternately, right click the custom folder, and from the menu, select Create Entity, then select Component.
- From the Create Entity: Component dialog box, select Database Template and click Continue.
 - Cloud Control displays the Create Database Template 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.

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

Pre/Post Scripts for Database Provisioning

Starting with Oracle Enterprise Manager 13.5 Release Update 1 you can now add Pre and Post scripts for specific actions when creating or provisioning different types of Databases.

You can run custom scripts before and after a database has been created or provisioned. These requests can range from some additional host commands on the machine on which the database was created or commands to perform additional book activities and other operations on the newly provisioned database. Scripts can be Shell scripts or Perl scripts (SQL scripts cannot be called directly but can be called internally from Shell or Perl scripts).

Pre/Post scripts need to be uploaded as EM Software Library Entity (Directive) prior to their usage with Database Provisioning operations. With the pre/post scripts in Software Library, the whole process is centralized and automated with the script staged at a common place within the Software Library, the script is staged and cleaned up on the target host as part of the deployment execution itself.

Before using Pre/Post scripts in Database Provisioning, you will first need to create the script, upload the script to the Software Library and obtain the URN for your script. The following resources show step by step how to perform these activities:

- Creating Pre/Post scripts: Pre and Post Request Creation / Deletion Scripts
- Upload Pre/Post scripts: Software Library Entity For Uploading The Script
- Obtain the URN for your script: How to Get URN for SWLIB Entity

The following is a complete list of Pre and Post scripts you can create depending on the type of database to be created or provisioned.

Table 3-4 Provisioning SIHA Databases

Selected Option	Script
Deploy Grid Infrastructure software only	Pre Database Home deployment script Post Database Home deployment script
	Root Pre Database Home deployment script
	Root Post Database Home deployment script
Deploy Database software only	Pre Database Home deployment script Post Database Home deployment script
	Root Pre Database Home deployment script
	Root Post Database Home deployment script



Table 3-4 (Cont.) Provisioning SIHA Databases

Selected Option	Script
Deploy Grid Infrastructure software and Configure	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script
	Root Pre Grid Infrastructure Home configuration script
	Root Post Grid Infrastructure Home configuration script
Deploy Grid Infrastructure software and Configure + Deploy Database software and create a database	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script
	Root Pre Grid Infrastructure Home configuration script
	Root Post Grid Infrastructure Home configuration script
	Pre Database Home deployment script
	Post Database Home deployment script
	Root Pre Database Home deployment script
	Root Post Database Home deployment script
Deploy Grid Infrastructure software and Configure + Deploy Database software only	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script
	Root Pre Grid Infrastructure Home configuration script
	Root Post Grid Infrastructure Home configuration script
	Pre Database Home deployment script
	Post Database Home deployment script
	Root Pre Database Home deployment script
	Root Post Database Home deployment script
Deploy Grid Infrastructure software only	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script
Deploy Grid Infrastructure software only + Deploy Database software and create a database	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script
	Pre Database Home deployment script
	Post Database Home deployment script
	Root Pre Database Home deployment script
	Root Post Database Home deployment script
Deploy Grid Infrastructure software only + Deploy Database software only	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script
-	Pre Database Home deployment script
	Post Database Home deployment script
	Root Pre Database Home deployment script
	Root Post Database Home deployment script

Table 3-5 Provisioning RAC Databases

Selected Option	Script
Deploy Grid Infrastructure software only	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script

Table 3-5 (Cont.) Provisioning RAC Databases

Selected Option	Script	
Deploy Grid Infrastructure software and Configure	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script	
	Root Pre Grid Infrastructure Home configuration script	
	Root Post Grid Infrastructure Home configuration script	
Deploy Database software only	Pre Database Home deployment script Post Database Home deployment script	
	Root Pre Database Home deployment script	
	Root Post Database Home deployment script	
Deploy Database software and create a RAC database	Pre Database Home deployment script Post Database Home deployment script	
	Root Pre Database Home deployment script	
	Root Post Database Home deployment script	
Deploy Database software and create a RAC One Node database	Pre Database Home deployment script Post Database Home deployment script	
	Root Pre Database Home deployment script	
	Root Post Database Home deployment script	
Deploy Grid Infrastructure software only + Deploy Database software only	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script	
	Pre Database Home deployment script	
	Post Database Home deployment script	
	Root Pre Database Home deployment script	
	Root Post Database Home deployment script	
Deploy Grid Infrastructure software only + Deploy Database software and create a RAC database	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script	
	Pre Database Home deployment script	
	Post Database Home deployment script	
	Root Pre Database Home deployment script	
	Root Post Database Home deployment script	
Deploy Grid Infrastructure software only + Deploy Database software and create a RAC One Node	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script	
database	Pre Database Home deployment script	
	Post Database Home deployment script	
	Root Pre Database Home deployment script	
	Root Post Database Home deployment script	
Deploy Grid Infrastructure software and Configure + Deploy Database software only	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script	
	Root Pre Grid Infrastructure Home configuration script	
	Root Post Grid Infrastructure Home configuration script	
	Pre Database Home deployment script	
	Post Database Home deployment script	
	Root Pre Database Home deployment script	
	Root Post Database Home deployment script	



Table 3-5 (Cont.) Provisioning RAC Databases

Selected Option	Script	
Deploy Grid Infrastructure software and Configure + Deploy Database software and create a RAC	Pre Grid Infrastructure Home deployment script Post Grid Infrastructure Home deployment script	
database	Root Pre Grid Infrastructure Home configuration script	
	Root Post Grid Infrastructure Home configuration script	
	Pre Database Home deployment script	
	Post Database Home deployment script	
	Root Pre Database Home deployment script	
	Root Post Database Home deployment script	
Deploy Grid Infrastructure software and Configure	Pre Grid Infrastructure Home deployment script	
+ Deploy Database software and create a RAC	Post Grid Infrastructure Home deployment script	
One Node database	Root Pre Grid Infrastructure Home configuration script	
	Root Post Grid Infrastructure Home configuration script	
	Pre Database Home deployment script	
	Post Database Home deployment script	
	Root Pre Database Home deployment script	
	Root Post Database Home deployment script	

Table 3-6 Customizing Pre/Post Script Anchor Phases

Procedure Name	Anchor Phase	Description	Position in Script
Create Database	Begin: Custom pre- scripts (Advanced)	Insert your custom prescripts in appropriate phases and steps after this step. Note: This is advanced usage of customization. Alternately, you may use the simplified hooks provided for your prescripts to avoid procedure customization and maintenance.	After- Initialize Deployment Procedure



Table 3-6 (Cont.) Customizing Pre/Post Script Anchor Phases

Procedure Name	Anchor Phase	Description	Position in Script
	End: Custom pre-scripts (Advanced)	Insert your custom pre- scripts in appropriate phases and steps above this step.	Before- Database creation Post-Script





Table 3-6 (Cont.) Customizing Pre/Post Script Anchor Phases

Procedure Name	Anchor Phase	Description	Position in Script
	Begin: Custom post- scripts (Advanced)	Insert your custom prescripts in appropriate phases and steps after this step. Note: This is advanced usage of customization. Alternately, you may use the simplified hooks provided for your prescripts to avoid procedure customization and maintenance.	Before- Schedule Target Collections



Table 3-6 (Cont.) Customizing Pre/Post Script Anchor Phases

Procedure Name	Anchor Phase	Description	Position in Script
	End: Custom post- scripts (Advanced)	Insert your custom pre- scripts in appropriate phases and steps above this step.	Before- Schedule Target Collections

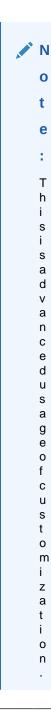




Table 3-6 (Cont.) Customizing Pre/Post Script Anchor Phases

Procedure Name	Anchor Phase	Description	Position in Script
Provision Oracle Database	Begin: Custom pre- scripts (Advanced)	Insert your custom pre- scripts in appropriate phases and steps after this step. Note: This is advanced usage of customization. Alternately, you may use the simplified hooks provided for your pre- scripts to avoid procedure customization and maintenance.	After- Initialize Deployment Procedure



Table 3-6 (Cont.) Customizing Pre/Post Script Anchor Phases

Procedure Name	Anchor Phase	Description	Position in Script
	End: Custom pre-scripts (Advanced)	Insert your custom pre- scripts in appropriate phases and steps above this step.	Before- Initialize Prerequisites and Fixups





Table 3-6 (Cont.) Customizing Pre/Post Script Anchor Phases

Procedure Name	Anchor Phase	Description	Position in Script
	Begin: Custom post- scripts (Advanced)	Insert your custom prescripts in appropriate phases and steps after this step. Note: This is advanced usage of customization. Alternately, you may use the simplified hooks provided for your prescripts to avoid procedure customization and maintenance.	Before Step- Cleanup



Table 3-6 (Cont.) Customizing Pre/Post Script Anchor Phases

Procedure Name	Anchor Phase	Description	Position in Script
	End: Custom post- scripts (Advanced)	Insert your custom pre- scripts in appropriate phases and steps above this step.	Before Step- Cleanup

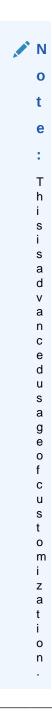




Table 3-6 (Cont.) Customizing Pre/Post Script Anchor Phases

Procedure Name	Anchor Phase	Description	Position in Script
Provision Oracle RAC Database	Begin: Custom prescripts (Advanced)	Insert your custom prescripts in appropriate phases and steps after this step. Note: This is advanced usage of customization. Alternately, you may use the simplified hooks provided for your prescripts to avoid procedure customization and maintenance.	After- Initialize Deployment Procedure



Table 3-6 (Cont.) Customizing Pre/Post Script Anchor Phases

Procedure Name	Anchor Phase	Description	Position in Script
	End: Custom pre-scripts (Advanced)	Insert your custom pre- scripts in appropriate phases and steps above this step.	Before- Initialize Prerequisites and Fixups





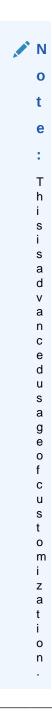
Table 3-6 (Cont.) Customizing Pre/Post Script Anchor Phases

Procedure Name	Anchor Phase	Description	Position in Script
	Begin: Custom post- scripts (Advanced)	Insert your custom prescripts in appropriate phases and steps after this step. Note: This is advanced usage of customization. Alternately, you may use the simplified hooks provided for your prescripts to avoid procedure customization and maintenance.	Before Step- Cleanup



Table 3-6 (Cont.) Customizing Pre/Post Script Anchor Phases

Procedure Name	Anchor Phase	Description	Position in Script
	End: Custom post- scripts (Advanced)	Insert your custom pre- scripts in appropriate phases and steps above this step.	Before Step- Cleanup





Pre/Post Script Examples for Database Provisioning

The following is a set sample of a Perl scripts for Pre/Post Database Provisioning operations.



Do not edit Pre/Post hooks sections of Pre/Post script.

```
#!/usr/local/bin/perl
        #Store the passed input file path in a variable
        my $fileName = $ARGV[0];
        #Print the input file content as it is...
        print "Printing the variable values received from cloud
request ...\n";
        open IN, "<", $fileName or die "Can't open file due to :$!";
        while(<>){
         print $;
        close IN;
        # Reads a text file with var=values pairs line delimited
        # Setting the values in the environment
        sub readDat {
                # Grab the passed parameter
                my filename = f[0];
                my @Data;
                # The ubiquitous $i counter
                my $i = 0;
                # Register a filehandle
                local (*DATFILE);
                # Open the file for read
                open (DATFILE, $filename) or print ( "Can't
open $filename: $1");
                # Read through the file one line at a time
                FORA:while (<DATFILE>) {
                        # Skip over any comments
                        if ( /#.*/ ) {
next FORA;
                        # Clean up any extraneous garbage
                        chomp;
                                               # no newline
                        s/^\s+//;
                                                # no leading white
                        s/\s+$//;
                                               # no trailing white
                    # If clean up eliminated any data worth reading
                    # lets skip to the next line
                    next unless length;
                        # We can't load the lines with $VARS as we will
                        # loose any values so we'll make sure to escape them
                        # (the $'s that is)
                        s/\$/\\\$/g;
                        # localizing $var and $value to make sure
                        # they are clean out on every read.
                        my (\$var, \$value) = split(/=/,\$);
```

```
# Load variable into enviroment
$ENV{$var}="$value";
print $value;
$i++;
}
print "Loading cloud variables into ENV..\n";
readDat($fileName);
```

Pre/Post Script With Additional Oracle Home, Oracle Base and SID Fields

Starting with Oracle Enterprise Manager 13.5 Release Update 21 (13.5.0.1), Oracle Home, Oracle Base and SID fields are available for additional customization in both pre and post scripts. The following script is an example:



You can provide a SQL script as a file or Software Library entity to be executed within a Post database creation PERL script as part of the script customization options. When specifying Oracle Home and SID fields in your custom scripts, keep in mind that Pre and Post scripts will be executed on all nodes, ensure that running the same script across the nodes multiple times will not cause an impact.

```
#!/usr/local/bin/perl
#Store the passed input file path in a variable
my $fileName = $ARGV[0];
#Print the input file content as it is...
print "Printing the variable values for OH post-scripts received from cloud
request ...\n";
open IN, "<", $fileName or die "Can't open file due to :$!";
       while (<>) {
          print $_;
        close IN;
        # Reads a text file with var=values pairs line delimited
        # Setting the values in the environment
        sub readDat {
                # Grab the passed parameter
                my filename = f[0];
                my $varName = $ [1];
                my @Data;
                # The ubiquitous $i counter
                my $i = 0;
                # Register a filehandle
                local (*DATFILE);
                # Open the file for read
                open (DATFILE, $filename) or print ( "Can't
open $filename: $1");
                # Read through the file one line at a time
                FORA:while (<DATFILE>) {
                       # Skip over any comments
                       if ( /#.*/ ) {
                       next FORA;
```

```
# Clean up any extraneous garbage
                       chomp;
                                         # no newline
                       s/^\s+//;
                                             # no leading white
                       s/\s+$//;
                                              # no trailing white
                    # If clean up eliminated any data worth reading
                    # lets skip to the next line
                    next unless length;
                        # We can't load the lines with $VARS as we will
                        # loose any values so we'll make sure to escape them
                        # (the $'s that is)
                        s/\$/\\\$/g;
                        # localizing $var and $value to make sure
                        # they are clean out on every read.
                        my (\$var, \$value) = split(/=/,\$);
                        if( defined($varName) and $var eq $varName ) {
                               return $value;
                        # Load variable into enviroment
                        #$ENV{$var}="$value";
                        #print $value;
                        $i++;
       print "\n";
       print "Loading cloud variables into ENV...\n";
       readDat($fileName);
       my $ORACLE_HOME= readDat($fileName,"ORACLE HOME LOC");
       my $DB SID = readDat($fileName, "DB SID");
       if( length $DB SID ) {
                $ENV{"ORACLE HOME"}=$ORACLE HOME;
                $ENV{"ORACLE SID"}=$DB SID;
                print "\nOracle home value is " . $ENV{"ORACLE HOME"};
                print "\nOracle SID is " . $ENV{"ORACLE SID"};
                print "Executing attached SQL ...\n";
                my $CMD = $ENV{"ORACLE HOME"} . "/bin/sqlplus / as sysdba
\@sample.sql";
                print "Firing SQL ...\n";
                print "$CMD\n";
                system($CMD);
                my $ERROR CODE = $?;
                print "Error code is $ERROR CODE \n";
                if ( $ERROR_CODE == 0 )
                        print "Script Completed\n";
                }
                else
                {
                        print "Error occured while executing \n";
                        return -1;
```

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:

- From the Enterprise menu in EM select Provisioning and Patching and then select Software Library.
- Click the Referenced File Locations tab and select NFS as the Storage type.
- 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
 - 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:

- 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:

 From the Enterprise menu in EM, select Provisioning and Patching, and then select Software Library.

- 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.
- 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.
 - 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.
- From the Actions menu, select Create Entity, then select Component. Alternately, 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.
- 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**. Alternately, right click the custom folder, and from the menu, select **Create Entity**, then select **Component**.
- From the Create Entity: Component dialog box, select Oracle Clusterware Clone and click Continue.
 - Cloud Control displays the Create Oracle Clusterware Clone: Describe page.
- 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.
- 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**. Alternately, right click the custom folder, and from the menu, select **Create Entity**, then select **Component**.
- 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.
- 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 if you need to get the latest version available, download from 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:

Note:

This is not an automated job, the update must be manually performed.

- 1. From the Targets menu, select Databases, then select Administration, then select Provisioning and Patching, and 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.
- Click OK. This will update the Software Library with the latest cluster verification utility binaries.



4

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 4-1 shows a typical Oracle Database (single-instance database) topology that you can provision using Cloud Control.



ASM Disk RAID Diskgroup named DATA consumed by Database for Datafiles Public Storage Network Network NIC **Database Host** Datafiles in ASM diskgroup named DATA SIHA Managed

Figure 4-1 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*

Note:

If you are going to provision an Oracle 19.6 database on Oracle Linux 8, you need to apply the following mandatory patches before you start the provisioning process:

- 19.6 DBRU Patch# 30557433
- 19.6 OJVM Patch# 30484981
- 19.6 OCWRU Patch# 30855259
- 19.6 Oneoff Patch# 30150710

If these patches are not applied, provisioning will fail when the Gold Image is created from the unpatched source. Sample code:

```
<19.3 on-prem_dbhome image>/runInstaller -applyRU <19.6 DBRU Patch# 30557433 Location> -applyOneOffs <19.6 OJVM Patch# 30484981 Location>, <19.6 OCWRU Patch# 30855259 Location>, <19.6 Oneoff Patch# 30150710 Location EMExpress>
```

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 4-1 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.
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).



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 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 Enterprise Manager, ensure that you meet the following prerequisites:

- Ensure that you meet the infrastructure requirements explained in Setting Up Your Infrastructure.
- 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 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 Pre or Post scripts for Database Creation, make sure they are created and loaded into the software library before starting the creation process. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.
- 8. If you are using Automatic Storage Management (ASM) as storage, ASM instances and disk groups must be configured prior to creating database.
- The Enterprise Manager user creating the database template must have CONNECT ANY TARGET privilege in Enterprise Manager.

Procedure for Creating an Oracle Database

To create an Oracle database, follow these steps:

- Access from Databases, then select Administration drop down menu and click 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.

 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.



Starting with Oracle Enterprise Manager 13.5 RU1 Oracle Database 21c is available as a Container Database only option.

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. Starting with Enterprise Manager 13c Release 5 Update 8 (13.5.0.8) you can specify a non-DBSNMP user for database monitoring, edit the DBSNMP field with the user name. If the user does not exist it will be created.

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 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.

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 Custom Properties, 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 SQL 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.

10. 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.

Select whether you wish to pause after the Prerequisite check to analyze the results and whether Prerequisite System Checks are to be performed.

Select the Pre/Post scripts you will use, these need to be previously uploaded into the Software Library. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.

Click Next.

11. 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.



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.

Creating Oracle Real Application Clusters Database

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



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.
- 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*.
- 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 Pre or Post scripts for Database Creation, make sure they are created and loaded into the software library before starting the creation process. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.



- 8. 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.
- 9. If you are using Automatic Storage Management (ASM) as storage, ASM instances and diskgroups must be configured prior to creating database.
- 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:

- Access from Databases, then select Administration drop down menu and click 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. Starting with Enterprise Manager 13c Release 5 Update 8 (13.5.0.8) you can specify a non-DBSNMP user for database monitoring, edit the DBSNMP field with the user name. If the user does not exist it will be created.

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.

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 Pre/Post and SQL scripts from the Software Library. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning. 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:

- Ensure that you meet the infrastructure requirements explained in Setting Up Your Infrastructure.
- 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 using Pre or Post scripts for Database Creation, make sure they are created and loaded into the software library before starting the creation process. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.
- 8. 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.
- 9. If you are using Automatic Storage Management (ASM) as storage, ASM instances and diskgroups must be configured prior to creating database.



10. 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:

- From Databases, then select Administration drop down menu and click 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.
- 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. Starting with Enterprise Manager 13c Release 5 Update 8 (13.5.0.8) you can specify a



non-DBSNMP user for database monitoring, edit the DBSNMP field with the user name. If the user does not exist it will be created.

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 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 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 Pre/Post and SQL scripts from the Software Library. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning. 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. 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 documentation 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 4-2 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 with Oracle Automatic Storage Management. 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:

- Access from Databases, then select Administration drop down menu and click 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.
- 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**.

- 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.
- In the Specify OS Users page, specify the operating system user for the Oracle Home for the database.



To use credentials other than root, 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)

If they do not exist, then either specify alternative groups that exist on the host or create new groups as described in: Oracle Database Install and Upgrade

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.

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.
- 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.

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 <code>ORACLE HOME/assistants/dbca/templates</code>.

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. Starting with Enterprise Manager 13c Release 5 Update 8 (13.5.0.8) you can specify a non-DBSNMP user for database monitoring, edit the DBSNMP field with the user name. If the user does not exist it will be created.

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 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.

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 Custom Properties, 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 SQL scripts. Select custom SQL 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.

Note:

When selecting a listener on a Single Instance Database provisioning make sure you select a listener that shares a location with the destination Oracle Home.

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. 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.

Select whether you wish to pause after the Prerequisite check to analyze the results and whether Prerequisite System Checks are to be performed.

If you are running a cluster provisioning, select the CVU (Cluster Verification Utility) options that you require. You can select whether to enable them, stop or continue the execution and from where the CVU will run.

Note:

CVU checks are only available for Cluster provisioning and will not be enabled for other instances.

Select the Pre/Post scripts you will use, these need to be previously uploaded into the Software Library. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.

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. 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 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 with Oracle Automatic Storage Management

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 with Oracle Automatic Storage Management

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

- Access from Databases, then select Administration drop down menu and click 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 ${\bf Add}$ 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.



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.

- 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 <code>ORACLE HOME/assistants/dbca/templates</code>.

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. Starting with Enterprise Manager 13c Release 5 Update 8 (13.5.0.8) you can specify a non-DBSNMP user for database monitoring, edit the DBSNMP field with the user name. If the user does not exist it will be created.



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.

Note:

Disk names must be provided in RAW format, for example: /prd/odb1, /prd/odb2. If the disk names are provided in ORCL: format Enterprise Manager will not let you submit and continue.

- 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 Custom Properties, 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 SQL scripts. 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. 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.

Select whether you wish to pause after the Prerequisite check to analyze the results and whether Prerequisite System Checks are to be performed.

If you are running a cluster provisioning, select the CVU (Cluster Verification Utility) options that you require. You can select whether to enable them, stop or continue the execution and from where the CVU will run.



Note:

CVU checks are only available for Cluster provisioning and will not be enabled for other instances.

Select the Pre/Post scripts you will use, these need to be previously uploaded into the Software Library. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.

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 Configure page, 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 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:

- Access from Databases, then select Administration drop down menu and click 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.
- In the Specify OS Users page, specify the operating system user for the Oracle Home for the database.



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.



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.

- From the Enterprise menu, select Provisioning and Patching, and then select Database Provisioning. Alternatively you can access from Databases, then select Administration drop down menu and click 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.



Do not select Provisioning Profile in this step.

- Click Next.
- 5. Click on **Setup Hosts** and provide the required details.
- 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.



 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:

- 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.



5

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, 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.



If you are provisioning an Oracle 21c Grid Infrastructure database there are additional steps to configuring the GIMR component.

Table 5-1 Getting Started with Provisioning Oracle Grid Infrastructure

Step Description	Reference Links
Step 1 Selecting the Use Case This chapter covers a few us provisioning Oracle Grid Infr Select the use case that bes your requirements.	astructure. configuring database with Oracle ASM, see

Table 5-1 (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:

- Log in as a designer, and access from Databases, then select Administration drop down menu and click 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.
- 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.
- To setup GIMR select Setup Local GIMR. Select the Grid Infrastructure Home location.



In Oracle Database 21c the GIMR must be configured from a separate Oracle Home. This is done as part of the provisioning process.

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.
- 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 https://docs.oracle.com/en/database/oracle/oracle-database/index.html



Note:

If you are setting up GIMR on an Oracle 21c database add the following groups:

- Backup DBA group for GIMR (OSBACKUPDBA)
- Oracle Data Guard group for GIMR (OSDGDBA)
- Encryption Key Management group for GIMR (OSKMDBA)
- Real Application Cluster administration group for GIMR (OSRACDBA)

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:

https://docs.oracle.com/en/database/oracle/oracle-database/index.html

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.

- 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.
- GIMR Database Home, a location on the destination host where the Grid
 Infrastructure Management Repository software can be provisioned. This is the Oracle
 home directory for GIMR only on 21C databases. Do not select a location that is a sub
 directory of the Oracle Base for Grid Infrastructure or database.

For Grid Infrastructure, Oracle Base is /u01/app/user and Oracle Home is %ORACLE_BASE%/sihahome. 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.

Starting with Oracle Enterprise Manager 13 Release 5 Update 23 (13.5.0.23) you can configure the failure group while creating ASM disk groups. The ASM disk within the disk group can belong to a different failure group or a quorum failure group. To set the failure group select the disks via the **Select for Suitable Disks** radio button and select all three disks from the list and set the **Failure Group** and if they are in **Quorum** enable the check mark. Or to manually select the disks, click on**Select Disks Manually** radio button and enter the **Disk List** paths separated by commas for example: /scratch/disks/disk_1, /scratch/disks/disk_3, and then enter the **Failure Group** for example: FGRP1 [QUORUM], FGRP2, FGRP3. Once entered 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.

You can also enable the Oracle ASM Filter Driver on ASM versions 12.2.0.1 or greater. To enable select the **Configure ASM Filter Driver** check box. For more information on ASM Filter Driver installation see: Configuring Oracle ASM Filter Driver During Installation in *Oracle Automatic Storage Management Administrator's Guide*.

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. Starting with Enterprise Manager 13c Release 5 Update 8 (13.5.0.8) you can specify a non-DBSNMP user for database monitoring, edit the DBSNMP field with the user name. If the user does not exist it will be created.

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.

Note:

Disk names must be provided in **RAW** format, for example: /prd/odb1, /prd/odb2. If the disk names are provided in ORCL: format Enterprise Manager will not let you submit and continue.

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 Custom Properties, 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 customSQL scripts. 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. 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.



Select whether you wish to pause after the Prerequisite check to analyze the results and whether Prerequisite System Checks are to be performed.

Select the Pre/Post scripts you will use, these need to be previously uploaded into the Software Library. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.

Click Next.

- 18. 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.
- 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.** 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**.
- 22. 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.

- 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 Operator role, launch the saved deployment procedure. Add targets for provisioning and provide values for configurable fields in the deployment procedure.
- 25. 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.
- 26. 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:

- Log in as a designer, and access from Databases, then select Administration drop down menu and click Database Provisioning.
- 2. In the Database Procedures page, select the Oracle Database Deployment Procedure and click **Launch**. The Oracle Database provisioning wizard is launched.
- 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.
- To setup GIMR select Setup Local GIMR. Select the Grid Infrastructure Home location.



In Oracle Database 21c the GIMR must be configured from a separate Oracle Home. This is done as part of the provisioning process.

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.
- 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 Install and Upgrade

Note:

If you are setting up GIMR on an Oracle 21c database add the following groups:

- Backup DBA group for GIMR (OSBACKUPDBA)
- Oracle Data Guard group for GIMR (OSDGDBA)
- Encryption Key Management group for GIMR (OSKMDBA)
- Real Application Cluster administration group for GIMR (OSRACDBA)

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.

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.
- GIMR Database Home, a location on the destination host where the Grid
 Infrastructure Management Repository software can be provisioned. This is the Oracle
 home directory for GIMR only on 21C databases. Do not select a location that is a sub
 directory of the Oracle Base for Grid Infrastructure or database.

For Grid Infrastructure, Oracle Base is /u01/app/user and Oracle Home is %ORACLE_BASE%/sihahome. 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 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.





If the Deployment Procedure fails, then review log files described in Reviewing Log Files.



6

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
- Oracle Real Application Clusters Database Topology
- Oracle Flex Clusters and Flex ASM
- 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
- Using No Root Credentials for Provisioning Oracle Real Application Clusters (Oracle RAC)
 Databases

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 6-1 Getting Started with Provisioning Oracle Grid Infrastructure and Oracle RAC Databases

Step	Description	Reference Links
Step 1		To learn about the topology, see Oracle Real Application Clusters Database Topologyand Oracle Flex Clusters and Flex ASM.



Table 6-1 (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 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 and ASM. 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 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 6-1 shows a typical Oracle RAC database topology that you can provision using Enterprise Manager.

ASM Disk RAID Diskgroup named -Diskgroup named QUORUM. DATA consumed Consumed by by RAC DB GI for OCR and for Datafiles VDSK NIC NIC Public Private Storage Network Network Network NIC NIC NIC 1 ... N Identical nodes Datafiles in ASM diskgroup named DATA OCR / Voting in ASM Disk Group named QUORUM

Figure 6-1 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 multihomed 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
- Overview of Oracle Flex ASM in Automatic Storage Management Administrator's Guide.

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 and ASM

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 and ASM

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

- Log in as a designer, and access from Databases, then select Administration drop down menu and click 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.
- 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.



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.

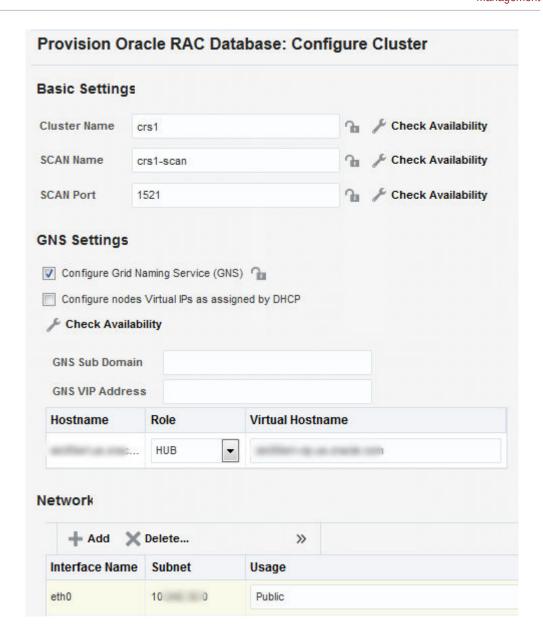
You can choose either Automatic Storage Management or File System for Oracle RAC Database. The storage type for Grid Infrastructure is set to Automatic Storage Management with Flex redundancy.

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.





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.

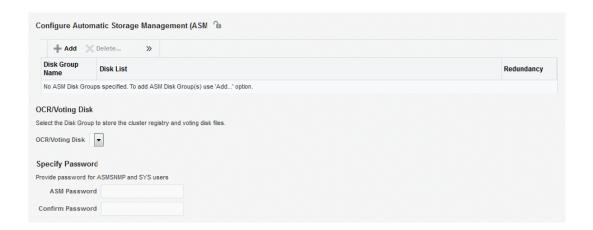
Click **Add** to add an interface and specify the **Interface Name** and **Interface Subnet** and click **OK**. Select the Usage as **Public**, **Private**, **ASM**, **ASM & Private**, or **Do Not Use** if you do not want to use the interface.

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 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**. 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.



Note:

ASMFD pre labelled disks (AFD:DISK1) are not supported by RDBMS, however you can configure ASMFD during grid Install.on ASM versions 12.2.0.0 or greater. To enable select the **Configure ASM Filter Driver** check box. For more information on ASM Filter Driver installation see: Configuring Oracle ASM Filter Driver During Installation in *Oracle Automatic Storage Management Administrator's Guide*.

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 <code>ORACLE HOME/assistants/dbca/templates</code>.

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. Starting with Enterprise Manager 13c Release 5 Update 8 (13.5.0.8) you can specify a non-DBSNMP user for database monitoring, edit the DBSNMP field with the user name. If the user does not exist it will be created.

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 Custom Properties page, select custom SQL 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. 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.

Select whether you wish to pause after the Prerequisite check to analyze the results and whether Prerequisite System Checks are to be performed.

Select the Pre/Post scripts you will use, these need to be previously uploaded into the Software Library. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.

Click Next.

- 19. 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.
- 20. Click the Compliance Standards link.
- **21.** In the Configuration Standards Target Association page, select a Compliance Standard to be associated with the database. Click **Next**.
- 22. In the Configure page, click Next.
- 23. 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**.
- 24. 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.

25. 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.

- **26.** In the Operator role, launch the saved deployment procedure. Add targets for provisioning and provide values for configurable fields in the deployment procedure.
- 27. 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.
- 28. 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.



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:

- Log in as a designer, and access from **Databases**, then select **Administration** drop down menu and click **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.
- 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.



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 <code>ORACLE HOME/assistants/dbca/templates</code>.

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. Starting with Enterprise Manager 13c Release 5 Update 8 (13.5.0.8) you can specify a non-DBSNMP user for database monitoring, edit the DBSNMP field with the user name. If the user does not exist it will be created.

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.

 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 Custom Properties page, select custom SQL 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. 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.

Select whether you wish to pause after the Prerequisite check to analyze the results and whether Prerequisite System Checks are to be performed.

If you are running a cluster provisioning, select the CVU (Cluster Verification Utility) options that you require. You can select whether to enable them, stop or continue the execution and from where the CVU will run.



CVU checks are only available for Cluster provisioning and will not be enabled for other instances.

Select the Pre/Post scripts you will use, these need to be previously uploaded into the Software Library. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.

Click Next.

- 16. 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.
- 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 Configure page, 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



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:

- Log in as a designer, and access from Databases, then select Administration drop down menu and click Database Provisioning.
- In the Database Procedures page, select the Provision Oracle RAC Database
 Deployment Procedure and click Launch. The Oracle RAC Database provisioning wizard
 is launched.
- 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

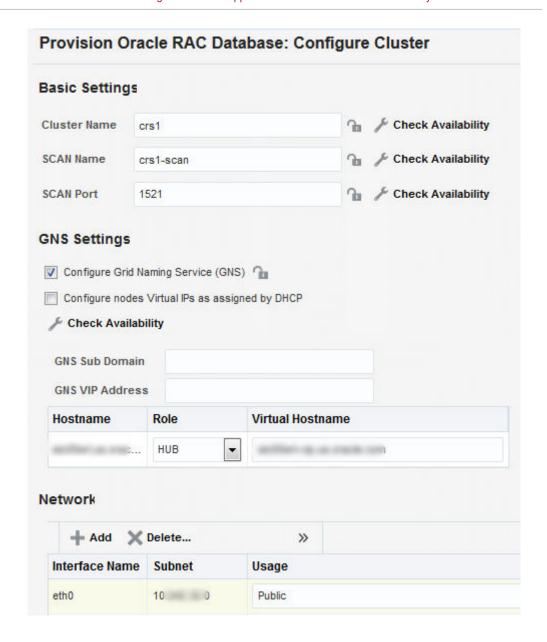
Click Next. You will come back to the Configure page. If you have configured the source and destination location for the software, the Configure Software task will have a Configured status.

- 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.





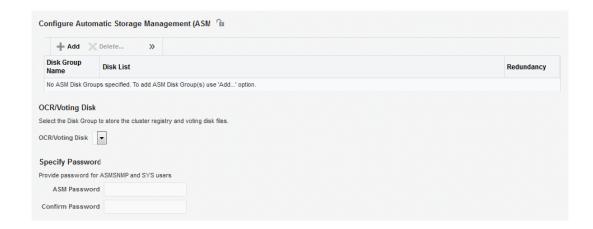
In the GNS Settings section, select **Configure GNS** and specify the **GNS Sub System** and **GNS VIP Address** if you want virtual host names outside the cluster to have dynamically assigned names.

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 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 <code>ORACLE HOME/assistants/dbca/templates</code>.

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. Starting with Enterprise Manager 13c Release 5 Update 8 (13.5.0.8) you can specify a non-DBSNMP user for database monitoring, edit the DBSNMP field with the user name. If the user does not exist it will be created.



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.

 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 Custom Properties page, select custom Pre/Post scripts from the Software Library or your local disk. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning. 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. 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.

Select whether you wish to pause after the Prerequisite check to analyze the results and whether Prerequisite System Checks are to be performed.

If you are running a cluster provisioning, select the CVU (Cluster Verification Utility) options

If you are running a cluster provisioning, select the CVU (Cluster Verification Utility) options that you require. You can select whether to enable them, stop or continue the execution and from where the CVU will run.



CVU checks are only available for Cluster provisioning and will not be enabled for other instances.

Select the Pre/Post scripts you will use, these need to be previously uploaded into the Software Library. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.

Click Next.

- 19. 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.
- 20. Click the Compliance Standards link.
- **21.** In the Configuration Standards Target Association page, select a Compliance Standard to be associated with the database. Click **Next**.
- 22. In the Configure page, click Next.
- 23. 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**.
- 24. 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.

- 25. 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.
- **26.** In the Operator role, launch the saved deployment procedure. Add targets for provisioning and provide values for configurable fields in the deployment procedure.
- 27. 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.

28. 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.



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:

- 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)
- After the command line is run manually using root user for each step, click Confirm. The
 database provisioning process then continues till it completes.



7

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, 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 7-1 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.
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 Enterprise Manager, 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:

- Access from Databases, then select Administration drop down menu and click 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.
- 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.

- 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 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 **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 \mathbf{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 <code>ORACLE HOME/assistants/dbca/templates</code>.

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. Starting with Enterprise Manager 13c Release 5 Update 8 (13.5.0.8) you can specify a non-DBSNMP user for database monitoring, edit the DBSNMP field with the user name. If the user does not exist it will be created. Select the type of Oracle RAC database, whether Policy Managed or Admin Managed. Specify the Service Name.



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 todata 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 SQL 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 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.

Select whether you wish to pause after the Prerequisite check to analyze the results and whether Prerequisite System Checks are to be performed.

If you are running a cluster provisioning, select the CVU (Cluster Verification Utility) options that you require. You can select whether to enable them, stop or continue the execution and from where the CVU will run.



CVU checks are only available for Cluster provisioning and will not be enabled for other instances.

Select the Pre/Post scripts you will use, these need to be previously uploaded into the Software Library. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.



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

This chapter explains how you can provision Oracle Real Application Clusters (Oracle RAC) using Oracle Enterprise Manager, 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 8-1 Getting Started with Provisioning Oracle RAC

Step	Description	Reference Links
Step 1	Understanding the Components Provisioned	To learn about the core components that are provisioned, see Core Components Deployed When Provisioning Oracle RAC.
	Understand the core components provisioned.	
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 8-1 (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
Step 4	Running the Deployment Procedure Run the Deployment Procedure to successfully provision Oracle RAC.	 Using Archived Software Binaries. 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).



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 Enterprise Manager.

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, you must meet the following prerequisites.

Prerequisites for Designers

- Ensure that you meet the prerequisites described in Setting Up Your 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. 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.
- 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.
- 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 unmask 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:



- Access from Databases, then select Administration drop down menu and click Database Provisioning.
- 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.



- 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.



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.
- Click Next.
- 5. On the Credentials/Schedule page, do the following:
 - 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.

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).

 In the Target Host(s) Credentials section, provide the credentials as described in Step 6 (a).



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.

- 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.



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. 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 Custom Properties 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 8-2.
- 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. 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.

Select whether you wish to pause after the Prerequisite check to analyze the results and whether Prerequisite System Checks are to be performed.

If you are running a cluster provisioning, select the CVU (Cluster Verification Utility) options that you require. You can select whether to enable them, stop or continue the execution and from where the CVU will run.



CVU checks are only available for Cluster provisioning and will not be enabled for other instances.

Select the Pre/Post scripts you will use, these need to be previously uploaded into the Software Library. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.

Click Next.

- 10. 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.
- **11.** After the Deployment Procedure ends successfully, instrument the database to collect configuration information.

Provisioning Oracle Real Application Clusters Using Gold Image

This section describes how you can provision a gold image of Oracle RAC.

Note:

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.

This section covers the following:

- Prerequisites for Provisioning Oracle Real Application Clusters Using Gold Image
- Procedure for Provisioning Oracle Real Application Clusters Using Gold Image

Prerequisites for Provisioning Oracle Real Application Clusters 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 you create gold images of existing Oracle RAC Database and Oracle Grid Infrastructure.
 - To understand how you can create a gold image, see Setting Up Database Provisioning.
- 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: https://docs.oracle.com/en/database/oracle/oracle-database/18/clustering.html
- 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.
- While selecting the source, remember to remove sqlnet.ora from the list of files mentioned in Files to Exclude.
- Ensure that the unmask 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:

- Access from Databases, then select Administration drop down menu and click 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.

- On the Select Source page, do the following:
 - 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.



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.

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*).



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.

- 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.



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.



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. 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 Custom Options page, do the following:



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 8-2.
- 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. 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.
 - Select whether you wish to pause after the Prerequisite check to analyze the results and whether Prerequisite System Checks are to be performed.

If you are running a cluster provisioning, select the CVU (Cluster Verification Utility) options that you require. You can select whether to enable them, stop or continue the execution and from where the CVU will run.



CVU checks are only available for Cluster provisioning and will not be enabled for other instances.

Select the Pre/Post scripts you will use, these need to be previously uploaded into the Software Library. For more information on Pre/Post scripts in Database Provisioning see: Pre/Post Scripts for Database Provisioning.

- 10. 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.

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.
- 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
- 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: Overview of Installing Oracle Clusterware in Oracle Clusterware Administration and Deployment Guide
- 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

Creating Database Templates.

- 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.
- Ensure that the unmask 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:

- Access from Databases, then select Administration drop down menu and click 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.



If you do not see the required component in the Software Library, then follow the workaround described in Troubleshooting Issues.

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.



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.



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 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.



- Click Next.
- 5. On the Credentials/Schedule page, do the following:
 - 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.

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*).

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*.

- 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.

 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.



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.

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:



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 8-2.
- 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 8-2 Configuration Page - Element Description

Element	Description		
Bonding Device Name	Specify the name of the bond to be created. For example, bond0.		
Subnet Mask	Specify the subnet mask for the IP address. For example, 255.255.255.0.		
Default Gateway	Specify the default gateway for the bonding device. For example, 10.1.2.3.		
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 /etc/resolv.conf 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, eth1,eth2,eth3.		
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 send out sequentially on each bonded slave interface beginning with the first one available. 		
Domain Name	Specify the domain name for the assigned host name. For example, example.com.		



Table 8-2 (Cont.) Configuration Page - Element Description

Element	Description
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-backup 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 milmon 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 milmon 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.co	nodel- priv.example.com	-	nodel- vip.example.com	-
node2.example.co	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:

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)
- 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, 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 9-1 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.
- 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:

- Access from Databases, then select Administration drop down menu and click Database Provisioning.
- 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 *.



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.

 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.

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*).

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.
- 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.



10

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 10-1 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 10-1 (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.
- REMOVE_ANY_TARGET Enterprise Manager privilege

Procedure for Deleting the Entire Oracle RAC

To delete the entire Oracle RAC, follow these steps:

- Access from Databases, then select Administration drop down menu and click 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.
- On the Delete/Scale down Oracle Real Application Clusters page, do the following:
 - In the Select Cluster section, click the torch icon for **Select Cluster** and select an Oracle Clusterware instance that you want to delete. Along with the selected Oracle Clusterware, the associated Oracle RAC and ASM instances will also be deleted. The table displays details about the member nodes that are part of the selected Oracle Clusterware.



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.

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*).

- In the Schedule section, schedule the Deployment Procedure to run either immediately or later.
- Click Review.
- On the Review page, review the details you have provided for deleting Oracle RAC, and click Submit.



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 Enterprise Manager. 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.

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:

 Access from Databases, then select Administration drop down menu and click 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.
- 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.

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 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.



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.
- On the Review page, review the details you have provided for deleting or scaling down Oracle RAC, and click Submit.



If the Deployment Procedure fails, then review log files described in Reviewing Log



Provisioning Oracle Database Replay Client

This chapter explains how you can provision Oracle Database Replay Client using Oracle Enterprise Manager, 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 11-1 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 11-1 (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 go 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 assign the correct privileges.
- 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:

- Access from Databases, then select Administration drop down menu and click Database Provisioning..
- 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.



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.



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 have the necessary privileges to switch to the locked account for performing the provisioning operation.

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.



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.
- On the Review page, review the details you have provided for provisioning an Oracle Database Replay Client, and click Submit.





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 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.
- 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:

- Access from Databases, then select Administration drop down menu and click 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.



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.



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.

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.



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**.



If the Deployment Procedure fails, then review log files described in Reviewing Log

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.



If you want to create a component for the software binaries of Oracle Database Replay Client, do not save the ship home 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.
- 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. Access from **Databases**, then select **Administration** drop down menu and click **Database Provisioning**.
- 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.



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 **Ship home**. 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.

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.



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
- Data Guard
- Fast-Start Failover REST API

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:

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.
- 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.

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 nondefault 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:



AES 256 is the only encryption mode available when using Oracle Database 23c deployments.

- 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.



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:



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.
- 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.

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:



AES 256 is the only encryption mode available when using Oracle Database 23c deployments.

- 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.



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.
- 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.



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.

- (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.
- On the Databases page, you see a list of databases. Select the database you want to manage an existing standby database.
- 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.

- 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.

Review the data for the configuration and standby database. Click Finish

Data Guard

Data Guard Switchover REST API



Table 12-1 Data Guard Switchover

Feature	Description		
URL	https:// <oms_host>:<oms_console_port>/em/websvcs/restful/db/dataguard/switchover</oms_console_port></oms_host>		
Request header	Authorization: basic		
	Content-Type: application/json		
Body			
	{		
	<pre>"primaryTargetName" : "<primary name="" target="">", "primaryTargetType" : "oracle_database rac_database", "standbyTargetName" : "<standby name="" target="">", "standbyTargetType" : "oracle_database rac_database", ["primaryDbCredsName" : "<primary credential="" database="" name="">",] ["primaryHostCredsName" : "<primary credential="" database="" host="" name="">",] ["standbyDbCredsName" : "<standby credential="" database="" name="">",] ["standbyHostCredsName" : "<standby credential="" database="" host="" name="">",] ["swapJobs" : "false true",] ["swapThresholds" : "false true",] ["tdeWalletCredsName" : "<transparent credentials="" data="" database="" encryption="" of="" primary="" the="" wallet="">"] ["tdeWalletPassword" : "<transparent data="" database="" encryption="" of="" password="" primary="" the="" wallet="">"] }</transparent></transparent></standby></standby></primary></primary></standby></primary></pre>		
Request method	POST		
Sample Body			
Jap.0 200)	{		
	<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			
Campio (Cosponac	<pre>"status": "Succeeded", "submissionName": "dbname_DataguardSwitchover_091418163542", "executionGUID": "75D3A124B36D1A4DE0536A8A7B0A80E0", "databaseName": "dbname", "message": "Deployment procedure for switchover to dbname:oracle_database successfully submitted." }</pre>		

Table 12-1 (Cont.) Data Guard Switchover

Feature	Description
Supported Since Release	EM DB 13.3.2.0.0



Parameters listed within a square bracket are optional.

Table 12-2 Data Guard Switchover Parameters

Parameter Description primaryTargetName Primary database Enterprise Manager target name. Can be either a single-instance database or a cluster database. primaryTargetType Primary database target type. Allowed values are either single-instance database (oracle_database) or cluster database (rac_database). standbyTargetName Standby database Enterprise Manager target name. Can be either a single-instance database or a cluster database. standbyTargetType Standby database target type. Allowed values Allowed values are either single-instance database (oracle_database) or cluster database (rac_database). primaryDbCredsName Primary database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used. standbyDbCredsName Primary database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used. standbyDbCredsName Standby database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used. standbyHostCredsName Standby database named credential for an operating system user who can access the standby database Oracle home. swapJobs 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 V		
single-instance database or a cluster database. primaryTargetType Primary database target type. Allowed values are either single-instance database (oracle_database) or cluster database (rac_database). Standby database Enterprise Manager target name. Can be either a single-instance database or a cluster database. StandbyTargetType Standby database target type. Allowed values Allowed values are either single-instance database (oracle_database) or cluster database PrimaryDbCredsName Primary database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used. 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. StandbyDbCredsName StandbyHostCredsName Standby database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used. Standby database named credential for an operating system user who can access the standby database Oracle home. StandbyHostCredsName Standby database host named credential for an operating system user who can access the standby database Oracle home. SwapJobs 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. Specifies that monitoring settings will be swapped between the primary and standby databases after the switchover is complete. Default Value: Not enabled Transparent Data Encryption wallet credentials for the primary database. Use create_named_credential -cred_name=WC1 -cred_type=GenericPassword -auth_target_type=' <system>' -attributes="GENERIC_PASSWORD:<pri>Password>"</pri></system>	Parameter	Description
Standby TargetName Standby TargetType Standby database Enterprise Manager target name. Can be either a single-instance database or a cluster database. Standby TargetType Standby database target type. Allowed values Allowed values are either single-instance database (oracle_database) or cluster database (rac_database). Primary DbCredsName Primary database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used. 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. StandbyDbCredsName StandbyDbCredsName Standby database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used. StandbyHostCredsName Standby database host named credential for an operating system user who can access the standby database Oracle home. SwapJobs 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 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=WCl - cred_type=GenericPassword -auth_target_type=' <system>' - attributes="GENERIC_PASSWORD:<primary database="" password="" tde="" wallet="">"</primary></system>	primaryTargetName	
single-instance database or a cluster database. standbyTargetType Standby database target type. Allowed values Allowed values are either single-instance database (oracle_database) or cluster database (rac_database). primaryDbCredsName Primary database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used. 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. StandbyDbCredsName Standby database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used. StandbyHostCredsName Standby database host named credential for an operating system user who can access the primary database Oracle home. StandbyHostCredsName Standby database host named credential for an operating system user who can access the standby database Oracle home. 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 swapThresholds Specifies that monitoring settings will be swapped between the primary and standby databases after the switchover is complete. Default Value: Not enabled 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>"</primary_database_tde_wallet_password></system>	primaryTargetType	
single-instance database (oracle_database) or cluster database (rac_database). primaryDbCredsName	standbyTargetName	
role. If this argument is not specified, preferred credential will be used. primaryHostCredsName 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. StandbyDbCredsName Standby database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used. StandbyHostCredsName Standby database host named credential for an operating system user who can access the standby database Oracle home. SwapJobs 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. Specifies that monitoring settings will be swapped between the primary and standby databases after the switchover is complete. Default Value: Not enabled 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>"</primary_database_tde_wallet_password></system>	standbyTargetType	single-instance database (oracle_database) or cluster database
can access the primary database Oracle home. If this argument is not specified, preferred credential will be used. StandbyDbCredsName Standby database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used. StandbyHostCredsName Standby database host named credential for an operating system user who can access the standby database Oracle home. SwapJobs 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. Specifies that monitoring settings will be swapped between the primary and standby databases after the switchover is complete. Default Value: Not enabled tdeWalletCredsName 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_walletpassword>"</primary_database_tde_walletpassword></system>	primaryDbCredsName	
role. If this argument is not specified, preferred credential will be used. standbyHostCredsName Standby database host named credential for an operating system user who can access the standby database Oracle home. 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. Specifies that monitoring settings will be swapped between the primary and standby databases after the switchover is complete. Default Value: Not enabled 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>"</primary_database_tde_wallet_password></system>	primaryHostCredsName	can access the primary database Oracle home. If this argument is not
who can access the standby database Oracle home. 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. Specifies that monitoring settings will be swapped between the primary and standby databases after the switchover is complete. Default Value: Not enabled 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>"</primary_database_tde_wallet_password></system>	standbyDbCredsName	
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. Specifies that monitoring settings will be swapped between the primary and standby databases after the switchover is complete. Default Value: Not enabled 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="" password="" tde="" wallet="">"</primary></system>	standbyHostCredsName	
and standby databases after the switchover is complete. Default Value: Not enabled 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="">"</primary_database_tde_wallet></system>	swapJobs	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
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="" password="" tde="" wallet="">"</primary></system>	swapThresholds	
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="" password="" tde="" wallet="">"</primary></system>		Default Value: Not enabled
<pre>cred_type=GenericPassword -auth_target_type='<system>' - attributes="GENERIC_PASSWORD:<primary database="" password="" tde="" wallet="">"</primary></system></pre>	tdeWalletCredsName	Use create_named_credential verb to create these credentials as shown
tdeWalletPassword Transparent Data Encryption wallet password for the primary database.		<pre>cred_type=GenericPassword -auth_target_type='<system>' - attributes="GENERIC_PASSWORD:<primary database="" pre="" tde="" wallet<=""></primary></system></pre>
	tdeWalletPassword	Transparent Data Encryption wallet password for the primary database.

Convert Standby Database REST API

Table 12-3 Convert Standby Database

Feature	Description		
URL	https:// <oms_host>:<oms_console_port>/em/websvcs/restful/db/dataguard/convertStandby</oms_console_port></oms_host>		
Request header	Authorization: basic		
	Content-Type: application/json		
Body			
	<pre>"dbTargetName" : "<standby name="" target="">", "dbTargetType" : "oracle_database rac_database", ["dbHostCredName" : "<standby credential="" database="" host="" name="">",] ["dbCredName" : "<standby credential="" database="" name="">",] ["convertToDbRole" : "PHYSICAL STANDBY SNAPSHOT STANDBY"] }</standby></standby></standby></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 Release	EM DB 13.3.2.0.0		

Table 12-4 Convert Standby Database Parameters

Parameter	Description
dbTargetName	Standby database Enterprise Manager target name. Can be either a single-instance database or a cluster database.
dbTargetType	Standby database target type. Allowed values Allowed values are either single-instance database (oracle_database) or cluster database (rac_database).
dbHostCredName	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.
dbCredName	Standby database named credential for a user with SYSDBA or SYSDG role. If this argument is not specified, preferred credential will be used.

Table 12-4 (Cont.) Convert Standby Database Parameters

Parameter	Description
convertToDbRole	Standby database role to be converted to given role. Possible values are Physical Standby and Snapshot Standby.

Fast-Start Failover REST API

This chapter provides details of the REST APIs for Fast-Start Failover and Observers supported in Enterprise Manager. When Fast-Start Failover is enabled, the Data Guard broker automatically fails over to a synchronized standby site in the event of a disaster at the primary site, requiring no intervention by the DBA. In addition, applications are automatically notified of the role transition.

- Fast-Start Failover Management REST API
- Observer Management REST API

Fast-Start Failover Management REST API

Fast-Start Failover is a feature that allows the Oracle Data Guard broker to failover a failed primary database automatically to a predetermined standby database. This feature increases the availability of the database by eliminating the need for user involvement as part of the failover process.

Fast-Start Failover Management REST API

Table 12-5 Fast-Start Failover

Feature	Description
URL	https:// <oms_host>:<oms_console_port>/em/websvcs/restful/db/dataguard/configure/fsfo</oms_console_port></oms_host>
Request header	Authorization: Basic
	Content-Type: application/json



Table 12-5 (Cont.) Fast-Start Failover

Feature Description Body "action" : "<edit | enable | disable | change mode>", "primary" : { "dbTargetName" : "<primary database target name>", "dbTargetType" : "oracle database | rac database", "dbCredName" : "<primary database credential name>", "dbHostCredName" : "<primary database host credential name>", "dbGIHostCredName" : "<grid infrastructure credential name>" "dbObserverConnectIdentifier" : "<ObserverConnectIdentifier value of the primary database>", "dbFastRecoveryArea" : "<Fast recovery area location of the database - default V\$RECOVERY FILE DEST>", "dbFastRecoveryAreaSize" : "<Fast recovery area size (MB) of the database - default 2x database size>", "dbFlashbackRetentionTime" : "<Flashback retention time hours - default 24h>" } **,** "candidateFSFOTargets" : "dbTargetName" : "<standby database target name>", "dbTargetType" : "<oracle database |</pre> rac database>", "dbCredName" : "<standby database credential name>", "dbHostCredName" : "<standby database host credential name>", "dbGIHostCredName" : "<grid infrastructure credential name>" "dbObserverConnectIdentifier" : "<ObserverConnectIdentifier value of the primary database>", "dbFastRecoveryArea" : "<Fast recovery area location of the database - default V\$RECOVERY FILE DEST>", "dbFastRecoveryAreaSize" : "<Fast recovery area size(MB) of the database - default 2x database size>", "dbFlashbackRetentionTime" : "<Flashback retention time hours - default 24h>" }, "observer" : { "name": "<observer name - default is the hostname>", "hostName" : "<observer host name>", "oracleHome" : "<observer oracle home>", "hostCredName" : "<observer host credential name>" "alternateObserver" : { "hostName" : "<alternate observer host name>", "oracleHome" : "<alternate observer oracle home>", "hostCredName" : "<alternate observer host credential name>" "failoverThreshold" : "<failover threshold>",

Table 12-5 (Cont.) Fast-Start Failover

Feature Description

```
"lagLimit" : "<lag limit>",
                             "failoverConditions" : "<failover conditions>",
                             "calloutScripts" : {
                                     "hostName" : "<host name>",
                                     "hostCredentials" : "<host credential>"
                                     "preCallout" : "<path of pre role change callout
                         script>",
                                     "postCallout" : "<path of post role change callout
                         script>",
                                     "sucFileName" : "<Name of the suc file that>",
                                     "errorFileName" : "<Name of the .err file>",
                                     "fsfoPreCalloutTimeout" : "<Value in seconds>",
                                     "fsfoActionOnPreCalloutFailure" : "<STOP |
                         CONTINUE>"
                                 },
                             "enableFlashBackDatabase" : "<Yes | No>",
                             "autoReinstatePrimary" : "<Yes | No>",
                             "autoShutdownPrimary" : "<Yes | No>"
                             "observeOnly" : "<Yes | No - Default Value: No>"
                             "includeOraError240" : "<Yes | No -Default Value: Yes>",
Body for Disable
                                 "action" : "disable",
                                 ["stopAllObservers" : "<Yes | No - Default value: Yes,
                         applicable only for 'disable' action>"]
                                 "primary" : {
                                         "dbTargetName" : "<Primary database Enterprise
                         Manager target name>",
                                         "dbTargetType" : "oracle database |
                         rac database",
                                         ["dbCredName" : "<Primary database named
                         credential for a user with SYSDBA or SYSDG role. Default
                         Value: Preferred credential will be used>",]
                                         ["dbHostCredName" : "<Primary database host
                         named credential for an operating system user who can access
                         the primary database Oracle home. Default Value: Preferred
                         credential will be used>",]
                                         ["dbGIHostCredName" : "<Grid Infrastructure
                         named credentials for an operating system user who can access
                         the grid infrastructure Oracle home of the Primary Database.
                         Default Value: Preferred credential will be used>"]
                                  }
```

Table 12-5 (Cont.) Fast-Start Failover

```
Feature
                         Description
Body for Change Mode
                                 "action" : "change mode",
                                 "primary" : {
                                         "dbTargetName" : "<Primary database Enterprise
                         Manager target name>",
                                         "dbTargetType" : "oracle database |
                         rac database",
                                         ["dbCredName" : "<Primary database named
                         credential for a user with SYSDBA or SYSDG role. Default
                         Value: Preferred credential will be used>",]
                                         ["dbHostCredName" : "<Primary database host
                         named credential for an operating system user who can access
                         the primary database Oracle home. Default Value: Preferred
                         credential will be used>",]
                                         ["dbGIHostCredName" : "<Grid Infrastructure
                         named credentials for an operating system user who can access
                         the grid infrastructure Oracle home of the Primary Database.
                         Default Value: Preferred credential will be used>"]
Notes
                         Observers and alternate observers are optional.
                         POST
Request method
Sample Response
                             "status" : "Succeeded",
                             "submissionName" : "DG FAST START FAIL OVER *",
                             "executionGUID" : "<execution guid>",
                             "databaseName" : "<primary database name>",
                             "message" : "Deployment Procedure for configuring fast-
                         start failover of the Data Guard configuration determined from
                         the primary database <primary database name> is successfully
                         submitted."
Supported Since Release
                         EM DB 13.4
```

Table 12-6 Fast-Start Failover

Feature	Description
URL	https:// <oms_host>:<oms_console_port>/em/websvcs/restful/db/dataguard/configure/fsfo</oms_console_port></oms_host>
Request header	Authorization: Basic
	Content-Type: application/json

Table 12-6 (Cont.) Fast-Start Failover

Feature Description Body "action" : "<edit | enable | disable | change mode>", "primary" : { "dbTargetName" : "<primary database target name>", "dbTargetType" : "oracle database | rac database", "dbCredName" : "<primary database credential name>", "dbHostCredName" : "<primary database host credential name>", "dbGIHostCredName" : "<grid infrastructure credential name>" "dbObserverConnectIdentifier" : "<ObserverConnectIdentifier value of the primary database>", "dbFastRecoveryArea" : "<Fast recovery area location of the database - default V\$RECOVERY FILE DEST>", "dbFastRecoveryAreaSize" : "<Fast recovery area size (MB) of the database - default 2x database size>", "dbFlashbackRetentionTime" : "<Flashback retention time hours - default 24h>" }, "candidateFSFOTargets" : "dbTargetName" : "<standby database target name>", "dbTargetType" : "<oracle database |</pre> rac database>", "dbCredName" : "<standby database credential name>", "dbHostCredName" : "<standby database host credential name>", "dbGIHostCredName" : "<grid infrastructure credential name>" "dbObserverConnectIdentifier" : "<ObserverConnectIdentifier value of the primary database>", "dbFastRecoveryArea" : "<Fast recovery area location of the database - default V\$RECOVERY FILE DEST>", "dbFastRecoveryAreaSize" : "<Fast recovery area size(MB) of the database - default 2x database size>", "dbFlashbackRetentionTime" : "<Flashback retention time hours - default 24h>" }, "observer" : { "name": "<observer name - default is the hostname>", "hostName" : "<observer host name>", "oracleHome" : "<observer oracle home>", "hostCredName" : "<observer host credential name>" "alternateObserver" : { "hostName" : "<alternate observer host name>", "oracleHome" : "<alternate observer oracle home>", "hostCredName" : "<alternate observer host credential name>" "failoverThreshold" : "<failover threshold>",

Table 12-6 (Cont.) Fast-Start Failover

```
Feature
                         Description
                             "lagLimit" : "<lag limit>",
                             "failoverConditions" : "<failover conditions>",
                             "calloutScripts" : {
                                     "hostName" : "<host name>",
                                     "hostCredentials" : "<host credential>"
                                     "preCallout" : "<path of pre role change callout
                         script>",
                                     "postCallout" : "<path of post role change callout
                         script>",
                                     "sucFileName" : "<Name of the suc file that>",
                                     "errorFileName" : "<Name of the .err file>",
                                     "fsfoPreCalloutTimeout" : "<Value in seconds>",
                                     "fsfoActionOnPreCalloutFailure" : "<STOP |
                         CONTINUE>"
                                 },
                             "enableFlashBackDatabase" : "<Yes | No>",
                             "autoReinstatePrimary" : "<Yes | No>",
                             "autoShutdownPrimary" : "<Yes | No>"
                             "observeOnly" : "<Yes | No - Default Value: No>"
                             "includeOraError240" : "<Yes | No -Default Value: Yes>",
                         }
Notes
                         Observers and alternate observers are optional.
                         POST
Request method
Sample Response
                             "status" : "Succeeded",
                             "submissionName" : "DG_FAST_START_FAIL_OVER_*",
                             "executionGUID" : "<execution guid>",
                             "databaseName" : "<primary database name>",
                             "message" : "Deployment Procedure for configuring fast-
                         start failover of the Data Guard configuration determined from
                         the primary database <primary database name> is successfully
                         submitted."
Supported Since Release
                         EM DB 13.4
```

Fast-Start Failover Parameters

Parameter	Description	Required	Value Default Value
action	Operation to be performed on the fast-start failover configuration. • enable: Enable fast-start failover. • disable: Disable fast-start failover and stop the observer. • edit: Change the fast-start failover target database and thresholds, and start or restart the observer process.	Required	edit enable disable
dbTargetName	Primary or standby database Enterprise Manager target names.	Required	N/A
dbTargetType	Primary or standby database target types.	Required	oracle_database rac_database
dbCredName	Primary or standby database named credential for a user with SYSDBA or SYSDG role.	Optional	Preferred credential will be used.
dbHostCredName	Primary or standby database host named credential for an operating system user who can access the database Oracle home.	Optional	Preferred credential will be used.
dbGIHostCredName	Grid Infrastructure named credential for an operating system user who can access the grid infrastructure Oracle home of the database.	Optional	N/A
hostName	Discovered host where Enterprise Manager will start the observer or alternate observer.	Required	N/A
oracleHome	Oracle home on the observer or alternate host (as specified by "hostName") from which Enterprise Manager will start the observer or alternate observer	Required	N/A



Parameter	Description	Required	Value Default Value
observer	A DGMGRL (Data Guard command line interface) client that continuously monitors the primary and target standby databases, evaluates whether failover is necessary, and initiates a fast-start failover when selected conditions occur.	Required	N/A
alternateObserver	Enterprise Manager supports the automatic restart of an alternate observer on a different host if it detects that the observer process had failed in the earlier host.	Required	N/A
hostCredName	Observer or alternate observer host named credential for an operating system user who can access the observer Oracle home.	Required	N/A
failoverThreshold	Amount of time in seconds the primary database must be out of contact with the observer and the standby database before a fast-start failover is initiated.	Optional	30 Seconds.
lagLimit	Amount of time in seconds the standby database is allowed to fall behind the primary database, beyond which a fast-start failover will not be allowed.	Optional	30 Seconds.
failoverConditions	A comma separated list of failover conditions which should cause a fast-start failover if detected on the primary database. 1 - Corrupted Controlfile 2 - Corrupted Dictionary 3 - Inaccessible Logfile 4 - Stuck Archiver 5 - Datafile Offline	Optional	1, 2, 5



Parameter	Description	Required	Value Default Value
autoReinstatePrimary	Causes the former primary database to be automatically reinstated if a fast-start failover was initiated because the primary database was either isolated or had crashed.	Optional	Yes No Default: Yes
autoShutdownPrimary	Causes the primary database to automatically shut down if fast-start failover is enabled.	Optional	Yes No
dbObserverConnectIden tifier	Primary database ObserverConnectIdentifi er configurable property value	Optional	N/A
dbFastRecoveryArea	Fast recovery area location of the database	Optional	V\$RECOVERY_FILE_ DEST
dbFastRecoveryAreaSiz e	Fast recovery area size of the database in MB	Optional	Twice the database size
dbFlashbackRetentionTi me	Flashback retention time in hours	Optional	24 hours
name	Name of the observer	Optional	hostname
hostName	host name of host containing the pre and post role change callout scripts	Optional	N/A
hostCredentials	host credential name of host containing pre and post role change callout scripts	Optional	N/A
preCallout	Full path of pre role change callout script	Optional	N/A
postCallout	Full path of post role change callout script	Optional	N/A
sucFileName	Name of the success (.suc) file that the precallout script creates	Optional	precallout.suc
errorFileName	Name of the error (.err) file that the precallout script creates	Optional	precallout.err
fsfoPreCalloutTimeout	Callout timeout, numeric value expressed in seconds	Optional	600 seconds
fsfoActionOnPreCalloutF ailure	Precallout action in case of failure. STOP or CONTINUE	Optional	STOP CONTINUE Default: Top
enableFlashBackDataba se	Enables flash back of database.	Optional	YES NO Default: NO



Observer Management REST API



This REST API is applicable to Oracle Database version 12.2 and higher.

Configure Multiple Observers for a Data Guard Configuration

Table 12-7 Start Observers

Feature	Description
URL	https:// <oms_host>:<oms_console_port>/em/websvcs/restful/db/dataguard/configure/observers</oms_console_port></oms_host>
Request header	Authorization: basic Content-Type: application/json



Table 12-7 (Cont.) Start Observers

Feature Description Body "action" : "<start | stop | setMaster | show | delete alternate observer | stopall | editCallout | removeCallout>", "primary" : { "dbTargetName" : "<Primary database Enterprise Manager target name>", "dbTargetType" : "oracle database | rac database", ["dbCredName" : "<Primary database named credential for a user with SYSDBA role. Default Value: Preferred credentials will be used>",] ["dbHostCredName" : "<Primary database host named credential for an operating system user who can access the primary database Oracle home. Default Value: Preferred credentials will be used.>"] }, "observers" : ["name" : "<Observer name>", "hostName" : "<Host name where observer will be started>", "oracleHome" : "<Oracle home location on the host where observer will be started>", "hostCredName" : "<Host named credential for user who can access the observer Oracle home>", "hostWalletCredName" : "<Named credential of the Oracle wallet configured in the observer Oracle home>", ["isMaster" : "<Y|N>",] ["runTimeDataFile" : "<Runtime data file path and name for the observer. If not specified, the file fsfo.dat will be created under the oracle base>",] ["logFile" : "<The full path of the observer log file>"1 }, { }, { . . . } 1 ["calloutScripts" : { "hostName" : "<Name of the host containing the pre and post role change callout scripts>", "hostCredentials" : "<Credential name of the host containing pre and post role change callout scripts>" "preCallout" : "<Absolute path of pre role change callout script on the hostName>", "postCallout" : "<Absolute path of post role change callout script on the hostName>", "sucFileName" : "<Name of the suc file that the precallout script would create on the observer host | Default - precallout.suc>", "errorFileName" : "<Name of the err file that the

Table 12-7 (Cont.) Start Observers

Feature Description

seconds of the time for whic the observer will wait for the suc/err file to be generated, default value - 600>",

"fsfoActionOnPreCalloutFailure": "<Action that indicates what the observer will do if the suc file does not exist even if after FastStartFailoverPreCalloutTimeout seconds, or the error file is detected before
FastStartFailoverPreCalloutTimeout seconds passed. Initiate the failover - CONTINUE, or not do failover - STOP. Values - STOP | CONTINUE, default value - STOP>"

```
},
```



Table 12-7 (Cont.) Start Observers

Feature

Description

Body for Start

```
"action" : "start",
    "primary" : {
        "dbTargetName" : "<Primary database Enterprise Manager
        "dbTargetType" : "oracle database | rac database",
        ["dbCredName" : "<Primary database named credential
for a user with SYSDBA or SYSDG role. Default Value: Preferred
credentials will be used>",]
        ["dbHostCredName" : "<Primary database host named
credential for an operating system user who can access the
primary database Oracle home. Default Value: Preferred
credentials will be used.>"]
    "observers" : [
            "name" : "<Observer name>",
            "hostName" : "<Host name where observer will be
started>",
            "oracleHome" : "<Oracle home location on the host
where observer will be started>",
            "hostCredName" : "<Host named credential for user
who can access the observer Oracle home>",
            "hostWalletCredName" : "<Named credential of the
Oracle wallet configured in the observer Oracle home>",
            ["isMaster" : "<Y|N>",]
            ["runTimeDataFile" : "<Runtime data file path and
name for the observer. If not specified, the file fsfo.dat
will be created under the oracle base>",]
            ["logFile" : "<The full path of the observer log
file>"]
        },
        {
        },
        {
            . . .
    ["calloutScripts" : {
            "hostName" : "<Name of the host containing the pre
and post role change callout scripts>",
            "hostCredentials" : "<Credential name of the host
containing pre and post role change callout scripts>"
            "preCallout" : "<Absolute path of pre role change
callout script on the hostName>",
            "postCallout" : "<Absolute path of post role
change callout script on the hostName>",
            "sucFileName" : "<Name of the suc file that the
precallout script would create on the observer host | Default
- precallout.suc>",
            "errorFileName" : "<Name of the err file that the
precallout script would create on the observer host | Default
- precallout.err>",
            "fsfoPreCalloutTimeout" : "<Numeric value in
```

Table 12-7 (Cont.) Start Observers

Feature Description seconds of the time for whic the observer will wait for the suc/err file to be generated, default value - 600>", "fsfoActionOnPreCalloutFailure" : "<Action that indicates what the observer will do if the suc file does not exist even if after FastStartFailoverPreCalloutTimeout seconds, or the error file is detected before FastStartFailoverPreCalloutTimeout seconds passed. Initiate the failover - CONTINUE, or not do failover - STOP. Values -STOP | CONTINUE, default value - STOP>" },] Body for Stop "action" : "stop", "primary" : { "dbTargetName" : "<Primary database Enterprise Manager "dbTargetType" : "oracle database | rac database", ["dbCredName" : "<Primary database named credential for a user with SYSDBA or SYSDG role. Default Value: Preferred credentials will be used>",] ["dbHostCredName" : "<Primary database host named credential for an operating system user who can access the primary database Oracle home. Default Value: Preferred credentials will be used.>"] } "observers" : [{"name" : "<observer 1 name>"}, {"name" : "<observer 2 name>"}, {"name" : "<observer 3 name>"} Body for Set Master "action" : "setMaster", "primary" : { "dbTargetName" : "<Primary database Enterprise Manager target name>", "dbTargetType" : "oracle database | rac database", ["dbCredName" : "<Primary database named credential for a user with SYSDBA or SYSDG role. Default Value: Preferred credentials will be used>",] ["dbHostCredName" : "<Primary database host named credential for an operating system user who can access the primary database Oracle home. Default Value: Preferred credentials will be used.>"] "masterObserver" : "<master observer name>"

}

Table 12-7 (Cont.) Start Observers

Feature

Description

Body for Show and Stop All

```
"action" : "show | stopall",
    "primary" : {
        "dbTargetName" : "<Primary database Enterprise Manager
target name>",
        "dbTargetType" : "oracle_database | rac_database",
        ["dbCredName" : "<Primary database named credential
for a user with SYSDBA or SYSDG role. Default Value: Preferred
credentials will be used>",]
        ["dbHostCredName" : "<Primary database host named
credential for an operating system user who can access the
primary database Oracle home. Default Value: Preferred
credentials will be used.>"]
    }
}
```



Table 12-7 (Cont.) Start Observers

Feature

Description

Body for Edit Callout

```
"action" : "editCallout",
    "primary" : {
        "dbTargetName" : "<Primary database Enterprise Manager
        "dbTargetType" : "oracle database | rac database",
        ["dbCredName" : "<Primary database named credential
for a user with SYSDBA or SYSDG role. Default Value: Preferred
credentials will be used>",]
        \hbox{["dbHostCredName" : "<Primary database host named}\\
credential for an operating system user who can access the % \left( 1\right) =\left( 1\right) \left( 1\right) 
primary database Oracle home. Default Value: Preferred
credentials will be used.>"]
    ["calloutScripts" : {
            "hostName" : "<Name of the host containing the pre
and post role change callout scripts>",
            "hostCredentials" : "<Credential name of the host
containing pre and post role change callout scripts>"
            "preCallout" : "<Absolute path of pre role change
callout script on the hostName>",
            "postCallout" : "<Absolute path of post role
change callout script on the hostName>",
            "sucFileName" : "<Name of the suc file that the
precallout script would create on the observer host | Default
- precallout.suc>",
            "errorFileName" : "<Name of the err file that the
precallout script would create on the observer host | Default
- precallout.err>",
            "fsfoPreCalloutTimeout" : "<Numeric value in
seconds of the time for whic the observer will wait for the
suc/err file to be generated, default value - 600>",
            "fsfoActionOnPreCalloutFailure" : "<Action that
indicates what the observer will do if the suc file does not
exist even if after FastStartFailoverPreCalloutTimeout
seconds, or the error file is detected before
FastStartFailoverPreCalloutTimeout seconds passed. Initiate
the failover - CONTINUE, or not do failover - STOP. Values -
STOP | CONTINUE, default value - STOP>"
        },
    ]
```



Table 12-7 (Cont.) Start Observers

Feature Description Body for Remove Callout "action" : "removeCallout", "primary" : { "dbTargetName" : "<Primary database Enterprise Manager "dbTargetType" : "oracle database | rac database", ["dbCredName" : "<Primary database named credential for a user with SYSDBA or SYSDG role. Default Value: Preferred credentials will be used>",] ["dbHostCredName" : "<Primary database host named credential for an operating system user who can access the primary database Oracle home. Default Value: Preferred credentials will be used.>"] "observers" : ["name" : "<Observer name>", "hostName" : "<Host name where observer will be started>", "oracleHome" : "<Oracle home location on the host where observer will be started>", "hostCredName" : "<Host named credential for user who can access the observer Oracle home>" }, { }, {] Notes A maximum of 3 observers can be assigned. Sample Response "status" : "Succeeded", "submissionName": "ConfigureObserver *", "executionGUID" : "<execution guid>", "databaseName" : "<primary database name>", "message" : "Deployment procedure to start the observer(s) for the Data Guard Configuration associated with the primary database <primary database name> is successfully submitted." **POST** Request method Supported Since Release EM DB 13.4



Table 12-8 Delete, Show, or Stop All Observers

```
Feature
                         Description
URL
                         https://<OMS HOST>:<OMS CONSOLE PORT>/em/websvcs/
                         restful/db/dataguard/configure/observers
                         Authorization: basic
Request header
                         Content-Type: application/json
Body
                             "action" : "<delete alternate observer | show | stopall>",
                             "primary" : {
                                 "dbTargetName" : "<primary database target name>",
                                 "dbTargetType" : "oracle database | rac database",
                                 "dbCredName" : "<primary database credential name>",
                                 "dbHostCredName" : "<primary database host credential
                         name>"
                             }
Sample Response
                         Show
                             "status": "Succeeded",
                             "submissionName" : "ShowObserver *",
                             "executionGUID" : "<execution guid>",
                             "databaseName" : "<primary database name>",
                             "message" : "Deployment procedure to list the observer(s)
                         for the Data Guard Configuration associated with the primary
                         database <primary database name> is successfully submitted."
                         Delete Alternate Observer
                             "status" : "Succeeded",
                             "submissionName" : "DeleteAlternateObserver *",
                             "executionGUID" : "<execution guid>",
                             "databaseName" : "<primary database name>",
                             "message" : "Deployment procedure to delete the alternate
                         observer for the Data Guard Configuration associated with the
                         primary database <primary database name> is successfully
                         submitted."
                         Stop All
                             "status" : "Succeeded",
                             "submissionName" : "StopObserver *",
                             "executionGUID" : "<execution guid>",
                             "databaseName" : "<primary database name>",
                             "message" : "Deployment procedure to stop the observer(s)
                         for the Data Guard Configuration associated with the primary
                         database <primary database name> is successfully submitted."
                         }
                         POST
Request method
Supported Since Release
                         EM DB 13.4
```

Table 12-9 Stop Observers

Feature	Description	
URL	https:// <oms_host>:<oms_console_port>/em/websvcs/restful/db/dataguard/configure/observers</oms_console_port></oms_host>	
Request header	Authorization: basic	
	Content-Type: application/json	
Body	<pre>{ "action" : "stop", "primary" : { "dbTargetName" : "<primary database="" name="" target="">", "dbTargetType" : "oracle_database rac_database", "dbCredName" : "<primary credential="" database="" name="">", "dbHostCredName" : "<primary credential="" database="" host="" name="">" } "observers" : [{"name" : "<observer 1="" name="">"}, {"name" : "<observer 2="" name="">"}, {"name" : "<observer 3="" name="">"} } }</observer></observer></observer></primary></primary></primary></pre>	
Sample Response	<pre>"status" : "Succeeded", "submissionName": "StopObserver_*", "executionGUID" : "<execution guid="">", "databaseName" : "<primary database="" name="">", "message" : "Deployment procedure to stop the observer(s) for the Data Guard Configuration associated with the primary database <primary database="" name=""> is successfully submitted." }</primary></primary></execution></pre>	
Request method	POST	
Supported Since Release	EM DB 13.4	

Table 12-10 Set Master Observer

Feature	Description
URL	https:// <oms_host>:<oms_console_port>/em/websvcs/restful/db/dataguard/configure/observers</oms_console_port></oms_host>
Request header Authorization: basic	
	Content-Type: application/json

Table 12-10 (Cont.) Set Master Observer

Feature	Description	
Body	<pre>"action" : "setMaster", "primary" : { "dbTargetName" : "<primary database="" name="" target="">", "dbTargetType" : "oracle_database rac_database", "dbCredName" : "<primary credential="" database="" name="">", "dbHostCredName" : "<primary credential="" database="" host="" me="">" } "masterObserver" : "<master name="" observer="">"</master></primary></primary></primary></pre>	
Sample Response	<pre>"status" : "Succeeded", "submissionName": "SetMasterObserver_*", "executionGUID" : "<execution guid="">", "databaseName" : "<primary database="" name="">", "message" : "Deployment procedure to set the master observer for the Data Guard Configuration associated with the primary database <primary database="" name=""> is successfully submitted." }</primary></primary></execution></pre>	
Request method	POST	
Supported Since Release	EM DB 13.4	



Table 12-11 Observer Parameters

Parameter	Description	Required	Value Default Value
action	The action the API will execute. All the operations are mutually exclusive. They can't be combined and must be executed separately. start: Starts one or more fast-start failover observers on the specified hosts and Oracle homes for the Data Guard configuration associated with the specified primary database. delete_alternate_observer: Deletes the Enterprise Manager alternate observer and its associated corrective action job. The Enterprise Manager alternate observer feature is not supported for Oracle Database version 12.2 and higher. show: Displays the list of observers configured for a Data Guard configuration associated with the specified primary database. stop: Stops the respective observer(s) for a Data Guard configuration associated with the specified primary database. stopall: Stops all the observers. setMaster: Configures the master observer for a Data Guard configuration associated with the specified primary database.	Required	start delete_a lternate _observe r show stop stopall setMaste r
dbTargetName	Primary database Enterprise Manager target name.	Required	N/A
dbTargetType	Primary database target type.	Required	oracle_d atabase rac_data base
dbCredName	Primary database named credential for a user with SYSDBA or SYSDG role.	[OPTIONAL]	Preferred credentials will be used.
dbHostCredName	Primary database host named credential for an operating system user who can access the primary database Oracle home.	[OPTIONAL]	
name	Observer name.	Required	N/A
hostName	Host name where observer will be started.	Required	N/A
oracleHome	Oracle home location on the host where the observer will be started.	Required	N/A
hostCredName	Host named credential for user who can access the observer Oracle home.	Required	N/A
hostWalletCredNa me	Named credential of the Oracle wallet configured in the observer Oracle home.	Required	N/A
isMaster	Defines if the observer is master.	[OPTIONAL]	Y N
masterObserver	Required for Set Master Observer Defines the master observer.	[OPTIONAL]	N/A



Table 12-11 (Cont.) Observer Parameters

Parameter	Description	Required	Value Default Value
runTimeDataFile	Runtime data file path and name for the observer. If not specified, the file fsfo.dat will be created under the Oracle base.	[OPTIONAL]	N/A
logFile	The full path of the observer log file.	[OPTIONAL]	N/A
hostName	host name of host containing the pre and post role change callout scripts	Optional	N/A
hostCredentials	host credential name of host containing pre and post role change callout scripts	Optional	N/A
preCallout	Full path of pre role change callout script	Optional	N/A
postCallout	Full path of post role change callout script	Optional	N/A
sucFileName	Name of the success (.suc) file that the precallout script creates	Optional	precallo ut.suc
errorFileName	Name of the error (.err) file that the precallout script creates	Optional	precallo ut.err
fsfoPreCalloutTime out	Callout timeout, numeric value expressed in seconds	Optional	600 seconds
fsfoActionOnPreCal loutFailure	Precallout action in case of failure. STOP or CONTINUE	Optional	STOP CONTINUE Default: Top



Managing Pluggable Databases Using Enterprise Manager

This chapter explains how you can manage pluggable databases (PDBs) using Oracle Enterprise Manager, 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.

If you are managing a PDB on OCI resources make sure that all OCI required prerequistes and discovery have been performed. For more information, see: (Optional) Deploying Agents on OCI Resources .

Table 13-1 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 13-1 (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 schema, schema objects, and non-schema 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.

Enterprise Manager 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.

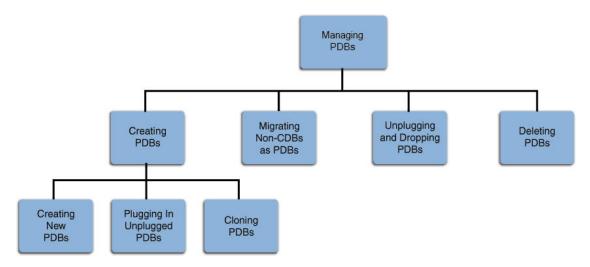
Note:

To manage the PDB lifecycle using Enterprise Manager, you must have the 12.1.0.3 Enterprise Manager for Oracle Database plug-in, or a later version, deployed. To delete PDBs using Enterprise Manager, 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 Administrator's Guide.*

Figure 13-1 provides a graphical overview of how you can manage the PDB lifecycle in Enterprise Manager.

Figure 13-1 Managing Pluggable Databases



For more information about PDBs and CDBs plug-ins, 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

If you are provisioning a PDB using TDE (Transparent Data Encryption) for additional information on configuring TDE see: Using Transparent Data Encryption in Oracle Database Advanced Security Guide



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 Enterprise Manager. In particular, it contains the following:

- Prerequisites for Creating a New Pluggable Database
- Creating a New Pluggable Database

If you are managing a PDB on OCI resources make sure that all OCI required prerequistes and discovery have been performed. For more information, see: (Optional) Deploying Agents on OCI Resources.

Prerequisites for Creating a New Pluggable Database

Before creating a new PDB using Enterprise Manager, ensure that you meet the following prerequisites:

- Oracle Software Library (Software Library) must be set up in Enterprise Manager.
 For information on how to set up Software Library in Enterprise Manager, see Setting Up Oracle Software Library.
- The CDB within which you want to create a PDB must exist, and must be a Enterprise Manager target.



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.
- Make sure to have bug 31303249 fix installed in the Oracle Home that the CDB belongs to.
 If it does not have the fix the PDBADMIN password set in Enterprise Manger won't be
 honored and the target CDB database will be created with PDB's without resetting the
 PDBADMIN password.
- If you are provisioning a PDB using TDE (Transparent Data Encryption) review additional information and actions on configuring TDE see: Using Transparent Data Encryption in Oracle Database Advanced Security Guide

Creating a New Pluggable Database

To create a new PDB in a CDB using Cloud Control, follow these steps:

 Access from Databases, then select Administration drop down menu and click Database Provisioning. In the Database Provisioning page, in the Related Links section of the left menu pane, click Provision Pluggable Databases.





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.

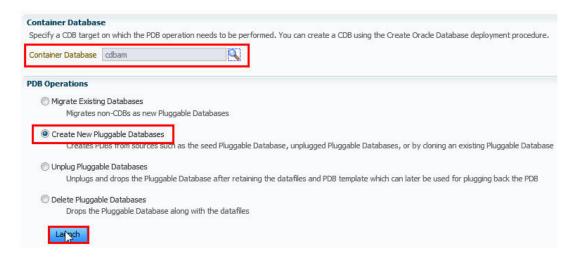
Note:

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 13-2 displays the Provision Pluggable Database Console.

Figure 13-2 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.

- 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 credentials with Enterprise Manager, you can select **Preferred** or **Named**. Otherwise, you can select **New** and enter the credentials.

Figure 13-3 displays the Creation Options page.

cdbam (Container Database) () Creation Options Identification Storage Schedule Review Create Pluggable Database: Creation Options Pluggable Database Creation Options Oreate a New PDB Plug an Unplugged PDB Clone an Existing PDB Full Clone Snap Clone Source PDB **Container Database Host Credentials** Specify credentials for host "abc.example.com" Credential Preferred Named New Preferred Credential Name Database Host Credentials 💌 Credential Details Default preferred credentials are not set.

Figure 13-3 Creating a New Pluggable Database: Creation Options Page

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.

WNOTE Host credentials are required to perform validations and initiate PDB creation on the CDB host or cluster.



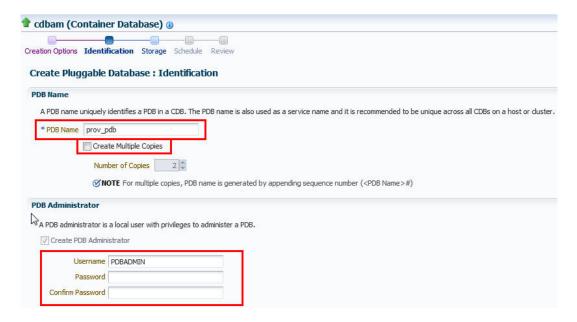
If you choose to create multiple PDBs, then the unique name you enter here is used as a prefix for all PDBs, and the suffix is a numeric value that indicates the count of PDBs.

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.

In the PDB Administrator section, enter the credentials of the admin user account you need to create for administering the PDB.

Figure 13-4 displays the Identification page.

Figure 13-4 Creating a New Pluggable Database: Identification Page

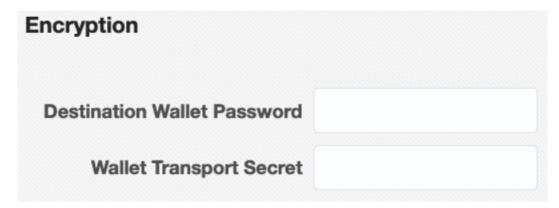




If you choose to create multiple PDB, then an admin user account is created for each PDB that you create, with the same set of the specified credentials.

If the CDB has TDE enabled encryption, enter the Wallet password and the Wallet Transport Secret for the CDB. The Wallet Transport Secret is a pass-phrase used to decrypt the Wallet Password when transporting a TDE encrypted PDB to a CDB that does not have the wallet password for that particular database.

Figure 13-5 TDE Options

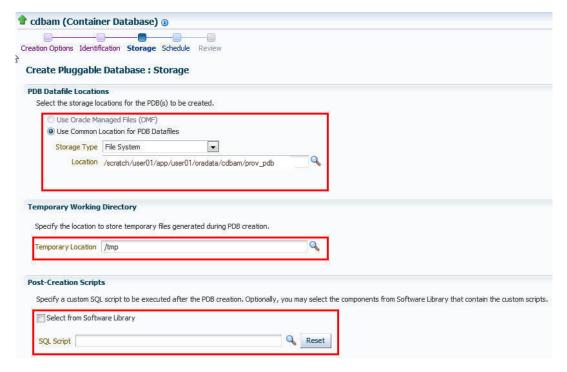


- 10. Click Next.
- 11. In the Storage page, in the PDB Datafile Locations section, select the type of location where you want to store the datafiles.
 - If the target CDB (CDB in which you are creating the PDB) is enabled with Oracle Managed Files and if you want to use the same, then select **Use Oracle 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 13-6 displays the Storage page.

Figure 13-6 Creating a New Pluggable Database: Storage Page



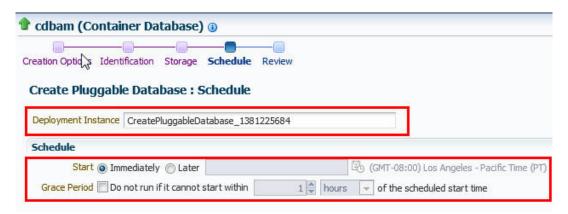
- 14. Click Next.
- **15.** In the Schedule page, enter a unique deployment procedure instance name and a schedule for the deployment. The instance name you enter here helps you identify and track the progress of this procedure on the Procedure Activity page.

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 13-7 displays the Schedule page.

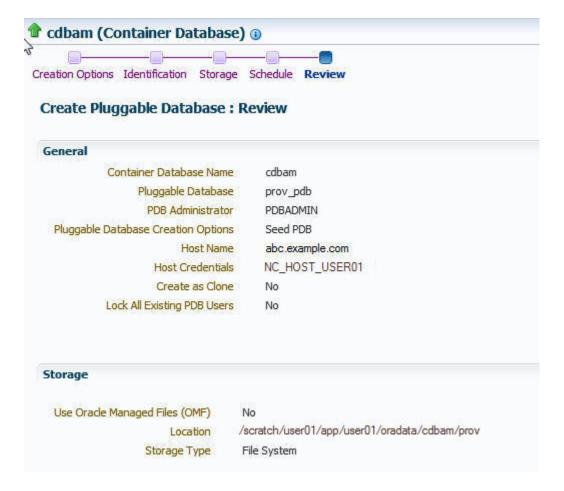
Figure 13-7 Creating a New Pluggable Database: Schedule Page



- 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 13-8 displays the Review page.

Figure 13-8 Creating a New Pluggable Database: Review Page

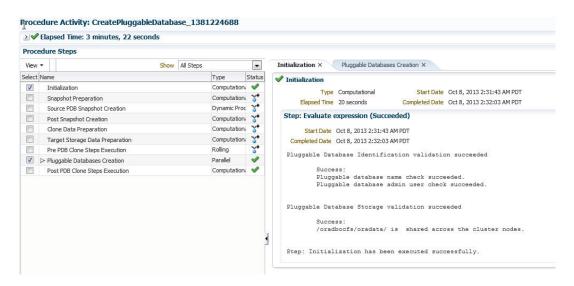




18. In the Procedure Activity page, view the status of the procedure. From the Procedure Actions menu, you can select Debug to set the logging level to Debug, and select Stop to stop the procedure execution.

Figure 13-9 displays the Procedure Activity page.

Figure 13-9 Creating a New Pluggable Database: Procedure Activity Page



When you create a new PDB, theEnterprise 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 Enterprise Manager. 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 Enterprise Manager, ensure that you meet the following prerequisites:

- Oracle Software Library (Software Library) must be set up in Enterprise Manager.
 For information on how to set up Software Library in Enterprise Manager, 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 Enterprise Manager target.



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.



For TDE Enabled Pluggable Databases ensure that the target CDB **must** be TDE enabled if the Unplugged PDB is TDE enabled. The **Wallet Location**, the **Wallet Password** and the **Wallet Transport Secret** of the Unplugged TDE enabled PDB must exist when plugged in a CDB that is already TDE enabled.

Plugging In an Unplugged Pluggable Database

To plug in an unplugged PDB to a CDB using Cloud Control, follow these steps:

 Access from Databases, then select Administration drop down menu and click 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:

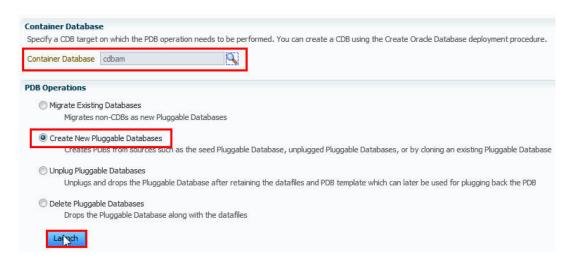
Skip this step if you have accessed the Provision Pluggable Database Console from the CDB's home page.

- 3. In the PDB Operations section, select Create New Pluggable Databases.
- 4. Click Launch.

Figure 13-10 displays the Provision Pluggable Database Console.



Figure 13-10 Provision Pluggable Database Console





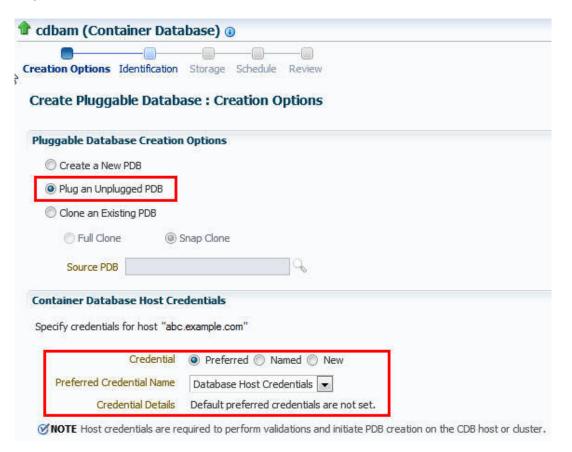
You will be prompted to log in to the database if you have not already logged in to it through Enterprise Manager. Make sure you log in using *sysdba* user account credentials.

- In the Creation Options page of the Create Pluggable Database Wizard, in the Pluggable Database Creation Options section, select Plug an Unplugged PDB.
 - If the destination CDB selected, is Data Guard enabled and there are multiple standby databases, you will be prompted to select a reference standby CDB. This will be used as a reference by the workflow to copy the datafiles to other standby databases.
- 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 13-11 displays the Creation Options page.



Figure 13-11 Plugging In an Unplugged Pluggable Database: Creation Options Page



- 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.



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.



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 13-12 displays the Identification page.

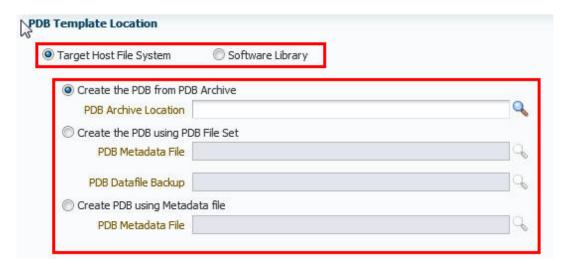
Figure 13-12 Plugging In an Unplugged Pluggable Database: Identification Page



- **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 13-13 displays the PDB Template Location section of the Identification page.

Figure 13-13 Plugging In an Unplugged Pluggable Database: PDB Template Location Section



If the PDB is TDE Enabled follow these additional steps:

- If the Unplugged PDB is TDE enabled and PDB Template Location is chosen as *Target Host File System*, Enter the **Wallet Password** and the **Wallet Transport Secret** of the Unplugged PDB.
- If the PDB Template Location is chosen as *Software Library*, enter only the **Wallet Password** for the Unplugged PDB.



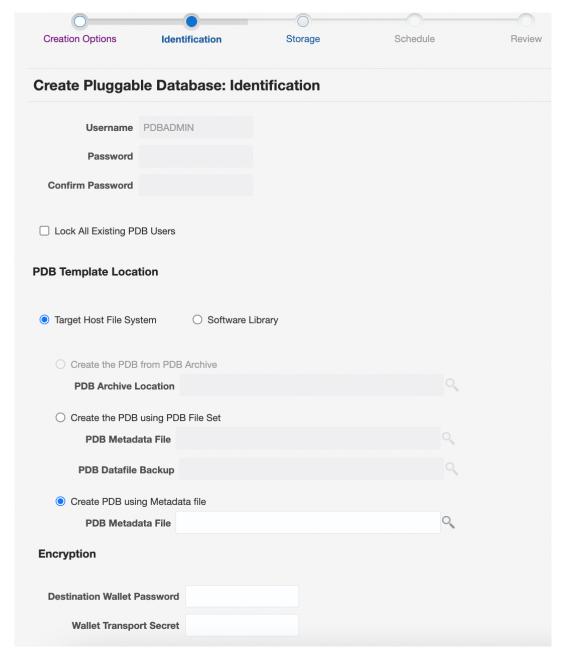


Figure 13-14 Plug in a TDE Enabled PDB

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 13-15 displays the Storage page.

Figure 13-15 Plugging In an Unplugged Pluggable Database: Storage Page





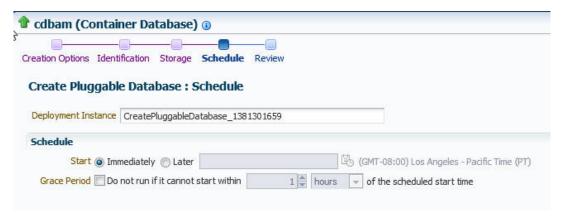
- 15. Click Next.
- **16.** In the Schedule page, enter a unique deployment procedure instance name and a schedule for the deployment. The instance name you enter here helps you identify and track the progress of this procedure on the Procedure Activity page.

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 13-16 displays the Schedule page.

Figure 13-16 Plugging In an Unplugged Pluggable Database: Schedule Page



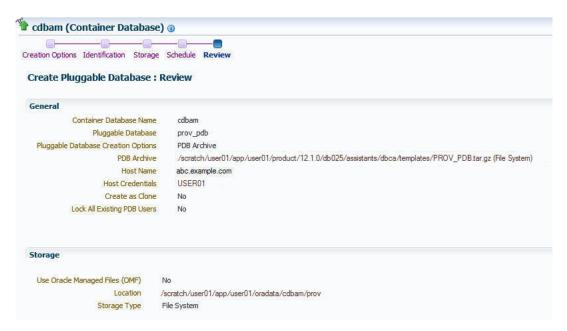
- 17. Click Next.
- **18.** In the Review page, review the details you have provided for the deployment procedure. If you are satisfied with the details, click **Submit.**

If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

Figure 13-17 displays the Review page.



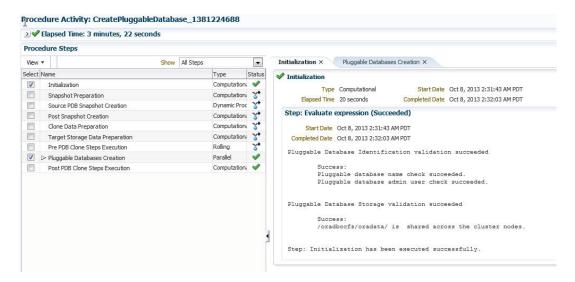
Figure 13-17 Plugging In an Unplugged Pluggable Database: Review Page



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 13-18 displays the Procedure Activity page.

Figure 13-18 Plugging In an Unplugged Pluggable Database: Procedure Activity Page



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.

If you are managing a PDB on OCI resources make sure that all OCI required prerequisites and discovery have been performed. For more information, see: (Optional) Deploying Agents on OCI Resources.



It is recommended that PDB cloning be performed with the Clone Wizard to utilize the latest features. For more information see: Creating a Full Clone Database

In particular, this section 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. 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. 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 latest Oracle Database and Storage Management Framework plug-ins must be 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.
- If the destination CDB has GLOBAL_NAMES initialization parameter set to trueand if the user has specified existing DB link, then the name of the DB Linkshould be same as that of the source PDB.



Due to the latest Perl stack (5.32), Snap Clone has also revised the version of the modules used to communicate with the storage. These modules will not work with the Perl available in agent versions prior to 13.5. In order to continue using the Snap Clone functionality, it is important to upgrade the agents used communicate with the storage filers to 13.5 along with the OMS upgrade.

If you are using the Snap Clone functionality with your current Enterprise Manager setup:

- Identify the agents that are registered to communicate with the storage for performing the Snap Clone operations.
- From the Setup menu, click on Provisioning and Patching, and then select Storage Registration.
- 3. Include these agents as part of the plan to upgrade the OMS.

For information on Snap Clone functionality, 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:



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).

 Access from Databases, then select Administration drop down menu and click Database Provisioning. In the Database Provisioning page, in the Related Links section of the left menu pane, click Provision Pluggable Databases.



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**.

In the Provision Pluggable Database Console, in the CDB section, select the CDB to which you want to add the cloned PDB.



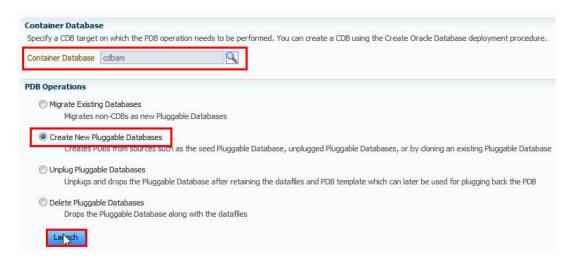
Skip this step if you have accessed the Provision Pluggable Database Console from the CDB's home page.

- In the PDB Operations section, select Create New Pluggable Databases.
- 4. Click Launch.

Figure 13-19 displays the Provision Pluggable Database Console.



Figure 13-19 Provision Pluggable Database Console





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.**

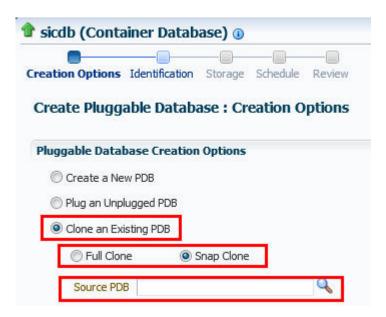
To clone a PDB using the traditional method of cloning the PDB datafiles, select **Full Clone.** Use this method if you want to clone a PDB for long term usage. This method is ideal for load testing, when you plan to make significant data updates to the PDB clone. However, this method takes a longer 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 13-20 displays the PDB Creation Options section of the Creation Options page.

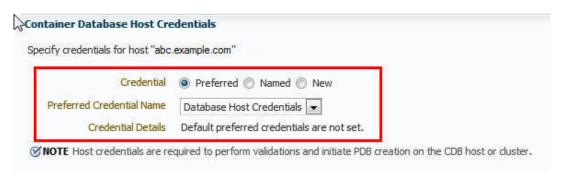
Figure 13-20 Cloning a Pluggable Database: Specifying Creation Options



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 13-21 displays the CDB Host Credentials section of the Creation Options page.

Figure 13-21 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:

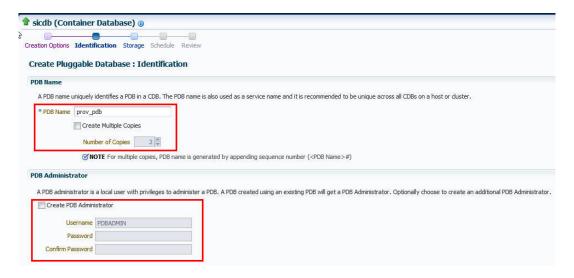
If you choose to create multiple PDBs, then the unique name you enter here is used as a prefix for all the cloned PDBs, and the suffix is a numeric value that indicates the count of PDBs.

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.

- 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 13-22 displays the Identification page.

Figure 13-22 Cloning a Pluggable Database: Identification Page



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.

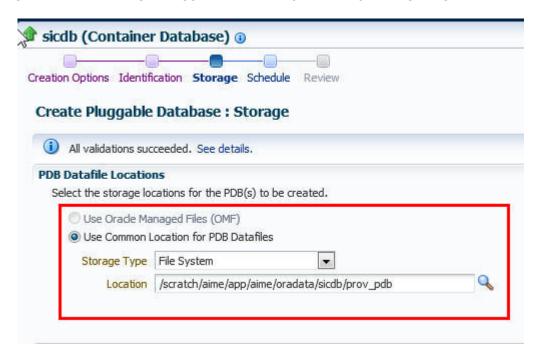
12. In the Storage page, specify the storage information.

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 13-23 displays the Storage page for the Full Clone method.

Figure 13-23 Cloning a Pluggable Database (Full Clone): Storage Page



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 13-24 displays the Storage page for the Snap Clone method.

Figure 13-24 Cloning a Pluggable Database (Snap Clone): Storage Page



- **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.

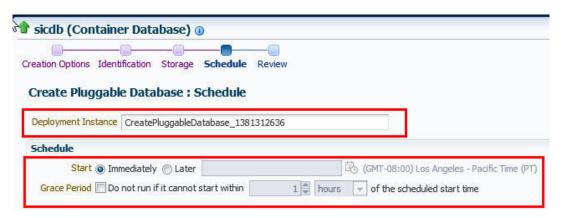
If you want to run the procedure immediately, then retain the default selection, that is, **Immediately.** Otherwise, select **Later** and provide time zone, start date, and start time details.

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 13-25 displays the Schedule page.



Figure 13-25 Cloning a Pluggable Database: Schedule Page



- 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 13-26 displays the Review page.

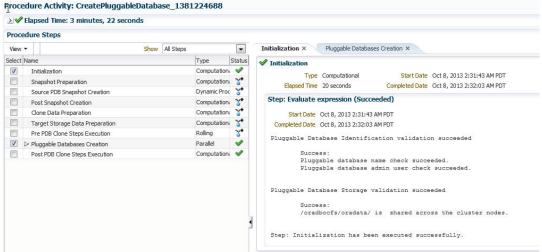
Figure 13-26 Cloning a Pluggable Database: Review Page



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 13-27 displays the Procedure Activity page.

Figure 13-27 Cloning a Pluggable Database: Procedure Activity Page



When you clone a PDB, the Enterprise Manager job system creates a Create Pluggable Databaa job, see Viewing Create Pluggable Database Job Details.



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

Snap Clone REST API

The following are examples for performing Snap Clone operations using REST API

Feature	Description
Request Method	POST
URI	/em/cloud/dbaas/dbplatforminstance/
Request Headers	Authorization: basic < generated credential value > Content-Type: application/json



Feature	Description
Payload	
,	{
	"zone": "/em/cloud/dbaas/zone/
	B5C22C56086C4B6E8A55851841494C6F",
	"name": "Snap Clone via REST
	API",
	"description": "sidb using rest
	api",
	"start date":
	"2023-02-15T16:25:00ZAsia/Calcutta",
	"end date":
	"2023-02-17T17:20:00ZAsia/Calcutta",
	"params":
	<u>.</u>
	"database_sid":"snrst1",
	"username": "oracle",
	"password": "oracle12",
	"service_name":
	"srvc_snrst1",
	"database_profile_version" :
	"Snapshot_02_07_2023_10_04_AM"
	}
	}



"uri" : "/em/cloud/dbaas/

Feature Description

Response

```
dbplatforminstance/byrequest/254",
  "name" : "Snap Clone via REST
API",
  "resource state" :
      "state" : "INITIATED" ,
      "messages" :
            "date" :
"2023-02-15T10:46:45+0000",
            "text" : "Delete Request
ID is : 255",
            "uri" : "/em/cloud/dbaas/
dbplatforminstance/byrequest/255"
      1
    } ,
  "media type" : "application/
oracle.com.cloud.common.DbPlatformIns
tance+json",
  "canonicalLink" : "/em/websvcs/
restful/extws/cloudservices/
service/v0/ssa/em/cloud/dbaas/
dbplatforminstance/byrequest/254",
  "status" : "SCHEDULED"
  "uri" : "/em/cloud/dbaas/
dbplatforminstance/byrequest/254",
  "name" : "Snap Clone via REST
API CREATE 46 44#254",
  "description" :
"DESCRITION NOT AVAILABLE VIEW" ,
  "resource state" :
      "state" : "CREATING" ,
      "messages" :
            "date" :
"2023-02-15T10:48:42+0000",
            "text" : "Delete Request
ID is : 255",
            "uri" : "/em/cloud/dbaas/
dbplatforminstance/byrequest/255"
      ]
    } ,
```



Feature	Description
	"context_id" : "254" ,
	<pre>"media_type" : "application/ oracle.com.cloud.common.DbPlatformIns tance+json" ,</pre>
	"canonicalLink": "/em/websvcs/ restful/extws/cloudservices/ service/v0/ssa/em/cloud/dbaas/ dbplatforminstance/byrequest/254", "status": "IN_PROGRESS", "created": "2023-02-15 10:46:44"

Migrating a Non-CDB as a Pluggable Database Using Enterprise Manager

This section provides information about migrating a non-CDB as a PDB using Enterprise Manager. 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 Enterprise Manager, ensure that you meet the following prerequisites:

- Oracle Software Library must be set up in Enterprise Manager.

 For information on how to set up Software Library in Enterprise Manager.
 - For information on how to set up Software Library in Enterprise Manager, 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 Enterprise Manager target.



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:

 Access from Databases, then select Administration drop down menu and click Database Provisioning. In the Database Provisioning page, in the Related Links section of the left menu pane, click Provision Pluggable Databases.



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.

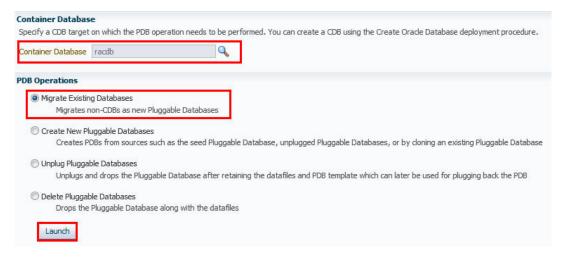


Skip this step if you have accessed the Provision Pluggable Database Console from the CDB's home page.

 In the PDB Operations section of the Provision Pluggable Database page, select the Migrate Existing Databases option and click Launch.

Figure 13-28 displays the Provision Pluggable Database Console.

Figure 13-28 Provision Pluggable Database Console



- 4. On the Database Login page, select the Credential Name from the drop-down list. Click **Login**.
- 5. 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.

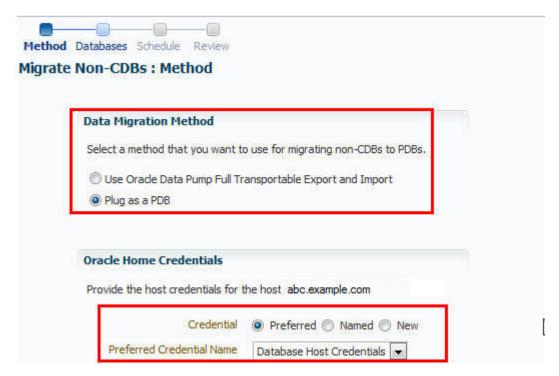
Enter the appropriate credentials for the Oracle Home Credential section.



Click Next.

Figure 13-29 displays the Method page.

Figure 13-29 Migrating a Non-CDB: Method Page



6. 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 13-30 displays how to select the non-CDB you want to migrate.

Figure 13-30 Migrating a Non-CDB: Selecting a Non-CDB



Figure 13-31 displays how to specify the database and database host credentials.

Figure 13-31 Migrating a Non-CDB: Specifying Credentials

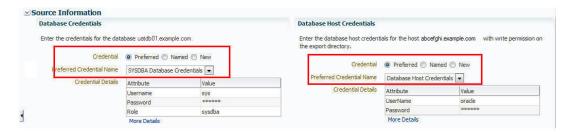


Figure 13-32 displays how to specify the stage location for migration.

Figure 13-32 Migrating a Non-CDB: Specifying Stage Location

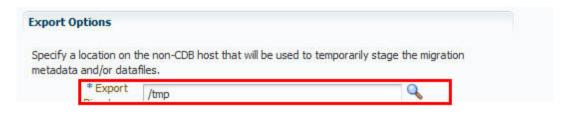


Figure 13-33 displays how to specify the PDB administrator details and datafile location.

Figure 13-33 Migrating a Non-CDB: Specifying PDB Properties and Datafile Location

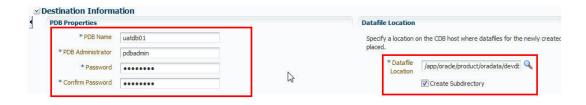


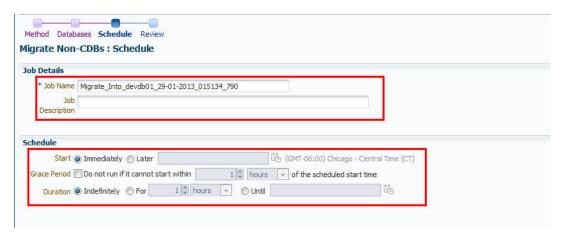
Figure 13-34 displays how to specify the import directory.

Figure 13-34 Migrating a Non-CDB: Specifying Import Directory



On the Schedule page, enter the appropriate job and scheduling details. Click Next.
 Figure 13-35 displays the Schedule page.

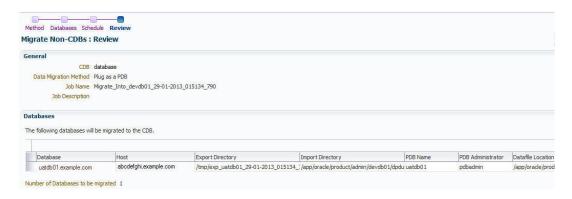
Figure 13-35 Migrating a Non-CDB: Schedule Page



On the Review page, review all details entered. If there are no changes required, click Submit.

Figure 13-36 displays the Review page.

Figure 13-36 Migrating a Non-CDB: Review Page



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.

If you are managing a PDB on OCI resources make sure that all OCI required prerequistes and discovery have been performed. For more information, see: (Optional) Deploying Agents on OCI Resources .



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, 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.
- If the destination CDB has GLOBAL_NAMES initialization parameter set to TRUE and if the user
 has specified existing DB link (when the mechanism of PDB relocation is via DB Link), then
 the name of the DB Link should be same as that of the source CDB.

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.
 - If the destination CDB selected, is Data Guard enabled and there are multiple standby databases, you will be prompted to select a reference standby CDB. This will be used as a reference by the workflow to copy the datafiles to other standby databases.
 - c. In the Credentials section, specify the SYSDBA credentials of the CDB and the CDB host credentials.
 - d. If the source database is TDE enabled and of version 12.2 and above, then enter the wallet password of the destination CDB. If the source database is of a version prior to 12.2, then you must also enter the source wallet password. Note that if the source PDB is TDE enabled, then you can only relocate it to a CDB that is also TDE enabled.
 - e. 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'.



This step is optional. If the existing database link is not selected, then a new database link is created automatically.



- f. 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.
 - Normal Availability 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 Enterprise Manager. In particular, it contains the following:

- Prerequisites for Unplugging and Dropping a Pluggable Database
- Unplugging and Dropping a Pluggable Database





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 Enterprise Manager, ensure that you meet the following prerequisites:

- Oracle Software Library (Software Library) must be set up in Enterprise Manager.
 For information on how to set up Software Library in Enterprise Manager, 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:

 Access from Databases, then select Administration drop down menu and click Database Provisioning. In the Database Provisioning page, in the Related Links section of the left menu pane, click Provision Pluggable Databases.



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.



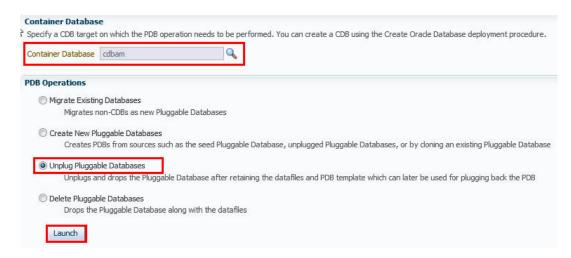
Skip this step if you have accessed the Provision Pluggable Database Console from the CDB's home page.

- In the PDB Operations section, select Unplug Pluggable Database.
- 4. Click Launch.

Figure 13-37 displays the Provision Pluggable Database Console.



Figure 13-37 Provision Pluggable Database Console





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 13-38 displays the Select PDB page.

Figure 13-38 Unplugging a Pluggable Database: Select PDB Page





- 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.



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 DBF 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,

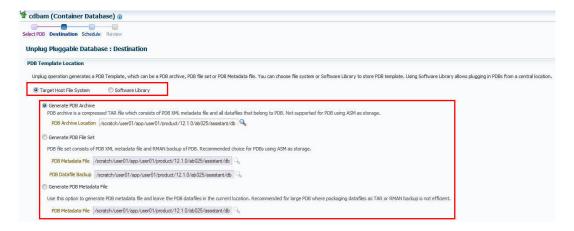
Database Configuration/12.1.0.0.2/unix/Database Templates

 Enter a temporary location where the archive can be created by Enterprise Manager before it is uploaded to the Software Library.

Figure 13-39 displays the Destination page.

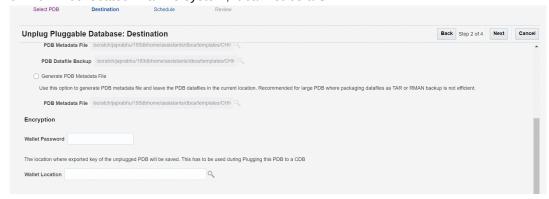


Figure 13-39 Unplugging a Pluggable Database: Destination Page



For TDE enabled PDB

If the PDB is TDE enabled, Enter the Wallet Password, Wallet Transport Secret and the Wallet Location where the Wallet will be stored after Unplug Operation. The wallet location where the stored exported encryption keys will be located cannot be an ASM Disk Group; or if it will be located in a File system. It cannot be a SYMLINK.





This step is applicable only if the PDB Template location is chosen as *Target Host File System* If the PDB Template location is chosen as *Software Library*, enter only the **Wallet Password** for the PDB.

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.

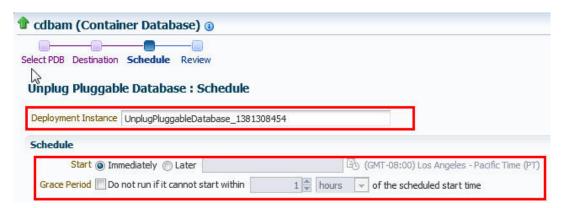
If you want to run the procedure immediately, then retain the default selection, that is, **Immediately.** Otherwise, select **Later** and provide time zone, start date, and start time details.

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 13-40 displays the Schedule page.



Figure 13-40 Unplugging a Pluggable Database: Schedule Page



- 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 13-41 displays the Review page.

Figure 13-41 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 Enterprise Manager.

If you are managing a PDB on OCI resources make sure that all OCI required prerequistes and discovery have been performed. For more information, see: (Optional) Deploying Agents on OCI Resources .

In particular, this section 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 Enterprise Manager, 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 ManagerAdministrator's Guide*.

- Oracle Software Library (Software Library) must be set up in Enterprise Manager.

 For information on how to set up Software Library in Enterprise Manager, 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:

 Access from Databases, then select Administration drop down menu and click Database Provisioning. In the Database Provisioning page, in the Related Links section of the left menu pane, click Provision Pluggable Databases.



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.



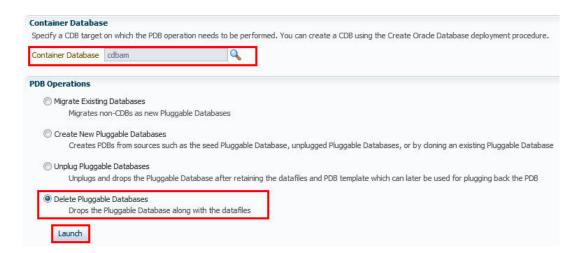
Skip this step if you have accessed the Provision Pluggable Database Console from the CDB's home page.

- 3. In the PDB Operations section, select **Delete Pluggable Databases**.
- 4. Click Launch.

Figure 13-42 displays the Provision Pluggable Database Console.



Figure 13-42 Provision Pluggable Database Console





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.

In the Select PDBs page of the Delete Pluggable Databases Wizard, click Add. Select the PDBs that you want to delete, then click Select.



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 13-43 displays the CDB Host Credentials section of the Select PDBs page.

Figure 13-43 Deleting Pluggable Databases: Specifying Credentials





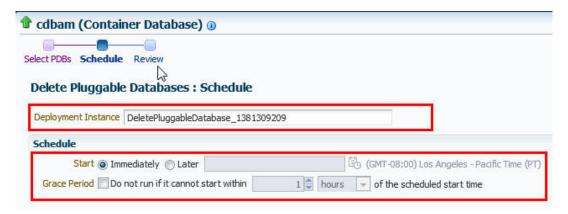
7. In the Schedule page, enter a unique deployment procedure instance name and a schedule for the deployment. The instance name you enter here helps you identify and track the progress of this procedure on the Procedure Activity page.

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 13-44 displays the Schedule page.

Figure 13-44 Deleting Pluggable Databases: Schedule Page

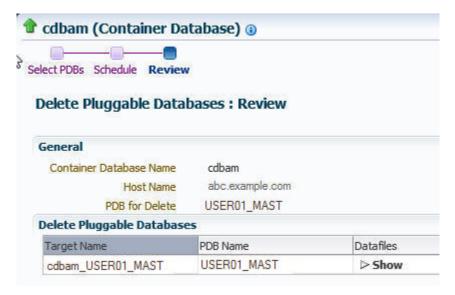


- 8. Click Next.
- In the Review page, review the details you have provided for the deployment procedure. If you are satisfied with the details, click Submit.

If you want to modify the details, then click **Back** repeatedly to reach the page where you want to make the changes.

Figure 13-45 displays the Review page.

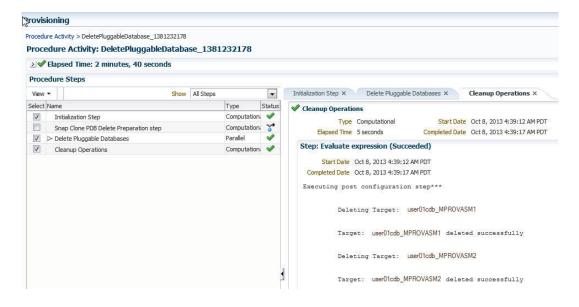
Figure 13-45 Deleting Pluggable Databases: Review Page



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 13-46 displays the Procedure Activity page.

Figure 13-46 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

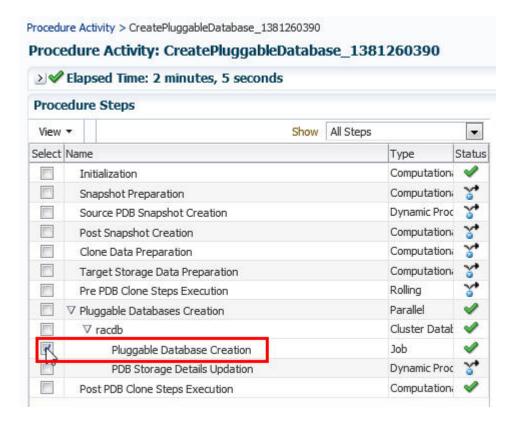
Viewing Create Pluggable Database Job Details

To view the details of a create PDB job, follow these steps:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and then select Procedure Activity.
- Click the deployment procedure that contains the required create PDB job.



3. Expand the deployment procedure steps. Select the PDB creation job.





4. Click Job Summary.



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.

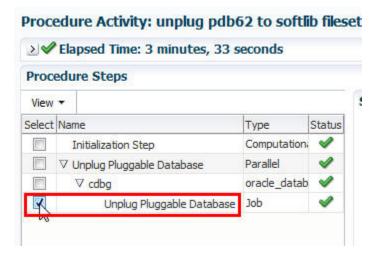
Viewing Unplug Pluggable Database Job Details

To view the details of an unplug PDB job, follow these steps:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and 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.

6. 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

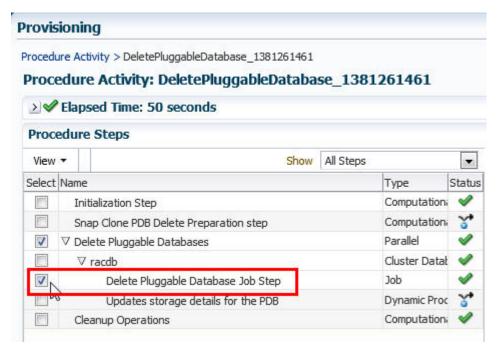
To view the details of a delete PDB job, follow these steps:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and then select Procedure Activity.
- 2. Click the deployment procedure that contains the required delete PDB job.

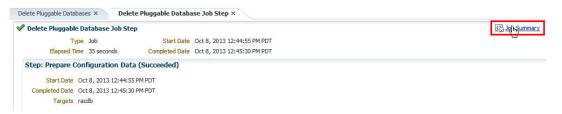


3. Expand the deployment procedure steps. Select the delete PDB job.





Click 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 Enterprise Manager. 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:

- 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.
- The page will update to show the System Queues.
- Click the context switcher to display the drop-down menu. If the menu does not show the PDBs you want, then select All Containers.
- 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 Database menu, select Control, then select Open/Close Pluggable Database.
- From the Open/Close PDB page, select a PDB from the list.
- 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.
- 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:

- 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.
- Once state change completes, the Open/Close PDB page will update to show the new state of the PDB.



14

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

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:

- 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.
- 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. If the source database is Transparent Data Encryption (TDE) enabled, then enter the wallet password and wallet transport secret.
 - 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:



Note:

The database clone wizard only supports the <code>{ORACLE_BASE}</code> variable for path simplification.

Example for a storage location:

{ORACLE BASE}/oradata

- 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.
- 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. Starting with Enterprise Manager 13c Release 5 Update 8 (13.5.0.8) you can specify a non-DBSNMP user for database monitoring, edit the DBSNMP field with the user name. If the user does not exist it will be created.
- If the source database is TDE enabled, then indicate if the wallet password should be
 the same for the cloned database by selecting the Use same wallet password as
 source check box and specifying the password.
 If the source database is not TDE enabled, then you can opt to configure wallet at
 destination by selecting the corresponding check box and specifying a password.
- 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.





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, see Using Custom Scripts Stored in the Software Library in *Oracle Enterprise Manager Middleware Lifecycle Management Administrator's Guide*.

You can view sample scripts in the Database as a Service/Samples/Pre or Post Database Script folder in the 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.

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

Database Cloning Using EM CLI

To create a full clone of a database you will need to execute the db_clone_management verb. See the following verb sample with parameters:

```
emcli db_clone_management
-createFullClone
-target_name="<DB_target_name>"
-target type="<oracle database/rac database>"
```



```
-clone_type="<LIVE/POINT_IN_TIME>"
-input file="<Input file path>"
```

Furthermore you can view the different parameters that need to be included in an input file for a live clone, or a point in time clone; and also the different parameters needed when cloning a Single Instance Database or a RAC Database Clone. To view this information you will need to run the db clone management with the -print properties flag:

```
emcli db_clone_management
-createFullClone
-target_name="<DB_target_name>"
-target_type="<oracle_database/rac_database>"
-clone_type="<LIVE/POINT_IN_TIME>"
-print properties
```



If the source database is TDE enabled, then you must also specify the SRC_WALLET_PWD, WALLET_SOURCE_SECRET and DEST_WALLET_PWD properties.

Database Cloning Using Pre and Post Scripts

The following is a list of accepted Deployment Procedure Variables that can be incorporated into Pre and Post scripts for Database cloning workflows. For more information on Pre and Post scripts see: Custom Pre and Post Scripts for Fleet Operations

Variable Name	Description
SRC_DB_TARGET_NAME	Source database target name
SRC_DB_TARGET_TYPE	Source database target type
COMMON_GLOBAL_DB_NAME	Global database name of clone
COMMON_DB_SID	Clone database SID



Use COMMON_DB_SID to obtain the Database Unique Name of a cloned database. Databases cloned via the workflow will have the Database Unique Name equal to the value of its SID.

DB_TARGET_NAME	Name of the clone database target
DB_TARGET_DISPLAY_NAME	Display name of the clone database target
DATABASE_TYPE	Name of the clone database. Values: dbTypeSI or dbTypeRAC



Variable Name	Description
TARGET_HOST_LIST	Destination host name
ORACLE_HOME_LOC	Destination Oracle Home location
ORACLE_BASE_LOC	Destination Oracle Base location
DB_STORAGE_TYPE	Destination storage type. Values: FS or ASM
DB_FILE_LOC	Destination database files storage location
FLASH_REC_AREA	Location at destination where recovery related files will be created
FLASH_REC_AREA_SIZE	Fast Recovery area size
ARCHIVE_LOG_MODE	Indicates if the clone database is in ARCHIVELOG mode
LISTENER_PORT	Listener port with which the clone database is configured
DB_TEMPLATE_STAGE	Temporary work directory at the destination
ENABLE_DATA_MASK	Indicates if data masking definition is applied on clone database

Sample Variables:

```
COMMON GLOBAL DB NAME=CLONEDB
COMMON DB SID=CLONEDB
DATABASE TYPE=dbTypeSI
DB_STORAGE_TYPE=FS
DB FILE LOC=/scratch/app/oradata
FLASH_REC_AREA=/scratch/user/app/fra
FLASH REC AREA SIZE=3780
ARCHIVE LOG MODE=YES
LISTENER PORT=<Port number>
DB TEMPLATE STAGE=/tmp
SRC DB TARGET NAME=xyz.example.com
SRC DB TARGET TYPE=oracle database
ORACLE BASE LOC=/scratch/app/base
ORACLE HOME LOC=/scratch/app/product/<DB Version>/dbhome 1
ORACLE_HOME_NAME=OraDBHome_Example
DB TARGET NAME=clonedb.xyz.example.com
TARGET_HOST_LIST=desthost.example.com
CONFIGURE WITH ORACLE RESTART=<Y/N>
FRA STORAGE TYPE=FS
DEGREE OF PARALLELISM=8
ENABLE DATA MASK=<True/False>
DB_TARGET_DISPLAY_NAME=SAMPLE_CloneDB
```

Verify Status of a Database Clone

To verify the status of the database clone creation, execute the verb emcli get instance status -instance={instance GUID}.



Refresh a Cloned Database

Prerequisites for Refreshing a Cloned Database

- The cloned database or Test Master should have been created using Enterprise Manager's deployment procedure. Procedure execution must have successfully completed.
- SSA plugin must be installed in Enterprise Manager and the *OracleCloud for Database* license pack must be enabled for the database target.
- If a database was marked as clone using the Add button in the Clone Management dashboard the refresh feature will not be available for said database.
- If a Test Master database is created on an Exadata machine (for Exadata Sparse Clones), the refresh feature will not be available for said database.

Refreshing a Cloned Database

Once the Enterprise Manager Plugins are deployed, the 'Refresh' option is enabled. To deploy the plugins perform the following steps:

- 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.
- Click Cloning Management, and then select a cloned database from the list.
- Click Refresh.

Cloning an Oracle Database Using Recovery Manager (RMAN) Backup

To clone an Oracle database using RMAN backup, follow these steps:

- From the Targets menu, select Databases.
- 2. On the Databases page, select a database that you want to clone.
- On the Database target page, from the Oracle Database menu, select Cloning and then click Clone Management.
- 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.
- 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.

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.

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.**

Creating a Full Clone Pluggable Database

There are two methods to create a Full Clone Pluggable Database, you can use the Clone Wizard in Enterprise Manager or via command line with EM CLI.

- Creating a Full Clone Pluggable Database Using the Clone Wizard
- Creating a Full Clone Pluggable Database Using EMCLI

If you are managing a PDB on OCI resources make sure that all OCI required prerequistes and discovery have been performed. For more information, see: (Optional) Deploying Agents on OCI Resources.



From Enterprise Manager 13.4 Release Update 3 onwards users will have the ability to clone and upgrade Pluggable Databases (PDB) by plugging a PDB into a higher version Container Database (CDB), the PDB will be automatically upgraded to the higher CDB version.

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**.
- 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.

Note:

If the destination CDB is a of higher version than the PDB clone, then the cloned PDB will be upgraded to the CDB version as part of the deployment procedure. The pre-upgrade SQL script will be executed on the source PDB. Before cloning a PDB version 12.1.0.2.0 to a higher CDB version using the Database Link Mechanism make sure you have deployed patch 18633374 as detailed in MOS NOTE 2090019.1.

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.



If the source database is TDE enabled and of version 12.2 and above, then enter the wallet password and Wallet Transport Secret of the destination CDB. If the source database is of a version prior to 12.2, then you must also enter the source wallet password. Note that if the source PDB is TDE enabled, then you can only clone it to a CDB that is also TDE enabled.

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 data files 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.

- 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 EMCLI

Pluggable Database Cloning using EMCLI

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.

• If the source database is TDE enabled and of version 12.2 and above, then you must also specify the <code>DEST_WALLET_PWD</code> property. If the source database is TDE enabled and of a version prior to 12.2, then you must specify the SRC WALLET PWD, WALLET TRANSPORT SECRET and DEST WALLET PWD properties.



Pluggable Database Cloning Using Pre and Post Scripts

The following is a list of accepted Deployment Procedure Variables that can be incorporated into Pre and Post scripts for Pluggable Database cloning workflows. For more information on Pre and Post scripts see: Custom Pre and Post Scripts for Fleet Operations

Variable Name	Description
SRC_HOST	Source host name
SRC_CDB_TARGET	Source CDB target name
SRC_CDB_TYPE	Source CDB target type. Values: oracle_database or rac_database
SRC_CDB_NAME	Source CDB name
SRC_CDB_SID	Source CDB SID
SRC_PDB_NAME	Source PDB name
SRC_DATABASE_VERSION	Source CDB version
SRC_ORACLE_HOME	Source Oracle Home location
DEST_HOST	Destination host name
DEST_CDB_TARGET	Destination CDB target name
DEST_CDB_TYPE	Destination CDB target type. Values: oracle_database or rac_database
DEST_PDB_NAME	Destination PDB name (clone)
DEST_CDB_NAME	Destination CDB name
DEST_CDB_SID	Destination CDB SID
DEST_PDB_DISPLAY_NAME	Display name of the clone PDB target
DEST_PDB_ADMIN_NAME	Destination PDB administrator
DEST_DATABASE_VERSION	Destination CDB version
CREATE_FILE_DEST	Destination PDB datafiles storage location
DEST_WORK_DIR	Work directory on the destination host
DEST_ORACLE_HOME	Destination Oracle Home location
LOGGING_TYPE	Destination PDB logging clause
SRC_PDB_ORIG_STATE	Original state of the Source PDB

Sample Variables:

```
SRC_HOST=xyz.example.com

SRC_CDB_TARGET=SAMPLE_CDB1

SRC_CDB_TYPE=oracle_database

SRC_PDB_TARGET=SAMPLE_PDB1

SRC_CDB_NAME=SAMPLE_CDB1

SRC_CDB_SID=CDB1

SRC_DB_SID=CDB1

SRC_DATABASE_VERSION=<DB_Version>

SRC_ORACLE_HOME=/scratch/app/product/<DB_Version>/dbhome_1

DEST_HOST=abc.example.com

DEST_CDB_TARGET=SAMPLE_CDB2

DEST_CDB_TYPE=oracle_database

DEST_CDB_NAME=SAMPLE_PDB2

DEST_CDB_NAME=CDB2

DEST_CDB_NAME=CDB2

DEST_CDB_SID=CDB2
```



```
DEST_PDB_DISPLAY_NAME=SAMPLE_Clone2_PDB
DEST_PDB_ADMIN_NAME=PDBADMIN
DEST_DATABASE_VERSION=<DB Version>
DEST_LOCATION=/scratch/app/oradata/EMEA_CL2
DEST_WORK_DIR=/tmp
DEST_ORACLE_HOME=/scratch/appk/product/<DB Version>/dbhome_1
CREATE_FILE_DEST=/scratch/app/oradata/SAMPLE_PDB2
LOGGING_TYPE=LOGGING
SRC_PDB_ORIG_STATE=READ_WRITE
```

Relocation Database Cloning Using Pre and Post Scripts

The following is a list of accepted Deployment Procedure Variables that can be incorporated into Pre and Post scripts for Pluggable Database relocation cloning workflows. For more information on Pre and Post scripts see: Custom Pre and Post Scripts for Fleet Operations

Variable Name	Description
SRC_PDB_TARGET_NAME	Source PDB target name
SRC_CDB_TARGET_NAME	Source CDB target name
SRC_CDB_TARGET_TYPE	Source CDB target type. Values: oracle_database or rac_database
SRC_PDB_NAME	Source PDB name
SRC_PDB_MODE	State of the Source PDB
SRC_ORACLE_HOME_LOC	Source Oracle Home location
SRC_CDB_VERSION	Source CDB version
RELOCATE_TYPE	Mechanism of PDB relocation
AVAILABILITY	Availability mode
DEST_PDB_NAME	Destination PDB name
DEST_PDB_DISPLAY_NAME	Display name of the relocated PDB target
DEST_CDB_TARGET_NAME	Destination CDB target name
DEST_CDB_TARGET_TYPE	Destination CDB target type. Values: oracle_database or rac_database
DB_LINK_NAME	Name of the Database Link used for PDB relocation
DEST_WORK_DIR	Work directory on the destination host
STORAGE_LOCATION	Destination PDB datafiles storage location
LOGGING_TYPE	Destination PDB logging clause
DEST_HOST_NAME	Destination host name
DEST_ORACLE_HOME_LOC	Destination Oracle Home location
DEST_CDB_VERSION	Destination CDB version

Sample Variables:

SRC_PDB_TARGET_NAME=PDB1_SAMPLE
SRC_CDB_TARGET_NAME=CDB1_SAMPLE
SRC_CDB_TARGET_TYPE=oracle_database
RELOCATE_TYPE=USING_DBLINK
AVAILABILITY=NORMAL
DEST_PDB_NAME=PDB2_SAMPLE
DEST_PDB_DISPLAY_NAME=PDB2_SAMPLE



DEST_CDB_TARGET_NAME=CDB2_SAMPLE

DEST_CDB_TARGET_TYPE=oracle_database

DB_LINK_NAME=LINK<LINK NUMBER>

DEST_WORK_DIR=/tmp/BKP_STAGE

STORAGE_LOCATION=/scratch/app/oradata

LOGGING_TYPE=LOGGING

SRC_PDB_NAME=PDB1_SAMPLE

SRC_PDB_MODE=READ_WRITE

SRC_ORACLE_HOME_LOC=/scratch/app/product/<DB_VERSION>/dbhome_1

SRC_CDB_VERSION=<DB_VERSION>

DEST_HOST_NAME=abc.sample.com

DEST_ORACLE_HOME_LOC=/s



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.



Enterprise 🔻 <u>Targets</u> ▼ <u>Favorites</u> ▼ Summary Auto Refr Monitoring Reports Configuration Compliance Provisioning and Patching Procedure Activity Quality Management Procedure Library My Oracle Support Software Library Cloud Bare Metal Provisioning Chargeback Database Provisioning Consolidation Middleware Provisioning Patches & Updates Saved Patches Linux Patching ORACLE Enterprise Manager Cloud Control 13c * · O. SYSMAN ▼ Database Provisioning Page Refreshed Nov 22, 2017 5:29:19 AM EST (*) **Profiles** ∡ Setup Logged in as Designer Profile Name Description Profile Status Software Library 🗸 Done This is a reference profile created for provisioning rac database using ASM for 11.2.0.3.0 on Windows RAC Database using ASM for 11.2.0.3.0 on Windows This is a reference profile created for provisioning rac database using File System for 12.1.0.1.0 on Unix RAC Database using File System for 12.1.0.1.0 on Unix RAC One Node using ASM for 11.2.0.3.0 on Windows Notifications Salus This is a reference profile created for provisioning database 11.2.0.3.0 using file system on unix Single Instance using File System for 11.2.0.3.0 on Unix Add Host Privilege Delegation Deployment Procedures ∠ Current Status 🗵 Launch 📗 Create Like... 📝 Edit 🐇 Edit Permissions 💢 Delete 🗷 Launch Upgradability Report Procedure Name Procedure Activity Provision Oracle Database ORACLE This procedure provisions the Oracle Orid Infrastructure for Standalone Server and Oracle ... Data Masking for Fusion SaaS POD Data Masking deployment procedure to run ADM Verify. Generate Masking Script and Run ... ORACLE Illi Uodate database 70 CLOU Procedure to upgrade Single Instance Database, Real Application Clusters, Grid Infrastruct... May Concepts Provision Oracle RAC Database This procedure provisions Oracle Grid Infrastructure and Oracle Real Application Clusters f... ORACLE Designer vs Operator Extend/Scale up Oracle Real Application Clusters This procedure will extend and configure an existing Oracle Clusterware and Oracle Databa. ORACLE Provisioning Profiles Database Gold Images Provision Cracle Clusterware / RAC for UNIX and RD(This procedure assists in installing/cloning and configuring a cluster database (a Real Appli... ORACLE

Figure 15-1 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 15-1 Predefined Roles for Designers

Roles	Description
EM_PROVISIONING_DESIGNE R	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.



For information about EM PROVISIONING DESIGNER see, Table 3-3.

• 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 15-2 Predefined Roles for Operators

Roles	Description
EM_ALL_VIEWER	Role Desc
EM_HOST_DISCOVERY_OPERA TOR	Role has privileges to execute host discovery
EM_PROVISIONING_OPERATO R	Role has privileges for provisioning operator
EM_TARGET_DISCOVERY_OPE RATOR	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.



For information about EM PROVISIONING OPERATOR see, Table 3-3.

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 Setting Up the Software Library in *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:



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.





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:
 - To disable a phase or step, select the phase or step you want to disable, and click Disable.
 - 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 selction 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:

- 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**.
- 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 <u>Target List</u>.
 - The advantage of this approach is that you can have multiple custom target lists, and assign it to the different phases in your procedure.
- 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:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and then select Procedure Activity.
- 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.



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 15-3 Field Description - Adding Rolling Phase

Field Name	Description
Select	Select Phase.



Table 15-3 (Cont.) Field Description - Adding Rolling Phase

Field Name	Description
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 name="" or="" step=""></phase> . To insert it inside the phase or step selected, select Inside<phase or="" step=""></phase> , Otherwise, select Before <phase or="" step=""></phase> .
Туре	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:

- On the Create page, specify general information about the step as described in Table 15-4.
- 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.
- 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 15-4.
- 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 15-4.
- 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.
- 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 15-4.
- 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.
- 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 15-4.
- 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 15-4.
- 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 15-4 Field Description - Adding Steps

Field Name	Description	
Select	Select Step.	
Name	Specify a name for the custom step.	
Description	Provide a description for the custom step.	
Condition	Leave this field blank.	
Insert Location	If you want to insert the custom step after the step you selected, then select After <step name="">. Otherwise, select Before <step>.</step></step>	
Туре	 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:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and then select Procedure Activity.
- 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.**

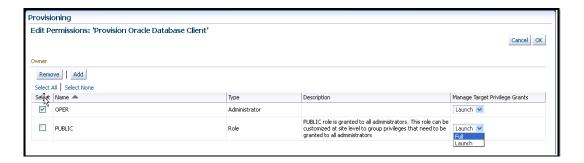
Editing and Saving Permissions of a Procedure

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. Access from **Databases**, then select **Administration** drop down, then select **Database Provisioning**, and then select **Procedure Activity**.
- 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.





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, access from **Databases**, then select **Administration** drop down, then select **Database Provisioning**, and then select **Procedure Activity.**.
- 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 15-5 Deployment Procedure Status

Status	Description
Scheduled	The Deployment Procedure is scheduled to be executed at the date and time that you have specified.
Running	The Deployment Procedure is currently being executed.
Action Required	The Deployment Procedure has stopped running as user interaction is required.
Suspended	The Deployment Procedure has been temporarily stopped from execution. You can resume from the paused step later.
Failed	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	The Deployment Procedure has successfully completed its execution.
Skipped	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	The Deployment Procedure has been permanently stopped from execution by the user.
Saved	The Deployment Procedure has not been submitted for execution, and has been saved.
Completed with Errors	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. You can filter by Search Text Fields, Type, and Status. Starting with Oracle Enterprise Manager 13.5 Release Update 17 (13.5.017) additional search filters Owner, Start Date From, and Start Date To, have been added.

- **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. Access from **Databases**, then select **Administration** drop down, then select **Database Provisioning**, and then select **Procedure Activity.**
- 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.
- 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:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and then select Procedure Activity.
- 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.



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.



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



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:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and then select Procedure Activity.
- 2. On the Provisioning page, from the list of actions, select **Create New**, and click **Go**.
- 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.
- 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.



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. Access from **Databases**, then select **Administration** drop down, then select **Database Provisioning**, and then select **Procedure Activity**.
- 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.**

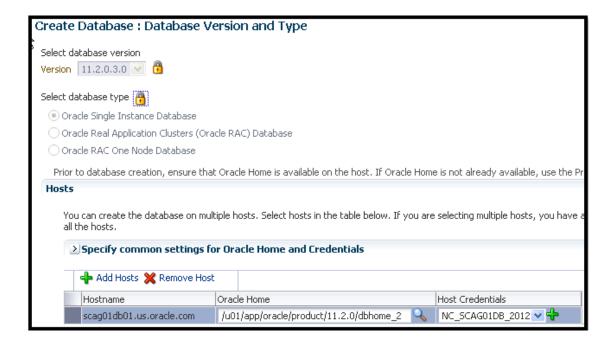
- 5. On the Set Credentials page, you need to set the credentials for the target host machines that you have included as a part of the host_target_list variable. Click **Next.**
- On the Set Schedule and Notification page, you can schedule the job to run immediately or at a later preferred time.
- 7. 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:





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

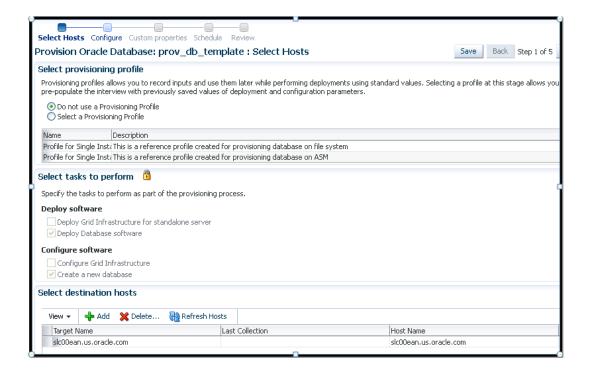
Step 1: Saving a Single Instance Database Deployment Procedure with Lock Downs

In the following section user logs in to Cloud Control as a designer (ATBARBOZ1), and provisions a Single Instance Database with lock downs as follows:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and then select Procedure Activity.
- 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.

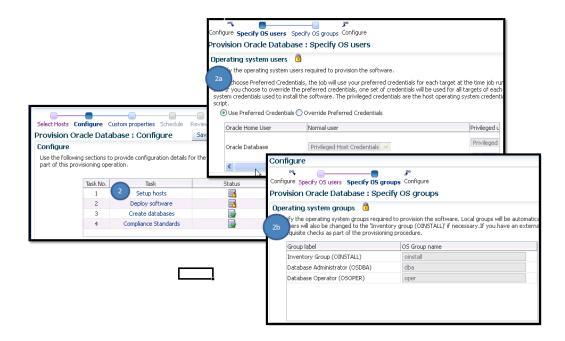
In the Select Tasks to Perform section, do the following and lock the values:

- 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.



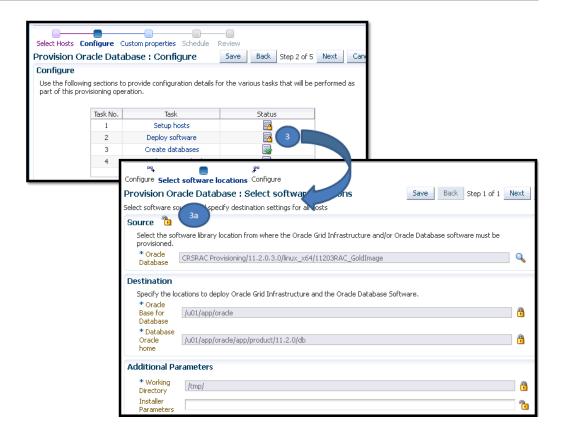


- 4. On the Configure page, the following configuration options appear:
 - On the Configure page, click Setup Hosts. On the Operating System Users page, specify the operating system user for the Oracle Home for the database. For Oracle Home User for the database, select the Normal User and Privileged 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.



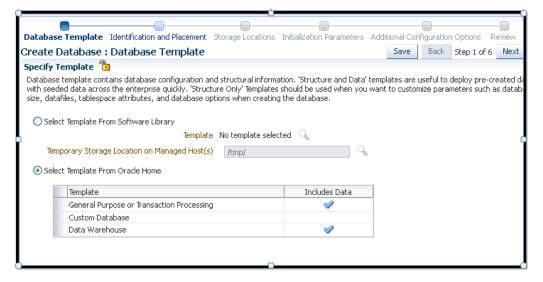
Click **Next** to come back to the Configure page.

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.

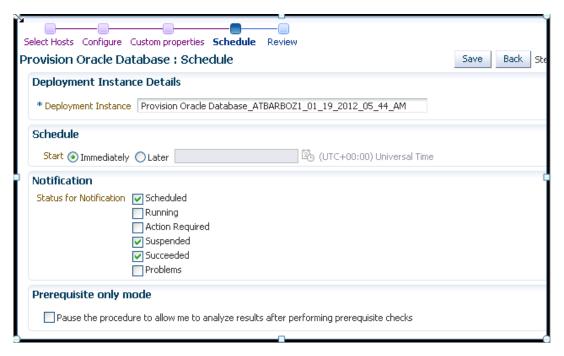
On the Configure page, click Create Databases, the following screen appears. Update
only the mandatory fields, and click Next to proceed. Do no lock any of the values in
this wizard:



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.



6. 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.



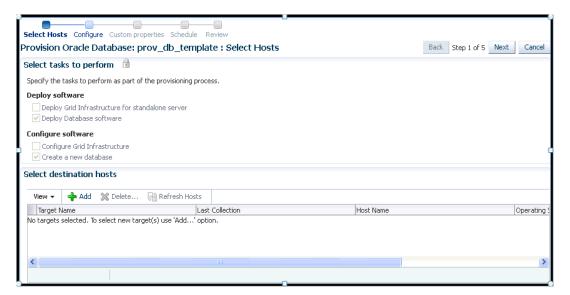
Step 2: Launching the Saved Single Instance Database Deployment Procedure

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.

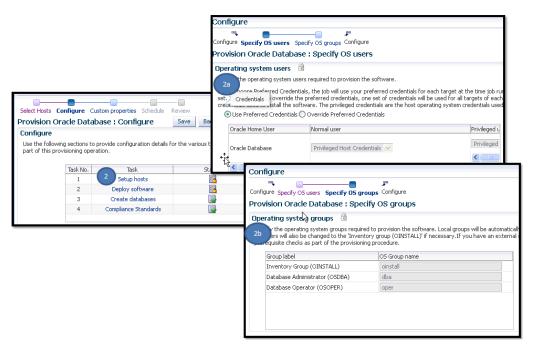
- 1. Access from **Databases**, then select **Administration** drop down, then select **Database Provisioning**, and then select **Procedure Activity.**
- 2. In the Procedure Library, select a procedure prov_db_template, and click Launch.



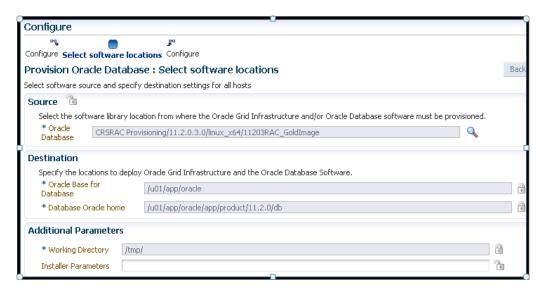
3. On the Select Hosts page, in the Select hosts section, click **Add** to select the destination host where you want to deploy and configure the software, and then click **Next.**



- 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.

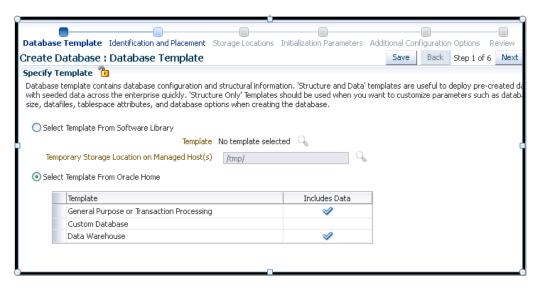


 On the Configure page, click **Deploy Software**. Since the values here are locked by the designer, you will not be able to edit them. Click **Next** to come back to the Configure page.



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.



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.



- 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.
- 6. 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:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and then select Procedure Activity.
- 2. On the Provisioning page, select the saved UDDP configuration template that you saved as a part of the previous step, and click **Launch**.



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.

- 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.**
- 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:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and then 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 Dump
- Viewing the Execution Time of a Deployment Procedure
- Searching for a Step
- Downloading a Step Output
- Accessing the Job Summary Page





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 15-6 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	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

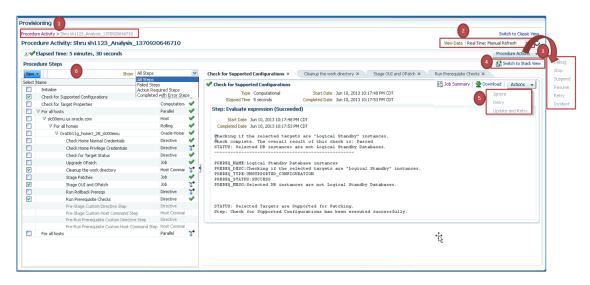
The Procedure Activity page primarily helps you track the execution of the procedure instance submitted. Click any instance running to view the Procedure Instance Execution page which is broadly divided into the following regions as shown in Figure 15-2:



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 15-2.
- 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 15-2 Procedure Execution Page



Note:

For information about the tasks that can be performed from the Procedure Instance Execution Page, see Procedure Instance Execution Page .

- Breadcrumb Trail
- 2. View Data
- 3. Procedure Actions Section
- 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:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and then select Procedure Activity.
- 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:

 Access from Databases, then select Administration drop down, then select Database Provisioning, and then select Procedure Activity.

- 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 Dump

To create an incident dump for a procedure execution, follow these steps outlined below. To create an Incident Dump for procedure steps or dynamic procedure steps where an inner incident execution ID is required follow the additional steps outlined in step 6.

- Access from Databases, then select Administration drop down, then select Database Provisioning, and then select Procedure Activity.
- On the Procedure Activity page, click the Deployment Procedure Execution to select the procedure.
- 3. Make a note of the Execution ID.
- 4. In the Procedure Instance Execution page, from the **Procedure Actions** menu, select **Incident**.
- 5. 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*.
- 6. Check the failed step:
 - If this not a procedure step: If the failed step is not a procedure step or dynamic procedure step, no further action is required. The incident dump created will contain all pertinent information.
 - If this is a procedure step:
 - Identify the Execution ID of the failed step (Step 3), this will be your outerDPExecutionGuid.
 - Run the following EM CLI command to obtain the deployment procedure execution details in XML format

```
emcli get_instance_status -exec=<outerDPExecutionGuid> -xml -details
```

- c. Check the failed deployment procedure inner step in the XML file and make note of the innerDPExecutionGuid ID.
- d. Create an incident dump for the inner procedure step using EM CTL:

```
emctl paf create incident -exec <innerDPExecutionGuid>
```

Entering the innerDPExecutionGuid from step c.

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.

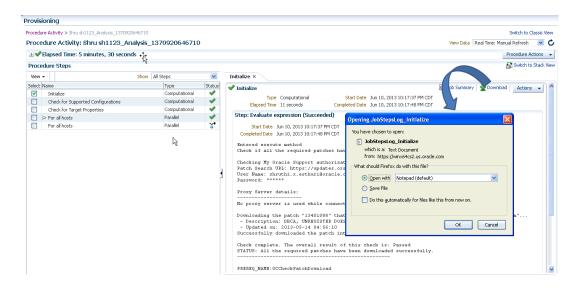
Searching for a Step

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:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and 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:

- Access from Databases, then select Administration drop down, then select Database Provisioning, and then select Procedure Activity.
- 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.

Part III

Database Maintenance: Updates (Patching) and Upgrades

Updates (patching) and Upgrades are some of the most important activities that need to be regularly performed as part of the database lifecycle management activities. With Enterprise Manager you can manage updates and upgrades for Oracle databases.

This section covers the following topics:

- Database Maintenance Concepts
- · Getting Started with Updates (Patching) and Upgrades
- Database Updates and Upgrades Using Fleet Maintenance
- Patching End-to-End Use Cases
- Upgrade databases and Grid Infrastructure Using Fleet Maintenance



Database Maintenance Concepts

Patching is an important phase of the product lifecycle, enabling you to keep your software updated with bug fixes. Oracle releases several types of patches periodically helping maintain your product.

This chapter describes how Oracle Enterprise Manager Fleet Maintenance solution addresses patch management challenges and covers the following:

- Patch Management Overview
- · Understanding the Patching Workflow

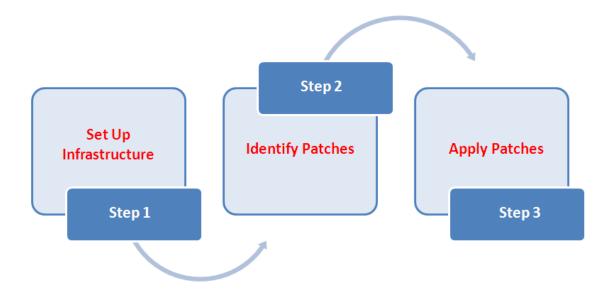
Patch Management Overview

This section describes the following:

- Downloading Patch Recommendations and Patches
- Understanding the Patching Workflow

Understanding the Patching Workflow

The following illustration describes the overall patching workflow offered within the Enterprise Manager console.



Step	Step Name	Description	Reference Links
Step 1	Set Up Infrastructure	Meet the prerequisites and set up the infrastructure for rolling out patches. Create administrator roles for patching, fulfill mandatory and optional prerequisites, and make online or offline patching settings.	Before you begin: Meeting the Prerequisites
Step 2	Identify the Patches	View the recommendations made by Oracle on patches to be applied, identifying the ones to apply.	Downloading Patch Recommendations and Patches
Step 3	Apply Patches	Apply patches using Fleet Maintenance in the Enterprise Manager user interface or EM CLI.	Database Updates and Upgrades Using Fleet Maintenance



17

Getting Started with Updates (Patching) and Upgrades

This chapter describes the actions and activities required to get started with updating and upgrading Oracle databases with Fleet Maintenance.

This section cover the following:

- Overview of Updates and Upgrades
- Before you begin: Meeting the Prerequisites

Overview of Updates and Upgrades

Your environment encounters challenges when trying to maintain databases at their latest version, leading to a dilemma where you either postpone your patching cycles or neglect to patch the databases altogether. This situation significantly heightens your vulnerability to potential threats.

Unpatched Systems	Misconfiguration and Sprawl	Compromised Administrative Privileges
High risk of breaches	45% of breaches were due to misconfiguration	74% of breaches leveraged weak credentials
21% of breaches are due to unpatched systems even though patches were available, but not applied	Preferred way to exploit are misconfigurations, insecure configuration changes, and sprawl. Home grown scripts increase vulnerability and maintenance cost	Lack of security policies with principles of least privileges to users contributes to breaches and security incidents

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.

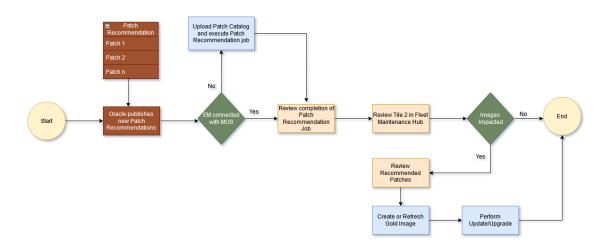


Figure 17-1 Fleet Maintenance Operations Lifecycle Overview

The video Patch Oracle Databases Using Oracle Enterprise Manager Fleet Maintenance, gives an overview of Oracle Database patching and features examples for single instance database patching: Video

About Database Fleet Maintenance

Database Fleet Maintenance allows administrators to maintain sets 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 Monthly Recommended Patches (MRPs).

Any new database patches and updates are made available in the form of images. A Gold image represents the software binary that is patched to the latest required level. Each newly created image from a new patch for a specific database version is a new version.

For example, a user may create an Image with the first version containing 19.18RU. Later, they may create a new version containing 19.18RU+MRP1. Then they can create a new version containing 19.19RU and so on. Once they create a version and they decide to roll it out, they should mark it as Current. This version is now the gold image for the databases subscribed to this target.

There are three main variations of patching:

- 1. Patch a specific (non-container) Oracle Database.
- Patch a Container Database (CDB) and it's Pluggable Databases (PDB) together.
- Patch PDBs independently by relocating them to a new CDB running from a patched Oracle home.

Database fleet maintenance can be used for:

Single instance, RAC, and RAC One Node databases.



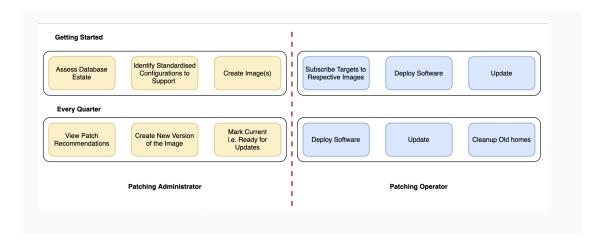


Support is limited to: Single instance database on Windows, Standby databases (Single instance and RAC), Grid Infrastructure for Clusters, Oracle Restart (Grid Infrastructure for Standalone Server), and associated databases.

- 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:

Figure 17-2 Fleet Maintenance Roles



- Identify configuration pollution using the Software Standardization Advisor (this is a one time activity).
- Create reference environments, this is a prerequisite for Gold Image creation.
- Create the Gold Images using the reference environments and versioning them for additional changes. This allows the reference environment to be patched with the desired set of patches to be rolled to the target databases.
- Subscribe databases and clusters to a gold image on which the appropriate patches have been applied. This is a one time activity, once subscribed this task is not required for rolling out the next update.
- Deploy the image and switching 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.

Fleet Maintenance Hub

Fleet Maintenance Hub, helps ease automated updates (patching), and upgrades of your database fleet. It also provides automated patch recommendations for individual images. This allows administrators to have an accurate list of patches that are missing and must be applied on a newer version to make images up to date.

Fleet Maintenance hub is split into three main tiles allowing for an easy identification of the required activity you as an administrator will need to perform. The three tile sections are: Targets Not Subscribed, Patch Recommendations For Images, and Patch Compliance with Images, each meant to perform a specific administrator activity.

To launch Fleet Maintenance Hub, navigate to Targets, then select Databases, select Administrator and click on **Fleet Maintenance Hub**.

Additionaly you can access Fleet Maintenance directly from the main Enterprise Manager Cloud Control page navigate to: **Target** then, **Databases**, then **Administration**, and click on **Fleet Maintenance**, here you will see the Getting Started screen. This will show you the basics of Fleet Maintenance, with direct links to points of interest within this book to get you started with Fleet Maintenance.

Terminology Used in Fleet Maintenance

- **Database Release**: It denotes a major release. The first numerical value in the Oracle Database Release Number denotes the release. Eg: 12c, 19c, 23c. Refer to release number section in Database Upgrade Guide for further information.
- Database Version: Any update to a major release is denoted by its second numerical value. Eg: 19.10, 19.20 or 19.22
- Gold Image: Represents a database release. A Gold Image consists of versions which are mapped to the Database version. When you create a Gold Image (sometimes referred as image), you create it with a version. You cannot have an empty image. Its is advisable that you create only one Gold Image for a particular database release. When Oracle Releases new database versions, you add on new versions to a particular Image. Eg: For a 19c Release, you can create a gold image named 19c_Release. Then you may add new versions like 1910DBRU, 1915DBRU and 1922DBRU. Oracle recommends that in a particular Gold Image, you keep only 3 versions, this will help you ease in maintenance and space utilization.
- **Version**: This represents a database version within a Gold Image. Versions can have multiple states: Draft, Active and Current. To deploy a particular version from Gold Image, ensure that you set desired version as Current.
- Current version: Current version is a state of version, which will be deployed by Fleet Maintenance. Current Version can be interpreted Gold Image, and hence can be used interchangeably.
- Update: It denotes a patching operation, where you update database from one version to higher within the same Release. Eg: A 19.15 to 19.22 operation is refereed as update operation.
- **Upgrade**: When you move database from one release to another, its referred as upgrade. Eq: A 12.2.0.4 to 19.22 operation is referred as upgrade operation.

Before you begin: Meeting the Prerequisites

This section describes how you can set up the infrastructure for patching. Meet these prerequisites before you start patching.

This section covers the following:

- · Meeting Basic Infrastructure Requirements for Patching
- User Roles and Privileges for Fleet Maintenance
- Preferred Credential Support for Fleet Maintenance
- Searching for Patches on My Oracle Support
- Search and View Patches in Oracle Software Library
- Downloading Patch Recommendations and Patches
- Security Models for running Root scripts
- Prerequisites for performing Fleet Maintenance on Windows

Meeting Basic Infrastructure Requirements for Patching

Before you begin Fleet Maintenance operations, you must ensure that the following prerequisites are met:

- The Enterprise Manager Software Library has been setup and configured, see Getting Started with Setting Up Your Infrastructure.
- The database and cluster targets to be patched have been discovered in Enterprise Manager, see: Discover Autonomous Databases and Discovering and Adding Database Targets.
- 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. For more information on creating reference environments see, Image Maintenance.
- Space is available on the hosts where the patching is to be carried out.
- The Agent Port on the OMS server needs to be open, allowing for communication between the target server and OMS server.
- Before attempting an Update or Upgrade remove the glogin.sql file from the target Oracle Home, it is located in the /sqlplus/admin sub-folder.
- Verify Agent tuning parameters to ensure performance and stability, see MOS note: 2449197.1.

User Roles and Privileges for Fleet Maintenance

This topic cover roles and privileges a user will need to have in order to be able to execute all Fleet Maintenance Operations.

The main role needed for Fleet Maintenance operations is: EM_PATCH_OPERATOR, to grant this role navigate from Enterprise Manager to **Setup**, select **Security** and finally select **Administrators**. Click on the **Create** button and provide a user name and password.



Table 17-1 User Roles and Privileges for Different Oracle Homes

Type of Oreals Horse	Polo	Target Privilege	Docouros Privilego
Single Instance Database	Role EM_PATCH_OPERATOR	Listener Privileges: Configure Target Blackout Target Used for: Database and listener on the destination targets.	View Credential for normal and privileged credentials that will be used for operations.
Oracle Restart (HAS)	EM_PATCH_OPERATOR	Listener Privileges: Configure Target Blackout Target Used for: Automatic storage management and Oracle High Availability Service targets running from the Oracle Home being patched.	View Credential for normal and privileged credentials that will be used for operations.
RAC Database	EM_PATCH_OPERATOR	Listener Privileges: Configure Target Blackout Target Used for RAC database instances and listener targets running from the Oracle Home being patched.	View Credential for normal and privileged credentials that will be used for operations.
Grid Infrastructure Cluster	EM_PATCH_OPERATOR	Listener Privileges: Configure Target Blackout Target Used for: Automatic Storage Management, Cluster and Oracle High Availability Service targets for the destination targets.	View Credential for normal and privileged credentials that will be used for operations.
Exadata Patching	EM_PATCH_OPERATOR	Listener Privileges: Configure Target Blackout Target Used for: Exadata database targets being patched.	View Credential for normal and privileged credentials that will be used for operations.

Preferred Credential Support for Fleet Maintenance

Preferred credentials are recommended for all Fleet Maintenance operations, the following is a list of preferred credentials based on the Fleet Maintenance operation.

Table 17-2 Preferred Credentials for Fleet Operations

Operation	Preferred Credential
Deploy SI Home	Normal Host Credentials for Oracle Home
	Privileged Host Credentials for Host

Table 17-2 (Cont.) Preferred Credentials for Fleet Operations

Operation	Preferred Credential
Migrate Listener	Normal Host Credentials for Oracle Home
	Privileged Host Credentials for Host
Update DB	Normal Host Credentials for Oracle Home
	Privileged Host Credentials for Host
ROLLBACK_DB	Normal Host Credentials for Oracle Home
	Privileged Host Credentials for Host
Deploy RAC Home	Normal Host Credentials for Oracle Homes (for both nodes)
	Privileged Host Credentials for Hosts
Update RAC DB	Normal Host Credentials for Oracle Homes (for both the nodes)
	Privileged Host Credentials for Hosts
Deploy GI	Host Credentials for Oracle High Availability Service
	Privileged Host Credentials for Hosts
Update GI	Host Credentials for Oracle High Availability Service
	Privileged Host Credentials for Hosts
DEPLOY_CDB	SYSDBA Database Credentials for CDB (Database Instance)
UPDATE_PDB	SYSDBA Database Credentials for CDB (Database Instance)
	Normal Database Credentials for PDB (Pluggable Database)

To set up preferred credentials see: Setting Up Credentials.

Privileged Access Management

Fleet Maintenance is compatible with PAM (Privileged Access Management), allowing you to eliminate the need for hard-coded application credentials embedded in applications, scripts, or configuration files. Allowing highly sensitive passwords to be centrally stored, logged, and managed within the Vault.

To set up PAM in your Enterprise Manager deployment see: (Optional) Setting a PAM User Account.

Searching for Patches on My Oracle Support

The search functionality allows you to perform flexible and advanced searches. It also offers capabilities like saving a frequently used search, and searching based on existing saved searches.

To search a patch on My Oracle Support, follow these steps:

- 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**.



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.

Use the **Saved** tab to search any previously saved searches. Use the **Recent** tab to access any recently performed search.

Once 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.

3. To search for patches using EMCLI, you can use the following command examples:

```
emcli search_patches -patch_name=<Patch number>
emcli search patches -patch name=<Patch number> -swlib -script
```

For more information on using the search_patches EM CLI verb see: search_patches in *Oracle Enterprise Manager Command Line Interface*.

Search and View Patches in Oracle Software Library

To search for a patch within the Software Library and verify it was correctly downloaded, follow these steps:

- 1. From the Enterprise menu, select Provisioning and Patching, then select Patches & Updates.
- On the Patches & Updates page, in the Software Library Patch Search region, enter the search parameters you want to use and click Search.



If you do not see the Patch Search region, click **Customize Page** from the topright 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.

Downloading Patch Recommendations and Patches

This section describes the different patching modes that can be used:

Patching in Online and Offline Mode

You have the flexibility to choose between Online and Offline patching.

 Online Mode: Online mode is useful when Enterprise Manager can connect to My Oracle Support (MOS) using an Internet connection, enabling Enterprise Manager to download the latest Oracle Patch Recommendations. Using this mode, search and download patches from MOS and create gold images to include these patches. In addition, you can access community information, knowledge articles, service requests, and automatically resolve patch conflicts with a merge patch directly from MOS.



Cloud Control does not upload data to MOS. It only uses MOS to download the latest updates.

For patching in online mode see: Setting up MOS Connection for Patch Recommendations and Patch Downloads

- Offline Mode: Offline mode is used when Enterprise Manager cannot connect to My Oracle Support. In this mode, you can search patches that were manually uploaded to the Software Library. 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 to offline, go to **Setup**, select **Provisioning and Patching**, and then select **Offline Patching**.

For patching in offline mode see: Setting Up the Infrastructure for Patching in Offline Mode (Not Connected to MOS)

Out-of-Place Patching

Out-of-Place patching is a mechanism where the gold image containing the required patches is deployed into a new home. Once complete, you migrate the database instances to run from the new home, ensuring minimal downtime.



If you have to apply patches to address critical issues without waiting for the new gold image to be published, the patching flow will clone the existing home and apply the required patches. The database can then be migrated to the new home. For more information see: Emergency Patching

If you choose to migrate only a few instances per session, make sure that you migrate the rest in the next session. This way, you can control database downtime in the data center by dividing the migration activity. This is particularly useful when you have multiple database instances running out of an Oracle home.



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), or Oracle Data Guard targets you can choose to patch the instances of the cluster either in rolling or parallel mode.



Non-rolling updates are deprecated and no longer available, starting with Enterprise Manager 13c Release 5 Update 17 (13.5.0.17).

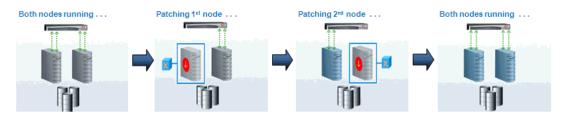
• **Rolling Mode**: In Rolling mode the nodes of the cluster are patched individually, one by one. For example, if you are patching a. Oracle Grid Infrastructure target that has five nodes, the first node is shut down, patched, and restarted. Then the process is rolled over to the next node until all the nodes are patched successfully.



The **ReadMe** of the patch states whether or not you can use Rolling Mode to apply your patches. Use this mode after ensuring that all patches in the gold image support rolling mode.

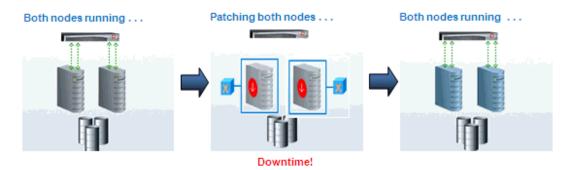
The following image illustrates how a two-node Oracle RAC target gets patched when rolling mode is used.

Figure 17-3 Rolling Mode of Patching



Parallel Mode: In Parallel mode all the nodes are patched at the same time, all the nodes
are shut down and the patch is applied on all of them at the same time.
 The following image illustrates how a two-node Oracle RAC target gets patched when
parallel mode is used.

Figure 17-4 Parallel Mode of Patching





Setting up MOS Connection for Patch Recommendations and Patch Downloads

In Online mode Enterprise Manager connects to MOS to download patches, release updates, monthly release patches, ARU seed data (products, platforms, releases, components, certification details, and patch recommendations). If you have setup a proxy server, register the proxy details using the My Oracle Support Proxy Settings page.

For online (connected to MOS) patching make sure you complete the following setup requirements:

- Enabling Online Mode for Patching
- Registering the Proxy Details for My Oracle Support

Enabling Online Mode for Patching



These steps are to be performed of you previously patched in Offline mode and need to revert to Online mode. Enterprise Manager does not upload data to MOS, it only downloads the latest updates.

Online mode has to be enabled to automatically download patch recommendations and to download patches to the software library. To patch the targets in Online mode, set the connection setting in Cloud Control to Online mode by logging in as a user that has the Patch Administrator role, and following these steps:

- From the Setup menu, select Provisioning and Patching, then select Online Patching.
- 2. For Connection, select Online.

Registering the Proxy Details for My Oracle Support

To register the proxy details for My Oracle Support (MOS), follow these steps:



My Oracle Support accesses support.oracle.com directly, provide network access to this URL or grant proxy access to the URL.

- From the Setup menu, select Proxy Settings. Then select My Oracle Support.
- If you want the OMS to connect to MOS directly, without using a proxy server, follow these steps:
 - 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.



- Specify the proxy server host name for HTTPS and an appropriate port number for Port.
- c. If the specified proxy server has been configured using a security realm select Password/Advanced Setup and provide values for Realm, User Name and Password.
- d. Click **Test** to verify 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 allow connectivity to aruakam.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 the requires sites, add the URLs mentioned above to the Unauthenticated Sites Properties of the proxy server.
- Details specified within the MOS Proxy Settings page apply to all OMS in a multi-OMS environment.

Setting Up the Infrastructure for Patching in Offline Mode (Not Connected to MOS)

What if my Enterprise Manager cannot connect to MOS?

For offline (disconnected to MOS) patching make sure you complete the following setup requirements:

- Download and Upload Enterprise Manager Catalog Zip File in Offline Mode
- Uploading Patches to Oracle Software Library

Download and Upload Enterprise Manager Catalog Zip File in Offline Mode

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 on a regular basis download the $em_catalog.zip$ file from My Oracle Support. To download the catalog file use URL: $https://updates.oracle.com/download/em_catalog.zip$.

Information about the targets affected by the latest patches, and patches that you have to download manually are available within the catalog zip file.

Uploading Enterprise Manager Catalog Zip File from your Host With No Internet Connectivity

After downloading transfer the $em_{catalog.zip}$ file back to your local host using FTP or other file transfer methodology. Then from your local host log in to Enterprise Manager to upload the zip file, follow these steps:

- 1. In the Setup menu, select Provisioning and Patching, then select Offline Patching
- 2. Click Browse and specify the location of the latest em catalog.zip file.



Click Upload.On clicking Upload, a refresh from My Oracle Support job is created and submitted.

Uploading Patches to Oracle Software Library

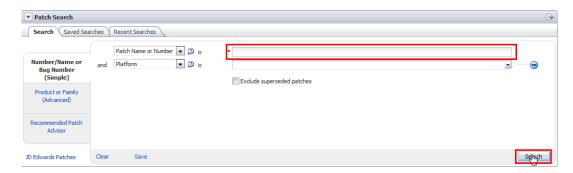
When patching targets with Enterprise Manager in offline mode you must first manually download from My Oracle Support, the patches and their metadata files. Once downloaded, add them manually into the Software Library so that they can be searched, selected, and added to the Gold Images for patching.

Downloading a Patch from My Oracle Support

To download a patch and its metadata file from My Oracle Support, follow these steps:

- 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 17-5, then click **Search**.

Figure 17-5 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 17-6.



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.

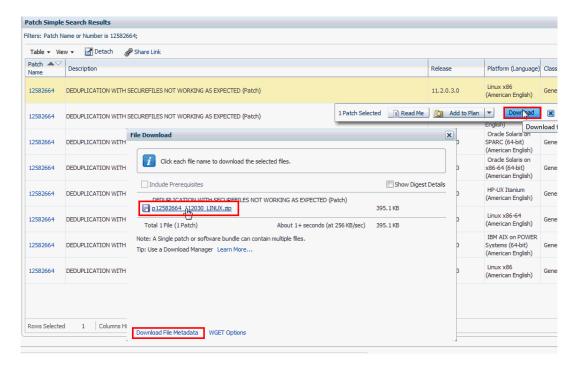


Figure 17-6 Downloading Patches from My Oracle Support

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**.
- 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, restart 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:

 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=2benot2be
```



Ensure that the EM CLI log in user has the ADD TARGET privilege.

Synchronize EM CLI:

```
<emcli install location>/emcli sync
```

4. Run the upload patches verb to upload the required patches to Software Library:

The parameters mentioned in [] are optional.

Use the <code>-from_host</code> option to specify the host on which the downloaded patch zip files and metadata files are present. You can use the <code>-location</code> option to specify the location of the

downloaded patch files on the host you specified using <code>-from_host</code>. When you specify <code>-location</code>, 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 <code>-patch_files</code> 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 <code>-cred_name</code> option to specify the named credentials that must be used to access the specified host, and the <code>-cred_owner</code> option to specify the owner of the specified named credential. If you do not specify the <code>-cred_name</code> option, the preferred normal credentials are used to access the host. If you do not specify the <code>-cred_owner</code> 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.



Ensure that you specify either the <code>-location</code> option or the <code>-patch_files</code> option with the <code>upload_patches</code> verb, but not both. Specifying the <code>-location</code> 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

Security Models for running Root scripts

Root Components for Fleet Maintenance (Patching)

- Manually Staging the Root Components For Fleet Maintenance
- Restricting the Root User Access For Fleet Maintenance

Prerequisites for performing Fleet Maintenance on Windows

Fleet Maintenance operations can be performed on Windows based deployments. This section covers specific prerequisites needed only when using Fleet Maintenance on Windows deployments.

 To perform Fleet Maintenance operations like: deploy, migrate, and/or rollback; you will need to use the same user account and credentials that the source database is using. For more information see: Setting Credentials for the Job System to Work with Oracle Enterprise Manager in *Oracle Database Installation Guide 19c for Microsoft Windows*.

- User details are passed to Fleet Maintenance operations via Enterprise Manager Named Credentials. There are two typical user configuration cases:
 - Installations where there is a single installation owner that is also the Oracle Home owner. This single owner's credentials will be passed to the fleet operations at the destination database.
 - Installations where there is an installation owner (local administrator) and there is a second local user with no administrator privileges who owns the Oracle Home. Both user's credentials will be passed to the fleet operations at the destination database.
- The Operating System users (mentioned above) need to be granted specific Windows OS level privileges on the specific host or hosts for Fleet Maintenance operations to properly function. For more information see: Creating Operating System Groups and Users for Enterprise Manager Cloud Control in Oracle Enterprise Manager Cloud Control Basic Installation Guide.
- CYGWIN needs to be installed for Gold Image creation in a Windows deployment. For more information see: Installing Cygwin and Starting the SSH Daemon in Oracle Enterprise Manager Cloud Control Basic Installation Guide.



Database Updates and Upgrades Using Fleet Maintenance

This chapter describes the procedures to standardize database environments by automatically patching a large number of databases with minimal downtime.

It covers the following sections:

- Get Organized With Standardization Advisor
- Use Fleet Maintenance Hub to get to Patch Compliance
- Patch Recommendations
- Image Maintenance
- Update Operations
- Fleet Maintenance Automation
- Patch History

Get Organized With Standardization Advisor

Discovering Configuration Pollution

Patching is one of the most challenging phases in the product lifecycle. With multiple patch cycles and various database configurations, different sets of patches are required to satisfy the patching needs of the users. This may require administrators to check for patch conflicts for each database and then request merge patch patches for the same. This can be cumbersome and time-consuming process.

In order to simplify this process, administrators must first identify their specific database configuration needs and select the standard configurations that are to be enforced in their environment.

With Enterprise Manager 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. This enables administrator to identify and reduce the number of configurations they need to manage. This also brings down the need to assess patch conflicts per database and, subsequently, reduce the number of merged patches they need to request from Oracle.

For example, in a Oracle 19c database on a 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.

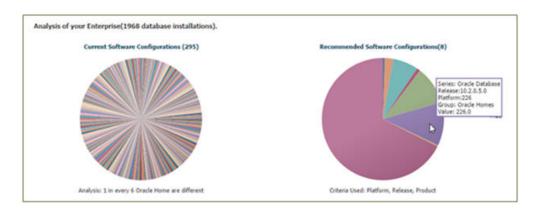
Standardize on Supported Images



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.
- 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 Configuratio n	No of Database s	Database
Oracle Database 12c	12.1.0.2.0	Linuxx86-64	2	19769480;20 299023;2041 5564.	1	SI: [salesCDB.local domain]
Oracle Database 12c	12.1.0.2.0	Linuxx86–64	3	19769480;20 299023.	2	SI: [fm1.oracle.com; HRdb0000]
Oracle Database 19c	19.0.0.0.1	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 Configuratio n	No of Database s	Database
Oracle Database 19c	19.0.0.0.6	Linuxx86-64	8	19769480;20 299023;2041 5564.	6	SI: [HRdb0000;fm2. oracle.com;sale sCDB.localdom ain;fm1.oracle.c om;FINdb000;d b01]



Use Fleet Maintenance Hub to get to Patch Compliance

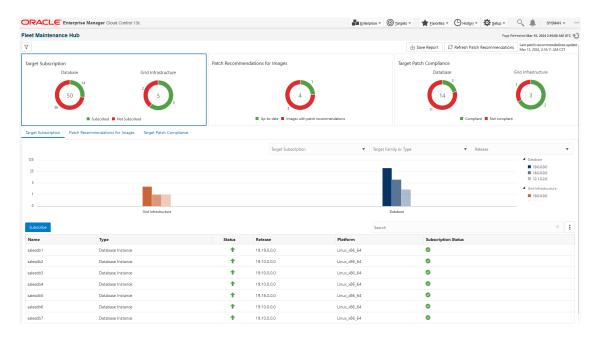
Use Fleet Maintenance Hub to subscribe targets, incorporate patch recommendations into images, and get targets to patch compliance.

For Oracle Enterprise Manager 13.5 Release Update 16 (13.5.0.16) and higher, you can access the Fleet Maintenance Hub page by navigating to: **Targets** then, **Databases**, then **Administration**, and click on **Fleet Maintenance Hub**.

Starting with Oracle Enterprise Manager 13.5 Release Update 21 (13.5.0.21) and higher, you can use Fleet Maintenance Hub to create Gold Images and add new versions to existing images. Fleet Maintenance Hub accommodates maintenance tasks at both target and image levels for administrators, while operators can utilize it for update or upgrade operations. As an operator you still have the option to use the existing Fleet Maintenance UI for these operations.

Target Subscription

Figure 18-1 Fleet Maintenance Target Subscription



The central section of the Target Subscription tab, shows two different sets of bar graphs, each set denoting release specific unsubscribed databases and Grid Infrastructures. By clicking on any bar, the table below the graph will show databases and/or Grid Infrastructure running on that release and not subscribed to any particular image.

You can further filter and refine the available data, by using the following filtering criteria: Target Type or family (database or Grid Infrastructure), Release (filter by version), and a text field to search by database name.

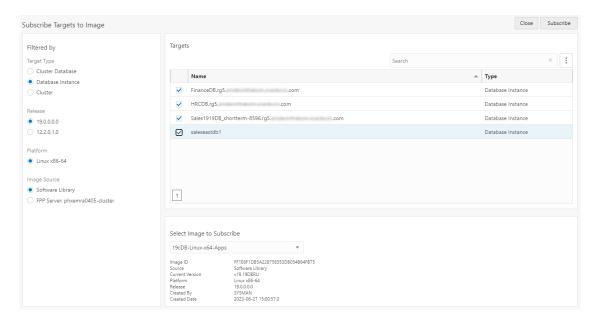
Note:

In order to subscribe a database to a image, it is recommended that you have already identified the image to subscribe the database to.

Subscribe Targets to Image

Click on the **Subscribe** button, a new slide out will appear. To the left, you will see a search function similar to the previous screen's, with one additional field Image Source. To select a image in the slideout, filter by Target Type, Release, Platform, and for Image Source select either **Software Library** or **FPP Server** depending on your environment. On the right side you will see the list of available targets based on the selected criteria, select the targets to be subscribed to the image. You can select multiple targets at once. Under Select Image to Subscribe, select the image to be used from the drop down menu, once selected you can see the image details.

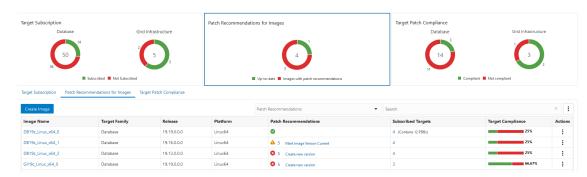
Figure 18-2 Subscribe Targets to Image



Click on **Subscribe** to associate the selected targets to the Gold Image. Once the operation completes you can close the slide out or select **Continue** to subscribe more targets to the images.

Patch Recommendations for Images

Figure 18-3 Patch Recomendations for Images





The second tile shows patch recommendations (Oracle published patch recommendations and additional one off patches) for already existing images. The pie chart allows you to see up to date images, and images with pending patch recommendations to be applied. Click on the image to view its details in the table (below the pie chart) form with the following fields:

Note:

If you have recently subscribed or unsubscribed a target, run a **Refresh Path Recommendations** before selecting **Patch Recommendations for Images** to ensure the latest target information details are pulled.

- **Image Name**: Name of the Gold Image, by clicking on it a pop-up will appear with image details.
- Release: Release version to which the image has been created (eg. Oracle Database 12.1.0.2 or 19.0.0.17).
- **Platform**: Platform operating system which was used to create the image, (eg. Linux x86-64).
- Target Type: Whether the image is for database or Grid Infrastructure.
- Patch Recommendation: This column indicates the health of the image.
 - A green check mark indicates that the current version of the image has all the required latest recommended patches published by Oracle, as well as additional one off patches. The image is in good shape and you should use this image to perform update or upgrade operations.
 - A yellow exclamation mark indicates that one of versions contained in the image has all the required latest recommended patches published by Oracle, but it is not marked current. Only current versions can be deployed. An administrator role is required to change the version of the image to CURRENT.
 - A red X mark indicates that no version contained in the image contains all the latest recommended patches. The administrator has to create a new version including the recommended patches and mark it CURRENT.
- Subscribed Targets: Shows details about the targets subscribed to the particular image.
 Click on the numeric to see more details about the targets that are subscribed to the image.
- **Patch Compliance**: This linear graph shows how many of the subscribed targets are on the current version. If you see any targets in red, this indicates that they are not on the current version of the image and need to be updated.
- **Action**: You can perform image management and update/upgrade operations on the subscribed targets from this menu.

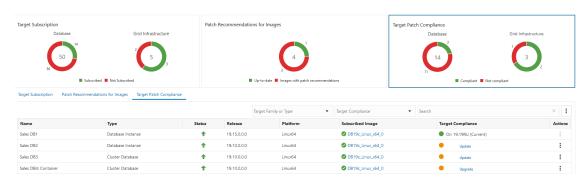
Note:

Images created directly in FPP and synced with Enterprise Manager will lack information about bugs fixes included in the image, impacting patch recommendations for these images. Oracle recommends creating the FPP images using Enterprise Manager to avoid this issue with patch recommendations.



Target Patch Compliance with Images

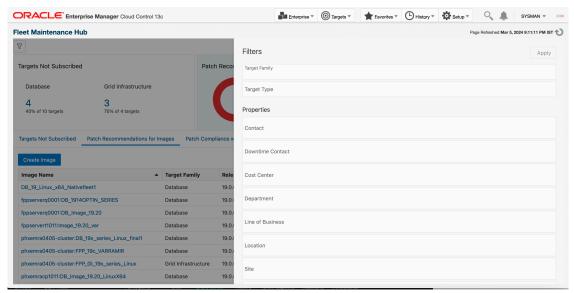
Figure 18-4 Target Patch Compliance with Images



In the Target Patch Compliance with Images chart, you can see compliant targets and non-complaint targets. Click on any section of the chart, and this will populate the grid accordingly. The table will show the following information:

- Name: Target name of the database or Grid Infrastructure.
- Type: Type of target, database or Grid Infrastructure.
- Release: Release version to which the target has been created (eg. Oracle Database 12.1.0.2 or 19.0.0.17).
- **Platform**: Platform operating system which was used to create the target, (eg. Linux x86-64).
- **Subscribed Image**: Image name to which this target has been subscribed. Click on the image name to see details about the image.
- Patch Level: Will show the patch status for a target, the following is a list of possible statuses:
 - If the target is on a current version it will show as Current and green.
 - If the target has additional patches pending it will show **Drifted** and yellow.
 - If outdated it will show the name of the Image Version it is on and will show as red.
- Action: Launches the UI to start the update operation.

Global Filters



With Global Filters you can further filter a search by filtering by target family, target type, and with Enterprise Manager user defined properties. To apply Global Filters click on the filter Icon located at the top left corner of the Fleet Maintenance Hub page, this will open the Filters slide-out. Select all additional filter options and click **Apply**.

The filters that have been applied can now be viewed in the Global Filters bar at the top of the Fleet Maintenance hub page

Patch Recommendations

Additional patch recommendation operations can be performed to keep your patch information up to date

Refresh Patch Recommendation Job

To refresh patch recommendation jobs, navigate to the Fleet Maintenance Hub page, and click on **Refresh Patch Recommendations**. This will update and refresh patch recommendations for all subscribed images.

Review and Save a Report

To save a CSV report of the patch recommendations navigate to the Fleet Maintenance Hub page, and click on **Save Report** located a the top right corner.

Subscribe Databases to the Gold Image

Once you have created a new up-to-date Gold Image, the next step is to subscribe databases to the image. Follow these steps to subscribe databases to images using the Fleet Maintenance Hub UI:

- In the Fleet Maintenance Hub screen, select the first tile **Target Subscription**, this will load a list of all non-subscribed database targets. Click on the **Subscribe** button above and to the left of the targets table.
- In the Subscribe Targets to Image screen select the databases you wish to subscribe, to assist your search you can apply filters and, or use the search bar. Once the database targets have been selected, select the image to subscribe from the drop down menu and click on Subscribe.



3. Once all targets have been subscribed, the Subscribe Targets to Image screen will show the status of the subscription operation. To continue adding targets click on Continue Subscribing, to return to the Fleet Maintenance Hub click on Close.
At this point the database target has been subscribed to the Gold Image.

Additionally you can verify that the target databases have been subscribed to the image in the Fleet Maintenance Hub screen, by selecting the first tile **Target Subscription**. The recently subscribed targets will no longer appear in the Target Subscription table.

Image Maintenance

To begin creating a Gold Image using the Fleet Maintenance UI log in to Enterprise Manager, navigate to **Targets**, and select **Databases**. Once in the Databases screen navigate to **Administration**, click on **Fleet Maintenance Hub**.

Create a Gold Image

- 1. In the Fleet Maintenance Hub page, select the **Patch Recommendations for Images** tile and click on the **Create Image** button
- 2. The Create Image screen appears at the first step Image Details, enter the following information:
 - Source Details: Select the reference Oracle Home, and work directory. Once the
 Oracle Home has been selected if the reference home has missing patch
 recommendations you can select whether to patch the selected Oracle Home, or clone
 the selected Oracle Home and patch the cloned home.
 If you select to clone the selected Oracle Home and patch the cloned home, enter the
 cloned home path and select whether you want to delete the cloned home after image
 creation.
 - Credentials: Enter the Normal and Privileged host credentials. Click the Named radio button if you are using named credentials.
 - New Image Details: Verify the information displayed and enter the Version Name. This field is not editable when creating a new image. To edit you will need to Create Image.
 - Storage: Select the destination, storage type for the Software Library (NFS or OMS Shared), enter the storage name for the Software Library, and the image location in Software Library.

Once complete, click Next.

3. In the Add Patches step select all patches that will be applied, and click **Next**.



Patches must be saved within Enterprise Manager's software library to be applied.

4. In the final Review screen verify that all the information captured is correct. You can preview the JSON file that will be used to submit the procedure by clicking on Preview JSON. Once you have reviewed and are ready to proceed with creating the Gold Image click Submit.





To verify the status of the Gold Image creation activity navigate to the **Procedure Activity** dashboard and monitor the status in real time.

The Gold Image has been successfully created.

For Enterprise Manger releases previous to 13.5 Release Update 20 (13.5.0.20), follow the EM CLI steps outlined here: Image Maintenance.

Create a new version of the image

- 1. In the Fleet Maintenance Hub page, select the **Patch Recommendations for Images** tile and select the image. Navigate to **Actions** and click on **New Version**.
- 2. The Create Image screen appears at the first step Image Details, enter the following information:
 - Source Details: Select the reference Oracle Home, and work directory. Once the
 Oracle Home has been selected if the reference home has missing patch
 recommendations you can select whether to patch the selected Oracle Home, or clone
 the selected Oracle Home and patch the cloned home.
 If you select to clone the selected Oracle Home and patch the cloned home, enter the
 cloned home path and select whether you want to delete the cloned home after image
 creation.
 - Credentials: Enter the Normal and Privileged host credentials. Click the Named radio button if you are using named credentials.
 - New Image Details: Verify the information displayed and enter the Version Name. This field is not editable when creating a new image. To edit you will need to Create Image.
 - Storage: Select the destination, storage type for the Software Library (NFS or OMS Shared), enter the storage name for the Software Library, and the image location in Software Library.

Once complete, click Next.

In the Add Patches step select all patches that will be applied, and click Next.



Patches must be saved within Enterprise Manager's software library to be applied.

4. In the final Review screen verify that all the information captured is correct. You can preview the JSON file that will be used to submit the procedure by clicking on Preview JSON. Once you have reviewed and are ready to proceed with creating the Gold Image click Submit.



To verify the status of the Gold Image creation activity navigate to the **Procedure Activity** dashboard and monitor the status in real time.



The Gold Image has been successfully created.

Retrieve a List of Available Gold Images

To retrieve a list of available Gold Images in the Fleet Maintenance Hub page, select the **Patch Recommendations for Images** tile. A list of images will populate the table, under **Patch Recommendations** all images with a **Green** icon will be at the Gold Image standard level.

Change a version status to current

Once the new Gold Image has been created with the latest patches, in the Fleet Maintenance Hub the Gold Image will show with a yellow exclamation icon. This means that the image is up to date, but is not yet marked as current. To mark as current follow these steps:

- Click on the Actions menu to the right of the Gold Image and click on Make Image Version Current.
- 2. A popup will show all versions with that image, select the version created in the previous steps and click **Update** on the upper right hand side. Once the update has successfully completed, click **Close** to return to the Fleet Maintenance Hub page.

The Gold Image will now show with a green check mark icon, and the image version has been marked current.

Editable parameters of Image and Version

To edit the parameters of an image or version within an image, select the **Patch Recommendations for Images** tile and choose the image that needs editing, click on the **action menu** and select **Manage Image**.

Image: Under Image Details, the Image Name and Description fields are editable

Version: To edit a version, locate the version to be edited, click the **action menu** and select **Edit**.

Speed up deployment with NFS

Starting with Oracle Enterprise Manager 13.5 Release Update 13 (13.5.0.13) Fleet Maintenance supports NFS storage type. You can create and import Fleet Maintenance Gold Images using NFS storage options. This allows for a more efficient deployment to multiple hosts, reducing the number of copies and files by storing the Gold Image in an NFS reference location. NFS integration is supported for all image creation types including reference home and emergency patching.

For NFS environments a referenced file location needs to be created. To create a Referenced File Location for NFS follow these steps:

- 1. From the Targets menu select Administration, then Database Provisioning, click on Software Library.
- Navigate to the Referenced File Locations tab, under Storage Type select NFS from the drop down menu and click Add.
- 3. In the Add NFS Location dialog box add a Name for the NFS location, the NFS Server IP or host name for the NFS server, and the absolute path Location of the NFS share where the Gold Images, Archives, and Software Binaries will be stored.



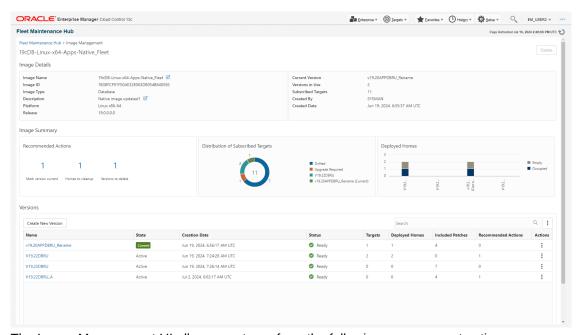


The same Gold Image can be used for deployments on both Oracle Red Hat Linux 7 and 8 deployments. This is due to the database media being identical in both versions, binaries will be relinked during the deployment process.

Image Management

Starting with Enterprise Manager 24.1 Release Update 2 (24.1.0.2), Fleet Maintenance now includes the Image Management UI, allowing you to create and delete images, review image utilization (based on associated targets per version), and assess recommendations at both the image and version levels.

Image Management Overview



The Image Management UI allows you to perform the following management actions:

- Image Details: The top-most section of the UI provides high-level image details, like name, ID, type, description, platform, release number, version, in use versions, targets subscribed, created by and date of creation.
 Image Name and Description are editable fields and can be modified directly from the UI, to edit click on the pencil and notepad icon to the right of the Image Name and Description fields.
- Image Summary: Shows the overview of a selected image allowing you to see recommended actions to be taken on the selected image, a donut chart with a distribution of subscribed targets within the image, and a bar graph detailing empty and occupied Oracle Homes within the image.
 - Recommended Actions: Allows you to see recommended actions to be taken on the selected image: pending marking a version as current, homes to be cleaned, and versions to be deleted from the image. Click on the numeric values next to these recommendations to see the tasks required, as well as the corresponding REST API calls to perform the recommended actions.



- Distribution of Subscribed Targets: A donut chart that visually represents how targets associated with this image relate to various versions. If you see "drifted" or "upgrade required" statuses, we recommend taking appropriate actions to keep targets updated.
- Deployed Homes: The bar graph chart displays information detailing occupied and empty Oracle Homes. Oracle recommends that older empty homes from previous versions be deleted based on your cleanup policy.
- Versions: Allows for a table view of all versions available within the image allowing you to view information, such as State and Creation Date. Within the table the Included Patches field lists the patches applied to this version. Finally the Recommended Actions shows version-level recommendations.



Version-level recommendations are distinct from the broader recommendations found in the **Image Summary** section.

To create a new version of the image click the **Create New Version** button, to navigate to the Create New Version flow, for complete set of steps within the flow see Create a new version of the image.

Update Operations

You can now perform update operations from the Fleet Hub Maintenance Hub page, using a simple intuitive process that guides you through setting up and performing Fleet Maintenance to your Oracle Database targets.

Follow the following for a successful Fleet Maintenance update operation:

- Update Databases Using Fleet Maintenance
- Track, Fix and Resume Update Procedures
- Use of Oracle Home Path Variables
- Custom Pre and Post Scripts for Fleet Operations

Note:

Do not perform lifecycle management activities during a patching window, as the monitoring credentials are locked during patching and cannot be modified, created or deleted. Examples of the activities to avoid performing during a patching window: provisioning, patching, cloning, migration and similar activities from Self-service portal for DBaaS/PDBaaS.

Update Databases Using Fleet Maintenance

You can now perform Migrate, Deploy and Update operations from the User Interface page using a simple intuitive process that guides you through setting up and performing Fleet Maintenance to your Oracle Database targets.



Note:

For DST Fleet Maintenance patching see MOS Note 2725653.1

- To begin performing updates using Fleet Maintenance
 - From the Fleet Maintenance Hub, navigate to the Target Patch Compliance tile. In
 the table below locate a database to update and click on Update under the Target
 Compliance column. This will preselect the database target, and move you into the
 Targets screen.
 - From the Fleet Maintenance Operator UI, in the Targets screen, select the databases you wish to perform the operations upon by check marking them. You can filter to find the databases by filtering via: Image, Target Type (Database Instance or Cluster Database), Operation (Update, Upgrade or Upgrade Non-CDB as Pluggable Database); or the Filter text box.
 Starting with Enterprise Manager 24.1 Release Update 2 (24.1.02), you can use Global Filters in the Fleet Maintenance Operator Wizard used for updates. Global Filters allow you to further filter a search by filtering by target family, target type, and with Enterprise Manager user defined properties. To use click on the filter icon, this will open the Filters slide-out. Set the filters you wish and click on Apply.

Once all databases you wish to perform Fleet Maintenance on are check marked, click **Next**.

2. Within the Operations screen you check mark the operation or operations to be performed under Maintenance Tasks:



You can also click on a link on each operation, this will also tell you what targets the operation is applicable to.

• Under Software Deployment: Enter the Destination Oracle Home, and enter the Home Name Prefix if your environment requires a custom Oracle Home name prefix, this allows for the ability to discover Oracle Homes with custom name post deployment. If the home name prefix is not specified, the default value will be used. This feature is available starting with Oracle Enterprise Manager 13.5 Release update 20 (13.5.0.20).

Set a **Start Schedule**. The schedule you enter must reflect across all applicable selected targets. You can select for the operation to start **Immediately** or **Later** at a time of your choosing.



If you select multiple databases running on different homes from the same host, the Fleet Management UI will only submit deployment for one database. This is because the Oracle Home path needs to be unique per source Oracle Home.

Migrate the Listener: Select a start time. Immediately/Auto the task will begin
immediately or as soon as the previous task finishes. Later, enter a start time of your



choosing that reflects the time usually taken for software deployment in your environment.



All times are set in the Management Server Time Zone. You cannot set a Maintenance Task to run immediately if there are preceding tasks.

- Update Database: Select a start time. Immediately/Auto the task will begin
 immediately or as soon as the previous task finishes. Later, enter a start time of your
 choosing that reflects the time usually taken for software deployment in your
 environment.
- 3. Enter the Credentials for Normal Host, Privileged Host and Database Credentials. You can select whether to use a Named or Preferred credential. Preferred credentials need to be set at the Oracle Home and target database before you can selecting them in the UI. For more information on preferred credentials in Fleet Maintenance see: Preferred Credential Support for Fleet Maintenance.



For Windows deployments use Windows Security Credentials.

- 4. Under **Options** you can set the following:
 - Work Directory: Enter the work directory where the files will be stored.
 - **Skip Cluster Verification Utility Checks**: Select to skip the Prerequisite Checks and/or Cluster Verification Utility Check.
 - Dispatch Location: Enter a read-only location to use pre-staged dispatcher and fleet scripts.
 - **Listener Options**: Select whether or not to Force the listener.



This option is only available if Migrate Listener is selected.

- Custom Scripts: Select the optional custom scripts you wish to run by clicking on Add, the Fleet Management UI is set to the default location for custom scripts. If your scripts are located in a different section Set the Stage, Script Location. Lastly select the check mark if you wish to Run as Root.
 - For more information on custon scripts see: Custom Pre and Post Scripts for Fleet Operations

Once all options have been entered click **Next**.

- 5. In the **Review** screen, verify all options entered in the previous screen. In this screen you can **View JSON** or **Save JSON** of the operations to be performed.
 - Click on **Validate** to perform a **Quick** or **Full** validation of your Fleet Maintenance Operation before submitting. Validation results are reported on a per target basis, for a **Full** validation a Deployment Procedure is submitted for each target. Each target will have a link available to see in detail the results of the Validate Deployment Procedure.



Note:

If you encounter a *Destination home not found* error when performing a Quick validation you will need to perform a Runbook Session. Navigate to **Enterprise**, then **Monitoring**, then **Runbook Session**. Select **Start Universal Runbook Session**, and initiate the session labelled **Fleet Maintenance Operation Failed: Home Deployment not recognized**. This will help identify the root cause of this issue.

6. To start your Fleet Maintenance operation, click **Submit**.

Note:

Starting with Enterprise Manager 13.5 Release update 19 (13.5.0.19), before and after a UPDATE/ROLLBACK operation on single instance or RAC databases, the deployment procedure activity logs include the number of invalid objects found.

Data Guard Update Operations Using the UI

1. To begin performing Fleet Maintenance operations on Data Guard enabled targets, in the Targets screen, select the databases you wish to perform the operations upon by check marking them. Data Guard enabled databases will show an additional Role category in the table, with Primary and Physical Standby values assigned to each respective database. You can filter to find the databases by filtering via: Image, Target Type, Operation, or the Filter text box.

Starting with Enterprise Manager 24.1 Release Update 2 (24.1.02), you can use Global Filters in the Fleet Maintenance Operator Wizard used for updates. Global Filters allow you to further filter a search by filtering by target family, target type, and with Enterprise Manager user defined properties. To use click on the **filter** icon, this will open the Filters slide-out. Set the filters you wish and click on **Apply**.

Once all databases you wish to perform Fleet Maintenance on are check marked, click **Next**.

Note:

You can select only Primary or only Standby databases in a Fleet Maintenance session but not both at the same time for the same session.

2. Within the **Operations** screen you check mark the operation or operations to be performed under **Maintenance Tasks**:

Note:

You can also click on a link on each operation, this will also tell you what targets the operation is applicable to.

a. Software Deployment: Enter the Destination Oracle Home and set a Start
 Schedule. The schedule you enter must reflect across all applicable selected targets.

You can select for the operation to start **Immediately** or **Later** at a time of your choosing.



If you select multiple databases running on different homes from the same host, the Fleet Management UI will only submit deployment for one database. This is because the Oracle Home path needs to be unique per source Oracle Home.

- b. Migrate Listener: Select a start time. Immediately/Auto the task will begin immediately or as soon as the previous task finishes. Later, enter a start time of your choosing that reflects the time usually taken for software deployment in your environment.
- c. Update Database: Select a start time. Immediately/Auto the task will begin immediately or as soon as the previous task finishes. Later, enter a start time of your choosing that reflects the time usually taken for software deployment in your environment.

Note:

All times are set in the Management Server Time Zone. You cannot set a Maintenance Task to run immediately if there are preceding tasks.

3. Enter the Credentials for Normal Host, Privileged Host and Database Credentials. You can select whether to use a Named or Preferred credential. Preferred credentials need to be set at the Oracle Home and target database before you can selecting them in the UI. For more information on preferred credentials in Fleet Maintenance see: Preferred Credential Support for Fleet Maintenance.

Note:

For Windows deployments use Windows Security Credentials.

- 4. Under **Options** you can set the following:
 - Work Directory: Enter the work directory where the files will be stored.
 - **Skip Cluster Verification Utility Checks**: Select to skip the Prerequisite Checks and/or Cluster Verification Utility Check.
 - Under Software Deployment:
 - Run CVU From Oracle Home: This will be set to default NO.
 - Custom CVU Component: Enter CVU component location in Software Library
 - Pre-staged Location for Root Scripts: Specify a read-only location to use prestaged dispatcher as well as fleet scripts
 - Data Guard Options: You can select Standby Auto Deploy for automatic deployment of software on standby host.
 - Migrate Listener allows you to set the Listener Options: Select whether or not to Force the listener.



Note:

This option is only available if Migrate Listener is selected.

- Update Database allows you to set the Data Guard Options with the following options:
 - Ignore Standby Prerequisite: Specify to disable verification check if the Standby is on the same image version where the Primary is being moved to.
 - Enable Data Guard: Specify if the standby is managed by Data Guard.
 - Enable FSFO: Enable Fast-Start Failover (FSFO). FSFO can be enabled when DG broker is on.
 - Startup Standby: Specify to startup the Standby databases after primary is patched. Used if you selected to keep the standby database down post patching. When selecting this, set the sandby startup option: Open, Mount, No Mount, Read Only, Read Only With Apply.
- Custom Scripts: Select the optional custom scripts you wish to run by clicking on Add, the Fleet Management UI is set to the default location for custom scripts. If your scripts are located in a different section Set the Stage, Script Location. Lastly select the check mark if you wish to Run as Root.

Once all options have been entered click Next.

- In the Review screen, verify all options entered in the previous screen.
 In this screen you can View JSON or Save JSON of the operations to be performed.
 - Click on **Validate** to perform a **Quick** or **Full** validation of your Fleet Maintenance Operation before submitting. Validation results are reported on a per target basis, for a **Full** validation a Deployment Procedure is submitted for each target. Each target will have a link available to see in detail the results of the Validate Deployment Procedure.
- 6. To start your Fleet Maintenance operation, click **Submit**.

Track, Fix and Resume Update Procedures

Oracle Enterprise Manager 13.5 allows you to Track, Fix, Pause and Resume Fleet Maintenance procedures directly from Enterprise Manager, allowing for greater oversight and control of all your Fleet operations.

Search for Fleet Maintenance Procedures

There are two methods to search for Fleet Maintenance procedures:

- If you are accessing after submitting a Fleet Maintenance Operations, click the Monitor Progress button, this will take you to the Procedure Activity page.
- If you are accessing from anywhere else in Enterprise Manager, navigate to the Procedure Activity page (Targets, then, Databases, then Provisioning and Patching, select Procedure Activity).

Track Operations

In the **Procedure Activity** page you can track in real time all ongoing Enterprise Manager procedures, Fleet Maintenance Operations will be flagged as **Fleet Maintenance Procedure**. Click on a specific Fleet Maintenance Operation to access detailed information for that activity. You can filter by Search Text Fields, Type, and Status.



Starting with Oracle Enterprise Manager 13.5 Release Update 17 (13.5.017) additional search filters have been added. You can further refine your deployment procedure activity searches with the following additional filters: Owner, Start Date From, and Start Date To.

Fix Operations

Fleet Operations that encountered issues will be presented with a red X mark, navigate to the **Procedure Activity** page. Click on the corresponding Fleet Maintenance Activity to open a new page with detailed information on said activity. Once you have reviewed the activity, click on the **Procedure Actions** button on the left hand side and select **Debug** to trouble shoot.

Resume Operations

In the **Procedure Activity** page, click on a specific Fleet Maintenance Operation. This will open a new page with detailed information for this activity. Click on the **Procedure Actions** button on the left hand side to access the available actions for said activity. You can perform Stop, Suspend, and Resume actions.

Use of Oracle Home Path Variables

Starting with Enterprise Manager 13.5 Release 15 (13.5.0.15) you can use predefined variables for an Oracle Home path, simplifying your EM CLI and REST API Fleet Maintenance operations.

The following is a list of variables available in Enterprise Manager 13.5 for Oracle Home path definitions. You can specify variable name using the format %<variable name>%.

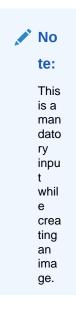
Table 18-1 Variable definitions for Oracle Home Path

Variable	Implementation	Comments		
ORACLE_BASE	Oracle Base location for the Database. Examples: Database deployment:	The Oracle Grid Infrastructure Oracle base should not contain the Oracle Grid Infrastructure binaries for an Oracle Grid Infrastructure for a cluster		
	%ORACLE_BASE%/19.0.0/dbhome	installation. Permissions for the file path to the Oracle Grid Infrastructure binary home are		
	Grid deployment:	changed to root during installation.		
	<pre>%ORACLE_BASE%//</pre>			
	19.0.0/grid			
OH_PARENT_DIR	Parent directory of the Oracle Home Examples: Database deployment:	Computed after retrieving the Oracle Home location for the target from EM repository using the property name OracleHome.		
	%OH_PARENT_DIR%/ dbhome_1			
	Grid deployment:			
	<pre>%OH_PARENT_DIR%/ gridhome</pre>			



Table 18-1 (Cont.) Variable definitions for Oracle Home Path

Variable	Implementation	Comments
VERSION	The VERSION variable will be replaced with Oracle Database base release value. Examples: Database deployment:	N/A
	%ORACLE_BASE%/ %VERSION%/dbhome	
	Grid deployment:	
	<pre>%ORACLE_BASE%// %VERSION%/gridhome</pre>	
VERSION_NAME	The VERSION_NAME variable will be replaced with Version Name value of that Image.	N/A



Examples:

Database deployment:

%ORACLE_BASE%/
%VERSION_NAME%/dbhome

Grid deployment:

%ORACLE_BASE%/../
%VERSION_NAME%/
gridhome



Table 18-1 (Cont.) Variable definitions for Oracle Home Path

Variable	Implementation	Comments	
U	This optional variable allows you to install the same product more than once in the same Oracle base directory. Examples: Database deployment:	This variable is to be used for bulk deployment. It recommended to use when there are several Homes in the same erallowing you to update single go.	t is this variable al Oracle nvironment
	%ORACLE_BASE%/ %VERSION%/dbhome%U%		
	Grid deployment:		№ No
	%ORACLE_BASE%// %VERSION%/grid%U%		te: Mult iple Ho me depl oym ent in par allel on the sam e host is curr entl y not sup port ed for Win dow s bas ed Ent erpr ise Man age r Dep loy men

Note:

For security purposes, the Oracle Home or Grid Home path folder should not contain binaries and data files of the existing home.

Table 18-2 Support for Variables in the Working Directory

Variables	Implementation	Comments
emStageDir	You can manually update the agent property file with the following commands:	N/A
	AGENT_HOME/agent_inst/ sysman/config/ emd.properties and run emctl reload	
	from	
	AGENT_HOME/ agent_13.5.0.0.0/bin	

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 and Post scripts for different operations is implemented. Pre and Post scripts need to be uploaded as EM Software Library Entity (Directive) prior to their usage with Fleet Maintenance operations. Pre and Ppost script usage is supported with DEPLOY, UPDATE, CLEANUP and ROLLBACK operations.

Creating and Using Custom Pre/Post Scripts:

- To develop your pre/post script, see: Custom Script Considerations.
- To upload custom scripts to the Software Library, see: Software Library Entity For Uploading The Script.
- Capture the URN for the Software Library Entity, see: How to Get URN for SWLIB Entity.
- 4. Specify the Pre and Post script Software Library URN entity in the EMCLI input file., see: Pre/Post Script 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 and 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.



that if there are swlib creds set, then they should be removed as we can run into privilege issues when we try to execute these scripts



Remove all Software Library credentials from your Fleet Maintenance Pre and Post scripts to avoid privilege issues during execution.

Additionally, you can run Pre and Post scripts as normal by setting the

RUN_PRE_SCRIPT_AS_ROOT=true and RUN_POST_SCRIPT_AS_ROOT=true flags respectively within the input_file. These flags are by default set to false and can be left out if not needed. These flags will only be honored if passed along with flags: CUSTOM_PRE_SCRIPTS_URN and CUSTOM_POST_SCRIPTS_URN respectively. The Pre and Post flags can be used for DEPLOY, UPDATE, ROLLBACK and CLEANUP Fleet Maintenance operations.



Do not edit Pre/Post hooks sections of Pre/Post script.

Sample input parameter file provided below:

Sample Code for CLEANUP_SOFTWARE

```
MAINTENANCE_PURPOSE=CLEANUP_SOFTWARE

CUSTOM_PRE_SCRIPTS_URN=oracle:defaultService:em:provisioning:1:cmp:COMP_Direct
ives:none:74A730047930C5FDE053DF0FC40A3E69:0.1

ORACLE_HOME_TARGET_DETAILS=abc01def.example.com:/scratch/cuser/app/cuser/
product/13.4.1/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

O['RACLE_HOME_ROOT_CRED_NAME=ORACLE_ROOT:SYSMAN

RUN_PRE_SCRIPT_AS_ROOT=true
```

Sample Code for DEPLOY_DB_SOFTWARE

```
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_Component:SUB_OracleDB:74AA263BE602CF3BE053DF0FC40A2ED6:0.1
ROOT_SCRIPT_LOC=/tmp/1311624666555972//
MASSDB-194177ce43d441b89a504584d146aafb
CUSTOM_PRE_SCRIPTS_URN=oracle:defaultService:em:provisioning:1:cmp:COMP_Directives:none:74A730047930C5FDE053DF0FC40A3E69:0.1
```



```
CUSTOM_POST_SCRIPTS_URN=oracle:defaultService:em:provisioning:1:cmp:COMP_Direc tives:none:74A730047935C5FDE053DF0FC40A3E69:0.1

ORACLE_HOME_TARGET_DETAILS=den01nre.example.com:OraDB12Home2_4_den01nre.exampl e.com_3012;

WORK_DIR_LOC=/tmp/1311624666555972//1311626758529303/

DEPLOY_TYPE=SIDB

CREDSET_TARGET_TYPE=OracleHome

ORACLE_HOME_ROOT_CREDSET_NAME=HostCredsPriv

RUN_POST_SCRIPT_AS_ROOT=true

NEW_ORACLE_HOME_LIST=/scratch/cuser/app/cuser/product/13.4.1/dbhome_Dep2
```

Software Library Entity For Uploading The Script

Create a Software Library entity of type Directive (for uploading the custom scripts to EM).

- 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.
- The INPUT_FILE to be added as mentioned in the list below. Click OK.
 - Argument Prefix="
 - Argument Suffix="
 - Property Name=INPUT_FILE
- Click Next.
- 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.
- 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:

- Select Enterprise > Provisioning and Patching > Software Library.
- 2. Click View > Columns > Internal ID. Select this option.
- 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.



Do not edit Pre/Post hooks sections.

```
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_dbOH.prop"

Contents of input_dbOH.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_Direct ives:none:74A730047930C5FDE053DF0FC40A3E69:0.1

CUSTOM_POST_SCRIPTS_URN=oracle:defaultService:em:provisioning:1:cmp:COMP_Direct tives:none:74A730047935C5FDE053DF0FC40A3E69:0.1
```

Fleet Maintenance Automation

Fleet Maintenance operations can also be automated using REST API or EM CLI.

To automate Fleet Maintenance operations see the following:

- Database Fleet Maintenance Update REST APIs
- Fleet Maintenance EM CLI Operations

Database Fleet Maintenance Update REST APIs

Database Fleet Maintenance operations can also be performed using REST API. To view all available REST APIs for Database Fleet Maintenance operations see: Database Maintenance-Updates (Patching) And Upgrades REST Endpoints.



Do not perform lifecycle management activities during a patching window, as the monitoring credentials are locked during patching and cannot be modified, created or deleted. Examples of the activities to avoid performing during a patching window: provisioning, patching, cloning, migration and similar activities from Self-service portal for DBaaS/PDBaaS.

Fleet Maintenance EM CLI Operations

With Fleet Maintenance you can perform various different operations on your entire fleet with simple EM CLI commands. Fleet Maintenance allows you to simplify standardizing software configurations across your environment.

General Flow for Fleet Maintenance EM CLI Operations

For a successful Fleet Maintenance operation it is recommended you follow these steps in the order prescribed in this section:

- First perform Image Maintenance tasks, such as: creating a Gold Image, creating a new version of the image, and setting the new version of an image to current.
- 2. Perform Subscription Tasks, such as: verifying image applicability, subscribing targets, and verifying the subscription.
- Once you have completed the image maintenance and subscription tasks, move forward to Deploy Operations.
- 4. With the images deployed, you will have to Migrate the Listeners.
- **5.** After you have performed the update operations, you can perform the Update Operations on your targets.
- 6. Should you encounter issues, you can perform Rollback Operations.
- 7. Oracle highly recommends after patching you perform Cleanup Operations.
- 8. For emergency one-off patching see: Emergency Patching

Note:

Do not perform lifecycle management activities during a patching window, as the monitoring credentials are locked during patching and cannot be modified, created or deleted. Examples of the activities to avoid performing during a patching window: provisioning, patching, cloning, migration and similar activities from Self-service portal for DBaaS/PDBaaS.

Available Target Types and Supported Operations

The following table shows the target types that are supported and the values you need to provide for the **target type** and **-purpose** fields in **EM CLI** commands:



Entity	Target Type	Deploy Operations	Update Operations	Rollback Operations	Cleanup Software Operation
Single Instance database	oracle_data base	deploy_db_soft ware	update_db migrate_lis tener	rollback_db	cleanup_softw are
	N o t e: AvailableforWindowsplatfor		For more information see: Migrate the Listeners		
	m s				

RAC database	rac_databas e	<pre>deploy_rac_sof tware</pre>	<pre>update_racd b</pre>	<pre>rollback_ra ckdb</pre>	<pre>cleanup_softw are</pre>
Oracle Restart	has	<pre>deploy_siha_so ftware</pre>	update_siha	<pre>rollback_si ha</pre>	<pre>cleanup_softw are</pre>
Grid Infrastructure	cluster	<pre>deploy_gi_soft ware</pre>	update_gi	rollback_gi	<pre>cleanup_softw are</pre>

Entity	Target Type	Deploy Operations	Update Operations	Rollback Operations	Cleanup Software Operation
PDB	oracle_pdb	deploy_db_soft ware deploy_rac_sof tware deploy_cdb attach_cdb	update_pdb	rollback_pd b	N/A
CDB/RAC CDB	base	deploy_db_soft ware deploy_rac_sof tware deploy_cdb attach_cdb	update_racd	_	

db software maintenance General Code Layout and Information

```
emcli db_software_maintenance -performOperation
-name=<User Specified Operation>
-purpose=<Operation Type>
-target_type=<target type>
-target_list=<list of targets>
-normal_credential=<credential name>
-privilege_credential=<credential name>
-rolling=<true/false>
-input file="data:<location of input file>"
```

- name: This is the unique name of the operation.
- purpose: Refer to the EM CLI Command Inputs Based on Entity Type table.
- target type: Refer to the EM CLI 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: 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.

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>.

For a detailed tutorial on patching a single instance database using Fleet Management see: Patching a Single Instance Database Using Oracle Enterprise Manager Fleet Maintenance.

Image Maintenance

The following is a list of image tasks that need to be performed for Fleet Maintenance.

- Create a Gold Image Using EMCLI
- · Create a New Version of the Image
- Retrieving a List of Available Gold Images
- Changing Version Status to Current
- Deleting an Image or Versions of an Image
- Migrating Fleet Maintenance Gold Images between Enterprise Manager Deployments

Create a Gold Image Using EMCLI

You can create Gold Images using EMCLI

To begin creating Gold Images using EMCLI, you will need to use the db software maintenance verb, see the following example:

```
emcli db_software_maintenance -createSoftwareImage -input_file="data:/tmp/
create gldimg.properties"
```

Where the <code>create_gldimg.properties</code> input file contains the variables explained below. For a complete description of input variables for <code>db_software_maintenance</code> family of verbs, see <code>db_software_maintenance</code> in the <code>Enterprise Manager Command Line Reference Guide</code>.

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 databases with the latest Release Update (RU), we need to specify the Oracle Home target name for a reference 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 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>

- 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. 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.



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.

- workingDirectory: 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, OMS Agent File system, or NFS (

Note:

NFS is available for Enterprise Manager deployments 13.5 Release Update 13 (13.5.0.13) and higher.

- STORAGE_NAME_FOR_SWLIB: The storage name for the Software Library. To retrieve the storage name: from the Enterprise menu,
 - 1. From the Enterprise menu select **Provisioning and Patching**, then select **Software Library**.
 - 2. In the **Software Library** page, from the **Actions** menu select **Administration**.
 - 3. Check the value in the **Name** column, this is the storage name for the Software Library.
- NFS_LOCAL_MOUNT: The absolute path for the NFS local mount, if this parameter is not specified the workingDirectory location will be used as default instead.
 NFS is available for Enterprise Manager deployments 13.5 Release Update 13 (13.5.0.13) and higher.



Tip:

This field is only required for NFS storage, it is not required for OMS Storage options.

• 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_TESTSUPERADMIN_12_08_2015_05_09_AM' and execution 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'

Creating Gold Images by Patching Oracle Homes

You can also create Gold Images by patching Oracle Homes. Once the patch list is identified and downloaded, you as administrator can create a Gold Image version of the in one of the following scenarios:

Create a new version from an existing deployment. As the administrator you have an
existing test environment in which has deployed all the required patches. You point to a
prepared Oracle home that contains all desired patches, and copies them into a new Gold
Image Oracle Home.

To begin, point to a prepared Oracle home that contains all desired patches:

```
emcli db_software_maintenance -createSoftwareImage -input_file="data:/home/
user/input file"
```

In this use case input file must contain the following:

```
IMAGE_ID=<ID of existing Gold Image in which a new version is to be
created>
REF_TARGET_NAME=<ORACLE_HOME source target>
IMAGE_SWLIB_LOC=<Image Location>
REF_HOST_CREDENTIALS=USER:TESTSUPERADMIN
workingDirectory=/u01/tmp
DESTN_SAME_AS_SRC=<True/False>
STORAGE_TYPE_FOR_SWLIB=OmsShared
STORAGE_NAME_FOR_SWLIB=swlib
PATCH_LIST=<List of patches to deploy, comma separated>
VERSION NAME=Version1
```

In this use case, you as the administrator do not want to the modify the test environment. For this use case you will have to clone the home, apply patches and then create a Gold Image. The original test environment home is NOT modified. In this use case, the task is to clone the test environment, apply the patches and then create the Gold Image. To begin, point to a prepared Oracle home that contains all desired patches:

```
emcli db_software_maintenance -createSoftwareImage -input_file="data:/home/
user/input file"
```

For this use case two examples for the $input_file$ will be given, one for RAC and one for GRID:

Input file for cloning, patching and creating Gold Image for a RAC home :

```
IMAGE_NAME=<Unique name to identify operation>
IMAGE_DESCRIPTION=<Description of image>
REF_TARGET_NAME=<Oracle_Home source target>
IMAGE_SWLIB_LOC=<Image Location>
```

```
REF_HOST_CREDENTIALS=USER:TESTSUPERADMIN
REF_HOST_ROOT_CREDENTIALS=USER_ROOT:TESTSUPERADMIN
workingDirectory=/u01/tmp
STORAGE_TYPE_FOR_SWLIB=OmsShared
STORAGE_NAME_FOR_SWLIB=sl
VERSION_NAME=<Patch version name>
PATCH_LIST=<List of patches to deploy, comma separated>
DESTN_HOME_LOCATION=<Location of the destination Oracle Home target>
```

• Input file for cloning, patching and creating Gold Image for a GRID Home:

```
IMAGE_NAME=<Unique name to identify operation>
IMAGE_DESCRIPTION=<Description of image>
REF_TARGET_NAME=<Oracle_Home source target>
REF_GI_CREDENTIALS=USER:TESTSUPERADMIN
IMAGE_SWLIB_LOC=<Image Location>
REF_HOST_CREDENTIALS=USER:TESTSUPERADMIN
REF_HOST_ROOT_CREDENTIALS=USER_ROOT:TESTSUPERADMIN
workingDirectory=/u01/tmp
STORAGE_TYPE_FOR_SWLIB=OmsShared
STORAGE_NAME_FOR_SWLIB=sl
VERSION_NAME=<Patch version name>
DESTN_HOME_LOCATION=<Location of the destination Oracle Home target>
PATCH_LIST=<List of patches to deploy, comma separated>
```

3. Modify the test environment itself. In this case, you as the administrator will have to apply patches to specified home and then create the Gold Image.
To begin, point to an Oracle home that will be patched with the new patches to create the new Gold Image

```
emcli db_software_maintenance -createSoftwareImage -input_file="data:/home/
user/input_file"
```

Input file where both source and destination will be modified:

```
IMAGE_NAME=<Unique name to identify operation>
IMAGE_DESCRIPTION=<Description of image>
REF_TARGET_NAME=<Oracle_Home source target>
IMAGE_SWLIB_LOC=<Image Location>
REF_HOST_CREDENTIALS=USER:TESTSUPERADMIN
REF_HOST_ROOT_CREDENTIALS=USER_ROOT:TESTSUPERADMIN
workingDirectory=/u01/tmp
STORAGE_TYPE_FOR_SWLIB=OmsShared
STORAGE_NAME_FOR_SWLIB=sl
VERSION_NAME=<Patch version name>
PATCH_LIST=<List of patches to deploy, comma separated>
DESTN_SAME_AS_SRC=true
```

4. Enterprise Manager deployments that have no exclusive Stage or Test environments. In this use case you will deploy the CURRENT (default) version of the image, apply patches and then create a Gold Image.

To begin, point to a prepared Oracle home that contains all desired patches:

```
emcli db_software_maintenance -createSoftwareImage -input_file="data:/home/
user/input file"
```

Input file sample for a no test environment use case:

```
IMAGE_ID=<ID of existing Gold Image in which a new version is to be
created>
DESTN_HOME_LOCATION=<Location of the destination Oracle Home target>
DESTN_HOST=<host target name>
DISPATCHER_LOCATION=<Dispatch location of the
scripts>
SKIP_PREREQUISITE_CHECKS=<true/false>
SKIP_CVU_CHECK=<true/false>
REF_HOST_CREDENTIALS=USER:SUPERADMIN
workingDirectory=/u01/tmp
IMAGE_SWLIB_LOC=<Image Location>
STORAGE_TYPE_FOR_SWLIB=OmsShared
STORAGE_NAME_FOR_SWLIB=swlib
VERSION_NAME=<Patch version name>
PATCH_LIST=<List of patches to deploy, comma separated>
```

Create a New 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 should be created for every patching cycle. This is the target version to which all the subscribed databases are to be patched in the next maintenance window. In order to create the new version, identify the list of patches to be included into the new image.

- 1. It should ideally be a super set of the CURRENT version, this ensures all existing patches are included in the new version. The command emcli db_software_maintenance describeImage -image_id=<Image Id> lists all the patches in the CURRENT version of the image.
- 2. Get the list of Oracle recommended patches.
- 3. Identify the one off patches applied to the subscribed databases out of sync.



It is considered best practice, that one-off patches be applied in emergency situations. This helps maintain standardization in your environment. For more information on emergency patching see: Emergency Patching.

The command emcli db_software_maintenance -checkApplicability - image_id=<image_id> -target_list=<target_list> -target_type=<target_type> provides information about all additional patches applied to the databases but not included in the CURRENT version.

- Get the list of patch requests from database owners.
- 5. From steps 3 and 4, identify the bug fixes that are part of the patch recommendations.

- For the rest of the pending bug fixes, identify the patches available from Oracle Recommended Patches and download them.
- All other pending issues, request patches and merged patches that meet your environments needs or from Support.
- 8. Create the new version using these patches. This will ensure that your new version is a super set of the CURRENT version and will meet the needs of all the subscribed databases. For more information on creating a new version see: Create a Gold Image Using EMCLI
- Once the Oracle home is created and tested, you may mark it CURRENT and begin your next maintenance cycle.



The same Gold Image can be used for deployments on both Oracle Red Hat Linux 7 and 8 deployments. The database media is identical in both versions, These binaries will be relinked during the deployment process.

Create a New Version of the Image using EM CLI

To create a new version of the image use the following EM CLI command:

```
emcli db_software_maintenance -createSoftwareImage -
input file=data:"<input file path>"
```

For a description of input file, see Create a Gold Image Using EMCLI.

Example:

```
emcli db_software_maintenance -createSoftwareImage -input_file="data:/home/
user/input rac"
```

Input File Example:

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/19.18.0/linux_x64
REF_HOST_CREDENTIALS=REF_HOST_CREDS:SYSMAN
WORKING_DIRECTORY=/tmp
STORAGE_TYPE_FOR_SWLIB=OmsShared
STORAGE_NAME_FOR_SWLIB=swlib
VERSION_NAME=PSU_19180_Version</pre>
```

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 execution guid='25F92D9A00164A45E053D903C40A9B4B'
```



Retrieving a List of Available Gold Images



To format the output of the EM CLI commands listed below, set the <code>emctl</code> property - <code>oracle.sysman.dbprov.gis.emcli.verbs.tableLength</code> to the width of the terminal. For example:

```
emctl set property -name
oracle.sysman.dbprov.gis.emcli.verbs.tableLength -value 160
```

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=<Comma-separated
variable list>]
Image Id Image Name Description
Version
72D9D7C656A11AB2E053BC3A41 RACS_19.3_RAC_19.3
19.19.0 19CRACS 0A7A43
```

Description of the Input Variables:

- -getImages: Shows a list of images with a Production status. To get a list of Failed or Inactive statuses use: -getImages
- columnName: 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. For example: -columnName: image_id,
 image_name

Sample Output

```
Image Id Image Name Description
Version Platform Name Creation Date Owner Modified By
FE55AD7AB28974EFE04313B2F00AD4A0 19.19.0 RAC DB Gold Image for 112040 RAC DB Homes
19.19.0 Linux x86-64 2023-12-28 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

 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

```
*****
**********
STATUS
         DATE CREATED
EXTERNAL ID
                   HASHCODE
*****************
                   277DF28F2D30393BE053D903C40AC6 CUSTOMER 19180
VERSION ACTIVE 2023-12-22 05:54:45.0
ORACLE: DEFAULTSERVICE: EM: PROVI C3448035451: B1400395227
SIONING:1:CMP:COMP COMPONENT:S 10
UB ORACLEDB:277DF28F2D2C393BE0
53D903C40AC610:0.1
******************
***********
2
                   277E39B74D684C6BE053D903C40A59 CUSTOMER 19190
VERSION CURRENT 2023-12-22 06:14:40.0
ORACLE: DEFAULTSERVICE: EM: PROVI C3448035451: B3881579444
IONING:1:CMP: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 EM CLI command. It shows the current version of the given image by default unless a specific version_id is provided.

```
emcli db_software_maintenance -describeImage -image_id=<Image Id> [-
version id=<version id>] [-bugs=true]
```

Sample Output

# PATCH#	PATCH DESCRIPTION
1 35042068	OCW Patch Set Update : 19.19.0 (35042068)
2 34765931	Database Patch Set Update: 19.18.0 (34765931)
3 34419443	Database Patch Set Update : 19.17.0 (34419443)

Changing Version Status to Current

After an updated version of the image has been created see Create a Gold Image Using EMCLI, 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.

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.
```

Deleting an Image or Versions of an Image

Delete a Version of an Image

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.

Delete an Image

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.

Migrating Fleet Maintenance Gold Images between Enterprise Manager Deployments

With Fleet Maintenance you can send updates, software images and their meta data across separate Enterprise Manager installations by migrating the Gold Images between fleet deployments. This use case is especially useful when having two Enterprise Manager installations, one for Development and Production.

You can test patching and software images on the development instance and once satisfied, move all the Software Images and their meta data to the production Enterprise Manager, avoiding having to recreate all the software images again for the production Enterprise Manager. Creating software images from installed homes can be a time-consuming affair.

When importing a Gold Image from an exported Gold Image from another Enterprise Manager patch, descriptions will appear within describeImage only if Enterprise Manager has patch information available. Patch descriptions are available once an Oracle home with the new patches is provisioned or discovered.

Steps to Migrate Fleet Maintenance Gold Images

Migrating a Fleet Maintenance Gold image, its software images and metadata, can be broken down into 3 main steps:

- 1. Export the Gold Image with the **Export Gold Image** command from the source Enterprise Manager Software Library to a staging location.
- 2. Optionally copy the Gold Image to a location accessible by the destination Enterprise Manager agent.
- Run Import Gold Image command to import the Gold Image files to the software library in the destination Enterprise Manager. The corresponding meta data written to the fleet bookkeeping tables and views.

Prerequisites for Exporting and Importing Fleet Maintenance Gold Images

Before starting to migrate Gold Images between Enterprise Manager installations, review the following:

- Make sure that the given origin and destination host are Enterprise Manager visible targets.
- Make sure that the destination host provided is a managed target within Enterprise Manager.
- Make sure the given storage location is valid.

Steps to Export and Import Fleet Maintenance Gold Image Fleet files

- First you will need to create the export gldimg.prop file with the following parameters:
 - IMAGE NAME: Name of the image needs to be exported
 - IMAGE ID: ID of the image needs to be exported



- VERSION NAME: Name of the image version needs to be exported
- VERSION ID: ID of the image version needs to be exported



If no version details are passed, the version with state as current will be exported. Version_Name or Version_ID are mandatory

DEST HOST NAME: Host name where to store the gold image bundle (Required)

- DEST HOST CRED: Host normal credential (Required)
- GOLD_IMAGE_BUNDLE_LOCATION: Location to store the gold image bundle (Required)
- GOLD IMAGE BUNDLE NAME: Name of the gold image bundle (Optional)
- workingDirectory: By default, /tmp is the working directory (Optional)

Destination host name, destination host credential, Image Id or Image Name, Gold Image Bundle Location are mandatory parameters. These parameters are checked at run time and will return an error if they are **Empty** or **Null**.

A sample of the export gldimg.prop file would look like this:

```
IMAGE_ID=8BBCE9315DEB1C27E05388947B0AF101
IMAGE_NAME=18.3 patched SI Home
VERSION_ID=8BBCC9E029AF7B18E05388947B0AB627
DEST_HOST_NAME=example.sample.com
DEST_HOST_CRED=USER1:SYSMAN
GOLD_IMAGE_BUNDLE_LOCATION=/scratch/<file_path>/goldimage
GOLD_IMAGE_BUNDLE_NAME=ExportGoldImage.zip
```

2. Once the export_gldimg.prop has been created run the db_software_maintenance EM CLI verb with the exportSoftwareImage flag. See the following example:

```
emcli db_software_maintenance -exportSoftwareImage
    -input_file="data:<file_path>/export_gldimg.prop"
```

Note:

It is recommended that you use short file paths on Windows Fleet Maintenance operations to avoid issues. You can pass a work directory location within the input file with a short file path like: workingDirectory=C:\temp.

Once completed, the Fleet Maintenance Gold Image has been successfully exported.

- 3. To start the import first you will need to create an import_gldimg.prop file with the following parameters:
 - IMAGE NAME: Gold Image name to be used for the imported image (Required).
 - HOST NAME: Host name where gold image bundle is stored (Required).
 - HOST CREDENTIAL: Host normal credential (Required).
 - GOLD IMAGE BUNDLE NAME: Gold image zip file bundle name (Required).
 - GOLD IMAGE BUNDLE LOCATION: Location of the gold image bundle (Required).
 - IMAGE DESCRIPTION: Imported Image description (Optional).
 - IMAGE SWLIB LOC: Image component location in the Software Library (Required).



- workingDirectory: By default, /tmp is the working directory, required if path is different (Optional).
- STORAGE TYPE FOR SWLIB: Storage type for the Software Library (Required).
- STORAGE NAME FOR SWLIB: Storage name for the Software Library (Required).
- VERSION NAME: Imported image version name (Optional).

Gold Image Bundle name, location, host name and credential, storage name, type and location, working directory are mandatory parameters. These parameters are checked at run time and will return an error if they are **Empty** or **Null**.

A sample of the import gldimg.prop file would look like this:

```
IMAGE_NAME=18.3 SI Home import1
HOST_NAME=example.sample.com
HOST_CREDENTIAL=USER1:SYSMAN
GOLD_IMAGE_BUNDLE_LOCATION=/scratch/<file_path>/goldimage183
GOLD_IMAGE_BUNDLE_NAME=ExportGoldImage183.zip
IMAGE_DESCRIPTION=Gold Image for 18.3 SI Home - import
IMAGE_SWLIB_LOC=Database ProvisioningProfiles/<db version>/linux_x64
workingDirectory=/tmp
STORAGE_TYPE_FOR_SWLIB=OmsShared
STORAGE_NAME_FOR_SWLIB=swlib
VERSION_NAME=18.3 SI home-import1
```

4. Once the import_gldimg.prop has been created run the db_software_maintenance EM CLI verb with the importSoftwareImage flag. See the following example:

```
emcli db_software_maintenance -importSoftwareImage
    -input file="data:<file path>/import gldimg.prop"
```

Once completed, the Fleet Maintenance Gold Image has been successfully imported.

Export and import a Gold Image with two different version

- 1. First you will need to create the export_apps_1917.prop file, this file will be used to export both images with the following parameters:
 - IMAGE NAME: Name of the image needs to be exported.
 - HOST NAME: Host name to be exported (Required).
 - HOST CREDENTIAL: Host normal credential (Required).
 - GOLD IMAGE BUNDLE LOCATION: Location to store the gold image bundle (Required).
 - GOLD IMAGE BUNDLE NAME: Name of the gold image bundle (Required).
 - IMAGE SWLIB LOC: Location of the image within the Software Library.
 - workingDirectory: By default, /tmp is the working directory (Optional).
 - STORAGE TYPE FOR SWLIB: Storage type defined in the software library (Required).

- STORAGE NAME FOR SWLIB: Location in the Software Library (Required).
- VERSION NAME: Name of the image version needs to be exported (Required).



Mandatory parameters are labeled as **Required**, these parameters are checked at run time and will return an error if they are **Empty** or **Null**.

A sample of the export apps 1917.prop file would look like this:

```
IMAGE_NAME=19cDB-Linux-x64-Apps
HOST_NAME=emcc.marketplace.com
HOST_CREDENTIAL=ORACLE:SYSMAN
GOLD_IMAGE_BUNDLE_LOCATION=/home/oracle/Downloads/
GOLD_IMAGE_BUNDLE_NAME=ExportAPPS_1917.zip
IMAGE_DESCRIPTION=19c DB Gold Image
IMAGE_SWLIB_LOC=Fleet Maintenance/
WORKING_DIRECTORY=/u01/tmp
STORAGE_TYPE_FOR_SWLIB=OmsShared
STORAGE_NAME_FOR_SWLIB=default_loc
VERSION_NAME=v19.17DBRU
```

2. Once the export_apps_1917.prop has been created run the db_software_maintenance EM CLI verb with the exportSoftwareImage flag. See the following example:

```
emcli db_software_maintenance -exportSoftwareImage
    -input file="data:<file path>/export apps 1917.prop"
```

Note:

It is recommended that you use short file paths on Windows Fleet Maintenance operations to avoid issues. You can pass a work directory location within the input file with a short file path like: workingDirectory=C:\temp.

Once completed, the Fleet Maintenance Gold Image has been successfully exported.

3. Create the gold image from the export_apps_1917.prop file by running the following command:

```
emcli db_software_maintenance -importSoftwareImage -input_file="data:/home/
oracle/fleet/export apps 1917.prop"
```

This will create the Gold Image with a 19.17 version.

- 4. To add a 19.18 version into this Gold Image, you will need to additionally pass the 19.18 image ID. Locate the image ID and create a new input file named <code>ExportAPPS_1918.inp</code> with the following parameters:
 - IMAGE ID: ID of the image needs to be exported.

For this portion of the example IMAGE ID is mandatory.

- HOST NAME: Host name to be exported (Required).
- HOST CREDENTIAL: Host normal credential (Required).
- GOLD IMAGE BUNDLE LOCATION: Location to store the gold image bundle (Required).
- GOLD IMAGE BUNDLE NAME: Name of the gold image bundle (Required).
- IMAGE SWLIB LOC: Location of the image within the Software Library.
- workingDirectory: By default, /tmp is the working directory (Optional).
- STORAGE_TYPE_FOR_SWLIB: Storage type defined in the software library (Required).
- STORAGE NAME FOR SWLIB: Location in the Software Library (Required).
- VERSION NAME: Name of the image version needs to be exported (Required).

A sample of the ExportAPPS 1918.inp file would look like this:

```
IMAGE_ID=FC6A1AA82CB23A62E05326000000AB0CE
HOST_NAME=emcc.marketplace.com
HOST_CREDENTIAL=ORACLE:SYSMAN
GOLD_IMAGE_BUNDLE_LOCATION=/home/oracle/Downloads/
GOLD_IMAGE_BUNDLE_NAME=ExportAPPS_1918.zip
IMAGE_SWLIB_LOC=Fleet Maintenance/
workingDirectory=/u01/tmp
STORAGE_TYPE_FOR_SWLIB=OmsShared
STORAGE_NAME_FOR_SWLIB=default_loc
VERSION_NAME=v19.18DBRU
```

5. To add the 19.18 version into the Gold Image run the following command:

```
emcli db_software_maintenance -importSoftwareImage -input_file="data:/home/
oracle/fleet/ExportAPPS_1918.inp"
```

At this point the Gold Image will contain both 19.17 and 19.18 versions.

Subscription Tasks

The following is a list of subscription tasks that need to be performed for Fleet Maintenance.

EMCLI Operations:

- Verifying the Subscription
- Subscribing the Targets to the Selected Image
- Unsubscribing Targets from an Image
- Create a New Version of the Image



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.

Sample code:

```
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 "EM CLI command inputs based on Entity type".

Examples:

RAC Database

```
emcli db_software_maintenance -checkApplicability -image_id=
FE55AD7AB28974EFE04313B2F00AD4A0
-target_list=RACDB1 -target_type=rac_database
```

Oracle Restart

```
emcli db_software_maintenance -checkApplicability -
image_id=2EA08B6339234B7E053BC3A410A2826
-target list=HAS1 -target type=has
```

Grid Infrastructure

```
emcli db_software_maintenance -checkApplicability -
image_id=72F8E4FCDE017B99E053BC3A410AF226
-target list=CLUSTER1 -target type=cluster
```

Single Instance

```
emcli db_software_maintenance -checkApplicability -
image_id=6E4EA63DDC00535CE0531636B10ABA1A
-target list=DB1 -target type=oracle 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 uptake the same.

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.

Syntax:

```
emcli db_software_maintenance -subscribeTarget -target_name=<target_name>
-target type=<target type> -image id=<image id>
```

Sample:

```
emcli db_software_maintenance -subscribeTarget -
target_name=ORCL1910.abc.sample.com
-target type=oracle database -image id=A344EFAD4A9D7
```

Where:

- target name: The name of the target that needs to be patched.
- target_type: Type of database target to be patched, for more information on target types
 see: Fleet Maintenance EM CLI Operations.
- 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.



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
```

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 an image and a target subscribed:

Syntax:

emcli db_software_maintenance -getImageSubscriptions -image_id=<image_id>

Sample Output:

```
TARGETNAME TARGETTYPE VERSIONNAME DATESUBSCRIBED

CLUSTER1 CLUSTER 19180 Version 2023-12-22 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: Type of target, for example database"
- 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_19190 48B78E3967D56BE PSU Jan23 <null> <null> Jan 2023-02-21

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



Unsubscribing Targets from an Image

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 database target that needs to be removed.
- target type: Refer to the table "EM CLI command inputs based on Entity type".
- image_id: The ID of the gold image to which this target is to be removed. You can get the
 image_id by running the emcli command. For information about the command, see
 Retrieving a List of Available Gold Images.

Deploy Operations

Deploy operations are performed within Fleet Maintenance, to simplify databases deployments. They 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. In this topic we cover several examples of different deployment operations.

For more information on deploy operators within the EM CLI verb **db_software_maintenance** see db_sm_performoperation_deploy in *Enterprise Manager Command Line Interface*.

The general layout of a Deploy instruction is as follows:

```
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=<Named Credential: Credential Owner>
-privilege_credential=<Named Credential: Credential Owner>
[-windows_sec_credential=<Named Credential: Credential Owner>]
-input_file="data:<location of input file>"
[-standbyAutoDeploy=true|false] [-procedure name prefix="CustomUpdate"]
```

- name: This is the unique name of the operation.
- purpose: The purpose of the instructions, in this case DEPLOY:
 - DEPLOY DB SOFTWARE
 - DEPLOY RAC SOFTWARE
 - DEPLOY SIHA SOFTWARE
 - DEPLOY GI SOFTWARE
 - DEPLOY CBD



.

- target_type: The type of target on which this operation is being performed. Refer to the table EM CLI 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: Unix only, 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: **Unix only**, 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.

Note:

These credentials are used to run scripts as root.

- windows_sec_credential: **Windows Single Instance database only**, 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.

Note:

Starting with Enterprise Manager 13.5 Release Update 2, three different credentials are supported for Windows Single Instance database: Normal/Local account, Virtual account and Built in account.

Windows examples:

Normal/Local account:

```
emcli db_software_maintenance -performOperation -name="Deploy a Patched
Home"
-purpose=DEPLOY_DB_SOFTWARE -target_type=oracle_database -
target_list="ORCL1910.abc.sample.com"
-input file="data:/tmp/deploywin.properties" -
```



```
normal_credential="LOCALADMIN:SYSMAN"
-windows sec credential="LOCALUSERCRED:SYSMAN"
```

Virtual account:

```
emcli db_software_maintenance -performOperation -name="Deploy a Patched
Home"
-purpose=DEPLOY_DB_SOFTWARE -target_type=oracle_database -
target_list="ORCL1910.abc.sample.com"
-input_file="data:/tmp/deploywin.properties" -
normal_credential="LOCALADMIN:SYSMAN"
-is_virtual_account="true"
```

Built in account:

```
emcli db_software_maintenance -performOperation -name="Deploy a Patched
Home"
-purpose=DEPLOY_DB_SOFTWARE -target_type=oracle_database -
target_list="ORCL1910.abc.sample.com"
-input_file="data:/tmp/deploywin.properties" -
normal_credential="LOCALADMIN:SYSMAN"
-is_builtin_account="true"
```

- start_schedule: The date on which the stage and deploy is to be started if it is to be started in the future with the following format: "start_time:yyyy/mm/dd HH:mm". 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 EM CLI
 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 and write access to this location.
 - When deploying with NFS, NFS uses the working directory as the local mount location. NFS is available with Oracle Enterprise Manager 13.5 Release Update 13 (13.5.0.13) and higher.
 - homeNamePrefix=<User defined Oracle Home name>: Allows you to define a custom Oracle Home name prefix, allowing for the ability to discover Oracle Homes with custom name post deployment. If the home name prefix is not specified, the default value will be used. This variable is available starting with Oracle Enterprise Manager 13.5 Release update 20 (13.5.0.20).
 - 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".
- dispatchLoc=<Absolute path to the custom location>: This parameter has been added for operations that stage and use a dispatcher location, using an existing, readonly location to use pre-staged dispatcher as well as fleet scripts. This is an added security measure to only execute as root through a custom script which is only allowed to execute a reduced number of files.
- isRootPreStaged=<true/false>: If this flag is set to true the root components need to be pre-staged into the dispatchLoc location. This is helpful in skipping the pre-staged step when the components have already been pre-staged manually.
 In order to pre-stage root, the following components need to be previously staged and submitted:
 - Root Dispatcher component (Unix only requirement)
 - * Root Scripts, requires an additional step of unzipping the scripts in the dispatchLoc location.
 - * Common Root component

For detailed instructions on how to pre-stage the Root Components, see Manually Staging the Root Components For Fleet Maintenance (Patching).



If the flag is set to false and the components pre-staged, the components will be overwritten

- RUN_CVU_FROM_OH=<true/false>: The Cluster Verification Utility runs by default from the Oracle Home. Setting this flag to false will allow you to run the CVU from the Software Library used in conjunction with CUSTOM_CVU_COMPONENT. This flag will also be set to false in the following two cases:
 - * Database version is 12.1 or below, Cluster Verification Utility is not part of oracle home binaries for these versions.
 - * Standalone SIDB (Single Instance Database), as Cluster Verification Utility is not part of SIDB oracle home binaries.
- CUSTOM_CVU_COMPONENT=<Path of Custom CVU in the Software Library>: This component used in conjunction with RUN_CVU_FROM_OH=false allows you to run the CVU from the Software Library instead of the Oracle home by setting a path from where your custom CVU is located.

Deploy an Oracle Home

The following example goes over the deployment of an Oracle Home. The DEPLOY_DB_SOFTWARE purpose can be used for the creation of a Single Instance Oracle Database and Pluggable Database:

emcli db_software_maintenance
-performOperation
-name="Deploy Home"



```
-purpose="DEPLOY_DB_SOFTWARE"
-target_type= <oracle_database or oracle_pdb>
-target_list=db1916
-normal_credential="NORMAL:SYSMAN"
-privilege_credential="ROOT:SYSMAN"
-input_file="data:/deploy_OH.txt"
-attach home=true
```

For Container databases and RAC Container database creation use verb DEPLOY CDB, instead of DEPLOY DB SOFTWARE.

Deploy RAC Database

The following code snippet shows how to deploy an RAC Database using the DEPLOY_RAC_SOFTWARE purpose within the db_software_maintenance emcli verb

```
emcli db_software_maintenance
-performOperation
-name="Deploy Home"
-purpose="DEPLOY_RAC_SOFTWARE"
-target_type= rac_database
-normal_credential="NORMAL:SYSMAN"
-privilege_credential="ROOT:SYSMAN"
-input_file="data:/deploy_OH.txt"
-attach home=true
```

Deploy an Oracle Restart

```
emcli db_software_maintenance
-performOperation -name="Deploy Home"
-name="Deploy Home"
-purpose="DEPLOY_SIHA_SOFTWARE"
-target_type= has
-normal_credential="NC_HOST_CREDS:SYSMAN"
-privilege_credential="HOST_PRIV:SYSMAN"
-input_file="data:/deploy_OH.txt"
-SKIP_CVU_CHECK=<true/false>
```

SKIP_CVU_CHECK: This flag allows skipping the CVU check if set to true. This is an optional flag only used for DEPLOY SIHA and DEPLOY GI operations.

Deploy an Oracle Grid Infrastructure Database

```
emcli db_software_maintenance
-performOperation
-name="Deploy-1916 GI Home"
-purpose=DEPLOY_GI_SOFTWARE
-target_type=input_file
-target list="CLUSTER1"
```



```
-normal_credential="NC_HOST_CREDS:TESTSUPERADMIN"
-privilege_credential="HOST_PRIV:TESTSUPERADMIN"
-inputfile="data:/usr/oracle/deploy.txt"
SKIP CVU Check=<true/false>"
```

SKIP_CVU_CHECK: This flag allows skipping the CVU check if set to true. This is an optional flag only used for DEPLOY SIHA and DEPLOY GI operations.

Use DEPLOY_GI_SOFTWARE if you are planning to upgrade your Grid Infrastructure from 12.X to 19C or 18C. Once the new 19C or 18C home is deployed, perform UPDATE_GI. Finally use the DEPLOY RAC SOFTWARE to deploy the RAC home binaries of 19C or 18C

Deploy GIMR Software



From Oracle Database 21C onward, the GIMR update is now separated from GI update or patching. Performing an update of both GI and GIMR will happen in 2 separate operations. Only use this command on 21C or higher databases.

As a prerequisite to deploying the GIMR software, GI has to be updated beforehand. Once GI has been updated follow these steps:

- Create Gold Image: Create a Gold image for RAC home target or SI home target (an SI Oracle home image can be used to deploy RAC software).
- 2. Subscribe the GIMR Target to the Image: GIMR is not discovered in Enterprise Manager as a separate target. You will need to subscribe the cluster target to be configured with GIMR to a RAC or SI home image. To do this use the following command:

```
emcli db_software_maintenance
-subscribeTarget
-target_name="GIMR-target-cluster"
-target_type=cluster
-image_id=<RAC-SI_IMAGE_ID>
-qimr=true
```

3. Deploy the GIMR Software:

```
emcli db_software_maintenance
-performOperation
-name="Deploy GIMR Home"
-purpose=DEPLOY_GIMR_SOFTWARE
-target_type=cluster
-target_list=den01nr-cluster
-normal_credential="NC_HOST_CREDS:TESTSUPERADMIN"
-privilege_credential="HOST_PRIV:TESTSUPERADMIN"
-inputfile="data:/usr/oracle/deploy.txt"
```

Deploy Container Databases

In case of container databases, users have flexibility to update one or more PDB 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 PDB 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.



A PDB target may also be patched using an existing Container Database at a higher patch level, in this case <code>ATTACH_CDB</code> operation must be performed and <code>DEPLOY_CDB</code> operations must be skipped.

```
emcli db_software_maintenance -performOperation -name="<name>"
-purpose=DEPLOY_CDB
-target_type=rac_database
-target_list="<CDB Name that is being patched>"
-db_name="<db_name>"
-normal_credential="<normal_credential>"
-privilege_credential="<privilege_credential>"
-database_credential="<database_credential>"
-input_file="<input_file>"
-templateName="<template_path_and_name>"
-dataFileLocation="<dataFileLocation>"
```

If a container database already exists with the required image, then the same can be attached for housing the PDB 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"
-destinationCDB="<Container database to which PDBs will be migrated>"
-normal_credential="<credential name>"
-privilege_credential="<credential name>"
-database_credential="<SYSDBA credential name>"
-templateName="<template_path_and_name>"
-dataFileLocation="<dataFileLocation>"
```

Consolidate Oracle Homes

You can now use the EM CLI **attach_home** flag to merge homes of databases running from different Oracle Homes, on the same host. A second Oracle Home can be merged to a previously deployed Oracle Home if the two homes are patched with the same version of the Gold Image.



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 deploysoftware and deploy-CDB operations must be skipped.

User can pass the <code>attach_home=true</code> flag and provide the Oracle Home location within the input file, where an Oracle Home with the same Gold Image version has already been deployed. This will skip the actual deployment and add the required lineage to the new Oracle Home.

```
emcli db_software_maintenance
-performOperation
-name="Deploy Home"
-purpose="DEPLOY_DB_SOFTWARE"
-target_type=oracle_database -target_list=db1916
-normal_credential="NORMAL:SYSMAN" -privilege_credential="ROOT:SYSMAN"
-input file="data:/scratch/deploy OH.prop" -attach home=true
```

In order for the attach_home=true flag to execute successfully, two validations will be performed before moving Databases from Oracle Homes:

- If a user provides attach_home=true flag, but the deployed version does not match between Oracle Homes, the deployment will not submit and the following message will appear: A different version is deployed at the given location, you can attach to the same version of an image deployed.
- If a user provides <code>attach_home=true</code> flag, for a location where there is no deployed Oracle Home using Fleet Management; the deployment will not submit and the following message will appear: You can attach to the same version of deployed home at the given location.

For more information on EM CLI verb **db_software_maintenance** see db software maintenance in *Enterprise Manager Command Line Interface*.

Reorganizing your Oracle Homes can be useful in grouping databases according to location (same host) and version (same Gold Image). By consolidating Oracle Homes using <code>attach_home=true</code> flag you can simplify how patches are applied throughout an entire host, helping keep all databases up to date in the most efficient manner. Consolidating databases within a host can also help free space, once databases are merged into a common Oracle Home, empty Oracle Homes can be cleaned up using the Oracle Home cleanup EM CLI string. For more information see Cleanup Operations.

Redeploy Oracle Homes

Should an Oracle Home become corrupted or removed completely from the host, you can redeploy an Oracle Home again on the same host. This feature is available starting with Oracle Enterprise Manager 13.5 Release Update 18 (13.5.0.18).

Before You Begin

Make sure that the Oracle Home that you are redeploying is not associated with any additional targets in Enterprise Manager before starting a redeploy. The redeploy process will fail at the validations step if it encounters additional Enterprise Manager targets associated to a home. Corrupt Oracle Homes can not be cleaned up using the automated Fleet Maintenance Cleanup Operations, and need to be manually removed from the host before starting a redeploy operation.

Redeploy an Oracle Home

To redeploy an Oracle Home using EM CLI use this sample code:

```
emcli db_software_maintenance
-performOperation
```



```
-name=<REDEPLOY_DBHOME>
-purpose=DEPLOY_DB_SOFTWARE
-target_type=oracle_database
-target_list=<list of targets>
-normal_credential="<CREDENTIAL>:<OWNER>"
-privilege_credential="<CREDENTIAL>:<OWNER>"
-input file="<INPUT FILE LOCATION>"
```

Where the input file must contain the following additional parameters:

```
NEW_ORACLE_HOME_LIST=<NEW ORACLE HOME PATH>
workingDir=<FULL WORKING DIRECTORY PATH>
SKIP_PREREQUISITE_CHECKS=true
SKIP_CVU_CHECK=true
IS REDEPLOY=true
```

To redeploy an Oracle Home the IS_REDEPLOY parameter must be set to **true**. Alternatively, a redeployment operational parameter can also be passed as a property value as seen in the example below:

```
emcli db_software_maintenance
-performOperation
-name=<REDEPLOY_DBHOME>
-purpose=DEPLOY_DB_SOFTWARE
-target_type=oracle_database
-target_list=<list of targets>
-normal_credential="<CREDENTIAL>:<OWNER>"
-privilege_credential="<CREDENTIAL>:<OWNER>"
-input_file="<INPUT_FILE_LOCATION>"
-is_redeploy="true"
```

Where the input file will contain:

```
NEW_ORACLE_HOME_LIST=<NEW ORACLE HOME PATH>
workingDir=<FULL WORKING DIRECTORY PATH>
SKIP_PREREQUISITE_CHECKS=true
SKIP_CVU_CHECK=true
```

Migrate 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.



If your agent environment has TNS_ADMIN environment variable set, you may encounter an issue where the listener is not started in a new home after a listener migration. To avoid this insure the variable is removed from agent environment at the time of a listener migration. For more information on the TNS_ADMIN environment variable see: TNS_ADMIN Environment Variable.

Sample:

```
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, use format: start_time:yyyy/mm/dd HH:mm. This is an optional parameter. If no date is provided, the fleet operation will start immediately.
- procedure_name_prefix: OPTIONAL use this if the listener name exceeds the character limit, this allows for you to enter a deployment procedure name to ease tracking of the listener rollback procedure.

For Windows Single Instance database listener migrations use windows_sec_credential="Named Credential: Credential Owner" instead of privilege_credential.

Example:

```
emcli db_software_maintenance -performOperation -name="Update
Listener" -purpose=migrate_listener
-target_type=oracle_database -target_list="ORCL1910.abc.sample.com"
-normal_credential="LOCALADMIN:SYSMAN" -
windows_sec_credential="LOCALUSERCRED:SYSMAN"
```

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
```

Rollback Operations

The Rollback command is used to switch the database back to the previous Oracle Home after the Update operation has been completed.

Note:

- Before performing a rollback operation, if your listener is in a custom location, be sure to first set the TNS_ADMIN environment variable to this location. Next, restart the agent and then proceed with these operations.
- If during a database or RAC rollback you need to update any **Static Entry** and your listener.ora is located in the Grid Oracle Home default location (GI_HOME/network/admin) normal credentials will not work (for example DB Home Owner). This is because the owner of the file is the Grid Home owner. Make sure the file has permissions to at least group level to be modified by both Database and Grid Oracle Home Users.
- Starting with Enterprise Manager 13.5 Release update 19 (13.5.0.19), before and after a UPDATE/ROLLBACK operation on single instance or RAC databases, the deployment procedure activity logs include the number of invalid objects found.

Sample Code for ROLLBACK_DB

```
emcli db_software_maintenance
-performOperation
-name="Rollback DB"
-purpose=ROLLBACK_DB
-target_type=<target type>
-target_list=<list of targets>
-normal_credential="NC_HOST_CREDS:SYSMAN"
-privilege_credential="HOST_PRIV:SYSMAN"
-rolling=[true/false]
-node list="host1.example.com"
```

- 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".

```
Note:
  For Windows Single Instance database rollback operations use
  windows sec credential="Named Credential:Credential Owner" instead of
  privilege credential. Starting with Enterprise Manager 13,5 Release Update 2
  you can also use either -is virtual account="true" or -
  is builtin account="true" accounts.
  Examples:
  emcli db software maintenance -performOperation -name="Rollback SIDB"
  -purpose=ROLLBACK DB
  -target type=oracle database -target list="ORCLTDE.abc.example.com" -
  rolling=false
  -normal credential="LOCALADMIN:SYSMAN" -
  windows sec credential="LOCALUSERCRED:SYSMAN"
  -database credential="SYSDBA:SYSMAN"
  emcli db software maintenance -performOperation -name="Rollback SIDB"
   -purpose=ROLLBACK DB
  -target_type=oracle_database -target_list="ORCLTDE.abc.example.com" -
  rolling=false
  -normal credential="LOCALADMIN:SYSMAN" -is virtual account="true"
  -database credential="SYSDBA:SYSMAN"
  emcli db software maintenance -performOperation -name="Rollback SIDB"
   -purpose=ROLLBACK DB
  -target type=oracle database -target list="ORCLTDE.abc.example.com" -
  rolling=false
  -normal credential="LOCALADMIN:SYSMAN" is builtin account="true"
   -database credential="SYSDBA:SYSMAN"
```

Sample Code for ROLLBACK_RACDB

```
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_credential="HOST_PRIV:SYSMAN"
-rolling=[true/false]
-node list="host1.example.com"
```

Sample Code For ROLLBACK_GI

```
emcli db_software_maintenance
-performOperation
-name="Rollback GRID"
-purpose=ROLLBACK_GI
-target_type=cluster
-target_list="sample01-cluster"
-normal_credential="NC_HOST_CREDS:SYSMAN"
-privilege_credential="HOST_PRIV:SYSMAN"
-SKIP_CVU_CHECK=<true/false>
```

SKIP_CVU_CHECK: This flag allows skipping the CVU check if set to *true*. This is an optional flag only used for ROLLBACK GI operations.

Sample Code for ROLLBACK_SIHA

```
emcli db_software_maintenance
-performOperation
-name="Rollback"
-purpose="ROLLBACK_SIHA"
-target_type=has
-target_list="has_abc"
-normal_credential="NC_HOST_CREDS:SYSMAN"
-privilege_credential="HOST_PRIV:SYSMAN"
-SKIP_CVU_CHECK=<true/false>
```

SKIP_CVU_CHECK: This flag allows skipping the CVU check if set to *true*. This is an optional flag only used for ROLLBACK GI operations.

Rolling Back Of a Failed Update Operation

Should a single instance or RAC database update operation fail use the following steps to perform a rollback to its initial state.

Rolling Back a Failed Update Operation

Before you consider performing a rollback of a failed update verify if the update procedure was successful or failed. You can verify via the UI by following these instructions: Track, Fix and Resume Update Procedures.

To rollback a failed update operation use the following EM CLI command:

```
emcli db_software_maintenance
-performOperation -name="Rollback failed DB update"
-description="Rollback DB"
-purpose=<ROLLBACK_DB|ROLLBACK_RACDB>
-target_list=<DB name>
-target_type=<DB type>
-normal_credential=<HOST_CRED>:<CRED_OWNER>
-privilege_credential=<HOST_ROOT_CRED>:<CRED_OWNER>
-force=true
```

 purpose: There are standard purposes that can be performed by Fleet Operations which can be:

- ROLLBACK DB
- ROLLBACK RACDB
- 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.
- 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.

• force: This is mandatory parameter for failed update cases requiring a rollback.

Example

This is a sample code for a rollback operation of a failed update:

```
emcli db_software_maintenance
-performOperation
-name="Rollback RAC DB"
-purpose=ROLLBACK_RACDB
-target_type=oracle_database
-target_list="ORCLTDE.abc.example.com"
-normal_credential="NC_HOST_CREDS:SYSMAN"
-privilege_credential="HOST_PRIV:SYSMAN"
-force=true
```

For PDB relocation and updates that fail before the patching step or within the patching step there is no action to be taken as the update will not have taken place. If the failure occurs after the patching step, rollback is required. Follow the steps outlined under Rollback Operations.

If the update was successful and you wish to perform a rollback, use the rollback steps outlined under Rollback Operations.

Non-Rolling Applicable Patches

If the gold image contains a patch that is non-rolling applicable (For example 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.

Update and Rollback 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).



Before performing an **UPDATE** or **ROLLBACK** operation, if your listener is in a custom location, be sure to first set the **TNS_ADMIN** environment variable to this location . Next, restart the agent and then proceed with these operations.

The STANDBY_START_OPTION parameter sets the standby database start option after patching. This parameter should be included in <code>-input_file</code> while patching the standby database with the <code>UPDATE DB EM CLI verb</code>.

- If STANDBY_START_OPTION is provided while patching standby databases:

 Both rolling and non rolling standby databases will be started in the specified mode set in the -input_file after patching. If the parameter value is set as READ ONLY WITH APPLY, MRP will be started up, even if the MRP was not running before patching. Otherwise, the MRP status will enabled as set in the srcvtl configuration.
- If STANDBY START OPTION is not provided while patching standby database:
 - Rolling applicable patch: The standby database's existing open_mode and MRP status will be retained after patching. If unable to fetch the existing open_mode and MRP status for the standby database, it will be started in default mode and MRP status will be enabled as set in the srcvtl configuration.

Note:

The <code>UPDATE_DB</code> procedure will not stop, if any issue is detected when fetching the <code>open_mode</code> and MRP status. An error message will be shown in the <code>Start DB</code> Instance from destination home step in the <code>UPDATE_DB</code> procedure.

 Non-Rolling applicable patch: Standby database will restart in a default mode and MRP status will be enabled as set in the srcvtl configuration.

Update a Standby Database

For rolling applicable patches like database Release Updates, follow these steps:

Update the standby database:

```
emcli db_software_maintenance
-performOperation
-name=<User Specified Operation Name>
-purpose=<UPDATE_DB or UPDATE_RACDB>
-target_type=oracle_database
-target_list=<List of targets>
-normal_credential=<credential name>
-privilege_credential=<Credential name>
-database credential=<SYSDBA credentials of the database>
```



-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: The location and name of the temporary directory.
 - * disableDG: True/False, default value is false. If the standby database is managed by Data Guard this needs to be set to true. Setting to true enables the data guard configurations..
- Update the primary database:

```
emcli db_software_maintenance
-performOperation
-name=<User Specified Operation Name>
-purpose=-purpose=<UPDATE_DB or UPDATE_RACDB>
-target_type=oracle_database
-target_list=<List of targets>
-normal_credential=<credential name>
-privilege_credential=<credential name>
-dataguard_role=primary
[-ignoreStandbyPrereq=true|false]
-input_file="data:<location of input file>"
[] Denotes optional parameter
```

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: The location and name of the temporary directory.
 - * enableDG: True/False, default value is false. If the standby database is managed by Data Guard this needs to be set to true. Setting to true enables the data guard configurations.

For non-rolling applicable patches such as OJVM (pre-Jan 2017 RUs), follow these steps:

Update the standby database

```
emcli db_software_maintenance
-performOperation
-name=<User Specified Operation Name>
-purpose=<UPDATE_DB or UPDATE_RACDB>
-target_type=oracle_database
-target_list=<List of targets>
-normal_credential=<credential name>
-privilege_credential=<credential name>
-dataguard role=standby
```



-startupAfterSwitch=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: The location and name of the temporary directory.
 - * disableDG: True/False, default value is false. If the standby database is managed by Data Guard this needs to be set to true. Setting to true enables the data guard configurations.
- Update the primary database with dataguard role=Primary

```
emcli db_software_maintenance
-performOperation
-name=<User Specified Operation Name>
-purpose=-purpose=<UPDATE_DB or UPDATE_RACDB>
-target_type=oracle_database
-target_list=oracle_database
-normal_credential=<credential name>
-privilege_credential=<credential name>
-dataguard_role=primary
-rolling=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 and location of the temporary directory.
- Start the standby database from a patched home

```
emcli db_software_maintenance
-performOperation
-name=<User Specified Operation Name>
-purpose=<UPDATE_DB or UPDATE_RACDB>
-target_type=oracle_database
-target_list=<List of targets>
-normal_credential=credential name>
-privilege_credential=<credential name>
-dataguard_role=standby
-startupDatabase=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: The name and location of the temporary directory.
 - * enableDG: True/False, default value is false. If the standby database is managed by Data Guard this needs to be set to true. Setting to true enables the data guard configurations.
 - * primary_dbhost_creds: 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.

For reference information on the db_software_maintenance verb and updateoperations see: db_sm_performoperation_update in *Oracle Enterprise Manager Command Line Interface*.

Rollback a Standby Database

The rollback process is the same with only the purpose value will be changing. For rolling back rolling pactches like database Release Updates, follow these steps:

Rollback the standby database:

```
emcli db_software_maintenance
-performOperation
-name=<User Specified Operation Name>
-purpose=<ROLLBACK_DB or ROLLBACK_RACDB>
-target_type=oracle_database
-target_list=<List of targets>
-normal_credential=<credential name>
-privilege_credential=<credential name>
-database_credential=<SYSDBA credentials of the database>
-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: The location and name of the temporary directory.
 - * disableDG: True/False, default value is false. If the standby database is managed by Data Guard this needs to be set to true. Setting to true enables the data guard configurations.
- Rollback the primary database:

```
emcli db_software_maintenance
-performOperation
-name=<User Specified Operation Name>
-purpose=<ROLLBACK_DB or ROLLBACK_RACDB>
-target_type=oracle_database
-target list=<List of targets>
```



```
-normal_credential=<credential name>
-privilege_credential=<credential name>
-dataguard_role=primary
[-ignoreStandbyPrereq=true|false]
-input_file="data:<location of input file>"
[] Denotes optional parameter
```

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: The location and name of the temporary directory.
 - * enableDG: True/False, default value is false. If the standby database is managed by Data Guard this needs to be set to true. Setting to true enables the data guard configurations.

For non-rolling applicable patches such as OJVM (pre-Jan 2017 RUs), follow these steps:

Rollback the standby database

```
emcli db_software_maintenance
-performOperation
-name=<User Specified Operation Name>
-purpose=<ROLLBACK_DB or ROLLBACK_RACDB>
-target_type=oracle_database
-target_list=<List of targets>
-normal_credential=<credential name>
-privilege_credential=<credential name>
-dataguard_role=standby
-startupAfterSwitch=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: The location and name of the temporary directory.
 - * disableDG: True/False, default value is false. If the standby database is managed by Data Guard this needs to be set to true. Setting to true enables the data guard configurations.
- Rollback the primary database with dataguard role=Primary

```
emcli db_software_maintenance
-performOperation
-name=<User Specified Operation Name>
```



```
-purpose=<ROLLBACK_DB or ROLLBACK_RACDB>
-target_type=oracle_database
-target_list=oracle_database
```

-normal_credential = < credential name >

-privilege_credential = < credential name >

-dataguard_role=primary

-rolling=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 and location of the temporary directory.
- Start the standby database

```
emcli db_software_maintenance
-performOperation
-name=<User Specified Operation Name>
-purpose=<ROLLBACK_DB or ROLLBACK_RACDB>
-target_type=oracle_database
-target_list=<List of targets>
-normal_credential=credential name>
-privilege_credential=<credential name>
-dataguard_role=standby
-startupDatabase=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: The name and location of the temporary directory.
 - * enableDG: True/False, default value is false. If the standby database is managed by Data Guard this needs to be set to true. Setting to true enables the data guard configurations.
 - * primary_dbhost_creds: 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.

For reference information on the db_software_maintenance verb and rollbackoperations see: performoperation_rollback in *Oracle Enterprise Manager Command Line Interface*.

Rolling Back Listeners

This command is used to switch the listeners back to the old Oracle Home after the Update operation has been completed.

Verb

emcli db_software_maintenance -performOperation -name="Update Listener" purpose=rollback_listener -target_type=oracle_database -target_list="DB1" normal_credential="NC_HOST_CREDS:SYSMAN" -privilege_credential="HOST_PRIV:SYSMAN"
where:

- name: This is the unique name of the operation.
- purpose: ROLLBACK 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 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.

 procedure_name_prefix: OPTIONAL use this if the listener name exceeds the character limit, this allows for you to enter a deployment procedure name to ease tracking of the listener rollback procedure.

For complete rollback definition and parameters see: db sm performoperation rollback.

Cleanup Operations

Oracle homes and Container Databases consume valuable storage space, they can be cleaned up to release space when they are no longer in use. This operation removes unused, empty no longer needed Oracle Homes or Container Databases (CDB) to help you better administrate space and resources.





If any target is associated with the given home, Cleanup will be not done for the given home. You have to patch all the targets within an Oracle Home before Cleanup can be done. For RAC Database Oracle Homes, you will need to pass all nodes of Oracle Home name by comma separated value and Cleanup cannot be done from only one of the nodes.

Cleaning Up Oracle Homes

Oracle Homes can have several older versions of the same home, and these can accumulate space. You can modify how many of these previous Oracle Home versions are retained at any given point using the following EMCTL property: 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 will be retained. The default value for how many versions are stored is three.

For example, there are four versions of an Oracle Home with the target running from the fourth version; then only the earliest of them can be deleted if the value is set to three. Similarly if the value is set to one, only the current Oracle home from which the target is running will remain and the other three homes will be deleted.

There are two ways to clean up Oracle Homes:

- You can cleanup old Oracle Homes of recently Fleet Maintenance patched databases by specifying a database or cluster target list. The Cleanup operation will find older eligible Oracle Homers for the patched database or cluster targets and perform a cleanup.
- If there are any orphan homes within the lineage path, these can be cleaned up by passing a list of Oracle Homes directly to the db software maintenance verb.

To perform an Oracle Home cleanup, the following command string will need to be entered in EM CLI:

```
emcli db_software_maintenance -performOperation
-name="cleanup"
-purpose="CLEANUP_SOFTWARE"
-target_list=<List of home targets>
-target_type=oracle_home
-normal_credential=NORMALCRED:SYSMAN
-privilege_credential=ROOTCRED:SYSMAN
-workDir=<working directory>
```

Description of Input Variables

- name: This is the unique name of the operation.
- purpose: CLEANUP SOFTWARE, use this purpose only for Oracle Homes.
- target_list: This is a comma separated list of databases, clusters or Oracle Home targets which need to be cleaned up.
- target type: The type of target on which this operation is being performed.
- normal_credential: This must be entered in the format <Named Credential: Credential
 Owner> where:



- <Named Credential>: Named credential for the host where Oracle home was deployed.
- <Credential Owner>: The Enterprise Manager user who owns this Named Credential.
- privilege_credential: This credential must have root privilege and must be entered in the format <Named Credential: Credential Owner> where:
 - <Named Credential>: Named credential for the host where Oracle home was deployed.
 - <Credential Owner>: The Enterprise Manager user who owns this Named Credential.

These credentials are used to run scripts as root.

• force: The default value for this optional flag is *false*. When set to *true* this flag forces an Oracle Home cleanup to occur even if there are potential conflicts, like a break within the lineage, with the databases contained within.

In its default value of *false*, should conflicts arise when performing an Oracle Home cleanup the process stops and you see a warning message in the EM CLI console.

Cleaning Oracle Homes on Windows Single Instance Database Example:

```
emcli db_software_maintenance
-performOperation -name="Clean"
-purpose=CLEANUP_SOFTWARE -target_list=ORCLTDE.abc.example.com
-target_type=oracle_database
-normal_credential="LOCALADMIN:SYSMAN"
-workDir="C:\Users\LocalAdmin\work"
```

Obtain a Report of Oracle Homes Available for Cleanup

Additionally, you can perform a reportantly operation where the eligible homes for cleanup will be displayed but no further actions will be taken. Credentials are not needed for reportantly operations.

There are two ways to get reports for Oracle Home cleanup:

List the Oracle Homes for given target list of given target type:

```
emcli db_software_maintenance
-performOperation
-name="cleanup"
-purpose="CLEANUP_SOFTWARE"
-target_list=<List of home targets>
-target_type=oracle_home
-reportonly=true
```

List the Oracle Homes Using a Filter:

```
emcli db_software_maintenance
-performOperation
-name="cleanup"
-purpose="CLEANUP_SOFTWARE"
-filter="target_type:host#target_list:<List of host targets>"
-reportonly=true
```

Description of Input Variables

filter: target_type (mandatory) and target_list (optional) can be specified using the following format: - target type:cluster#target list: <List of cluster targets>

 reportOnly: The default value is false. If set to true, it prints the eligible Oracle Homes or CDBs for cleanup but these are not deleted.



This flag will show No CDB or Oracle Home is eligible for cleanup if you have Pluggable Databases (PDB) associated with an old CDB. Make sure the PDB is pointing to the new Oracle Home or CDB.

Cleaning Up Container Databases

When Fleet Maintenance patching operations occur, new Oracle Homes, and new Container Databases are deployed. As the individual PDB instances are patched, they are relocated to the newly deployed CDBs. Once all the PDBs in a CDB are patched, the old CDBs need to be deleted. The CLEANUP operation, part of the db_software_maintenance verb, automates the removal of old CDBs upon completion of a patching cycle.



Cleaning up container databases will clean all eligible homes corresponding to a CDB.

To perform an Oracle Container Database cleanup, the following command string will need to be entered in EM CLI:

```
emcli db_software_maintenance
-performOperation
-purpose="CLEANUP"
-target_list =<target CDB names>
-target_type=oracle_pdb
-name="Cleanup CDBs"
-normal_credential=NORMALCRED:SYSMAN
-privilege_credential=ROOTCRED:SYSMAN
-workDir=<working directory>
```

You can cleanup lost lineage Oracle Container Databases with the following command string:

```
emcli db_software_maintenance
-performOperation
-purpose="CLEANUP"
-target_list =<target CDB names>
-target_type=<oracle_database/rac_database>
-name="Cleanup orphan CDBs"
-normal_credential=NORMALCRED:SYSMAN
-privilege_credential=ROOTCRED:SYSMAN
-workDir=<working directory>
```

Description of Input Variables for Cleaning Container Databases

- name: This is the unique name of the operation.
- purpose: CLEANUP, use this purpose only for CDB cleanup.

- target_type: The type of target on which this operation is being performed.
- target_list: This is a comma separated list of targets which need to be cleaned up.
- normal_credential: This must be entered in the format <Named Credential: Credential
 Owner> where:
 - <Named Credential>: Named credential for the Oracle Home where the CDB is 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 the CDB was deployed.
 - < Credential Owner>: The Enterprise Manager user who owns this Named Credential.

These credentials are used to run scripts as root.

force: Optional flag, its default value is false. Cleanup of Oracle home affects the
databases running from that home. By default the cleanup is skipped if processes
(databases, listeners, cluster) are found running from this home. You can provide the
force=true flag to skip the verification and remove all processes running from this home.
This removes all processes running from the home which were not discovered in
Enterprise Manager as well.



This is a severely destructive operation, the procedure will start after a 5 minute delay, providing you with the opportunity to abort the procedure.

 reportOnly: The default value is false. When set to true, it prints the eligible CDBs for deletion but the Datbasese are not deleted.

The flag will show *No CDB or Oracle Home is eligible for cleanup* if you have PDBs associated with the old CDB. Make sure the PDBs are pointing to the new Oracle Home or CDB.

Note:

When using reportOnly credentials are not required.

Note:

For further information on CLEANUP_SOFTWARE and database software maintenance tasks and using db_software_maintenance see the following:

- CLEANUP_SOFTWARE see db_sm_performoperation_cleanup in *Oracle Enterprise Manager Command Line Interface*.
- db_software_maintenance see db_software_maintenance in *Oracle Enterprise*Manager Command Line Interface.



Emergency Patching

You can now perform emergency patching with Fleet Maintenance. Emergency patching does not replace regular maintenance cycle patching, non-critical patches should be flagged for inclusion into the next version of the image.

Emergency patched databases will be labeled as *Drifted* databases and will need to be reconciled in the next maintenance cycle either by rolling the changes to a new version of the image or override it with the latest version.

Emergency patching with Fleet Maintenance will not require the creation of a new version of a Gold Image. Databases that require an emergency patch will be switched out to new homes similar to the process used for scheduled maintenance. By switching the problem databases to new homes they can be isolated and patched without impacting functioning databases in the original home.

Overview of Steps For Emergency Patching

- 1. Prepare a complete list of emergency patches to be applied.
 - a. Identify required emergency patches.
 - Identify and merge conflicts between patches applied and the list of patches to be applied.
 - c. Request merge patches to Oracle Support (if required).
 - d. Once all desired patches are available, download all patches into the software library.
- Deploy software along with emergency patches into a new home.
- 3. Switch the emergency patched database to the new home.

Deploy Software Along with Emergency Patches

Single Instance Database

```
emcli db_software_maintenance
-performOperation
-name="Deploy - Critical one-off 12345"
-purpose=EMERGENCY_DEPLOY_DB_SOFTWARE
-target_type='oracle_database' -target_list="DB1"
-normal_credential="<normal credential>"
-privilege_credential="<host privilege credential>"
-inputfile="data:/<user>/oracle/deploy.txt"
-patch list=<List of conflict-free patches in the software library>
```

Where:

- PURPOSE must be set with emergency deploy db software.
- PATCH_LIST must follow the patchID:ReleaseID:PlatformID:LanguageID format. ReleaseID, PlatformID, LanguageID are optional arguments.
- INPUTFILE must contain:
 - NEW ORACLE HOME LIST=<Emergency home location>
 - DISPATCHER LOC=<Dispatcher location>
 - PATCH LIST=<List of conflict-free patches in the software library>

RAC Database

```
emcli db_software_maintenance -performOperation -name="Deploy - Critical one-
off 12345"
-purpose=EMERGENCY_DEPLOY_RAC_SOFTWARE
-target_type='rac_database' -target_list="DB1"
-normal_credential="<normal credential>"
-privilege_credential="<host privilege credential>"
-inputfile="data:/<user>/oracle/deploy.txt"
-patch list=<List of conflict-free patches in the software library>
```

Where:

- PURPOSE must be set with EMERGENCY DEPLOY RAC SOFTWARE.
- patch_list Can be either passed as command line argument or as the input file parameter. In case of input file parameter the parameter should be defined in upper case.

Switch Database to Emergency Home

To switch database you will need to run an <code>UPDATE_DB</code> command, this will move the database to be patched to the new empty home.

```
emcli db_software_maintenance -performOperation -name="Update DB"
-purpose=UPDATE_DB
-target_type=oracle_database
-target_list="DB1"
-normal_credential="<normal credential>"
-privilege_credential="<host privilege credential>"
-destinationOracleHomeLocation=<path of the new Oracle Home Location>
-force=<true/false>
```

The force flag is by default set to false. The flag is mandatory if the target is subscribed to an image and you still want to switch the database to the emergency home.

Publish The New Version of the Image

The uploaded image is in DRAFT by default. As the administrator you have to explicitly mark it CURRENT in order to publish it. This is achieved by the following command:

```
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.



Rollback Emergency Patches

Should a rollback be required, the database moves back into its original untouched home. This avoids any confusion whether the rollback correctly reverted the software and its file permissions or not.

```
emcli db_software_maintenance
-performOperation -purpose="ROLLBACK_DB"
-target_list="DB1"
-target_type=oracle_database
-name="RollbackDB"
-description="Relocates the DB from inactive CDB to newly active CDB"
-normal_credential="<normal credential>"
-privilege_credential="<host privilege credential>"
-database credential="<Database credential>"
```

Consolidate Drifted Databases Into A Oracle Home

Over time, several databases will find a need to apply emergency patches. This results in the databases that were previously sharing the same Oracle home to get moved into their own Oracle homes with these additional patches. These databases have now drifted from the published gold image and are called *Drifters*. For more information see: Emergency Patching.

When you need to patch these databases, new Oracle homes will be deployed for each of these databases resulting in an Oracle home sprawl. The recommended practice would be to get these databases back into the shared Oracle home as before.

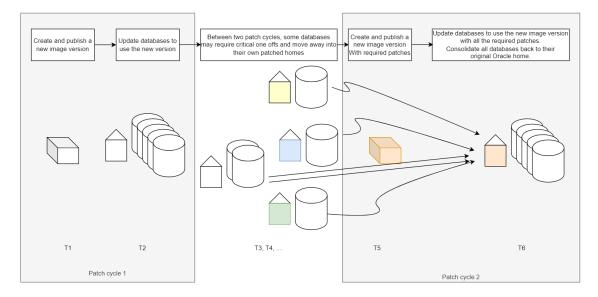


Figure 18-5 Database Drift Use Case

To now update and bring back into drifters into their correct Oracle homes during your next update follow these steps:

 Create and publish a new image version with all patches required to consolidate the homes. For more information see: Create a Gold Image Using EMCLI.

- 2. Update the gold images version status to CURRENT. For more information see: Changing Version Status to Current.
- 3. Update the database to use the new version. For more information see: Update Databases Using Fleet Maintenance and Update Operations.
- 4. Consolidate all databases back to the same new home as before. For more information see: Consolidate Oracle Homes.

Update Operations

Update operations are the most common operations performed within Fleet Maintenance. There are two types of supported **UPDATE** scenarios: updates and upgrades. An update is a minor improvement operation, moving an Oracle Home from one patch set to another (Ex: 19.17.0 to 19.19.0) whereas an upgrade, is a major improvement, moving an Oracle Database from one major release to another (ex: 12.2.0.1 to 19.19 or 18.x to 19.x). All the steps such as Gold Image creation, image subscription and deploy image remain the same.

Before you begin an update operation using EM CLI ensure the current version of the subscribed image for a target is deployed.



Note:

- Currently we do not support node-wise update for Upgrades. The update command for "Database Upgrade" requires database credentials to be provided as part of the command.
- When performing RAC or GI updates and upgrades it is a best practice to specify a working directory that is not a shared location between two nodes. This helps avoid overwriting checkpoint files and potential update failures.
- Before performing an UPDATE operation, if your listener is in a custom location, be sure to first set the TNS_ADMIN environment variable to this location. Next, restart the agent and then proceed with these operations.
- If during a database or RAC update you need to update any **static entry** and your listener.ora is located in the Grid Oracle Home default location (\$GI_HOME/network/admin) normal credentials will not work (for example DB Home Owner). This is because the owner of the file is the Grid Home owner; make sure the file has permissions to at least group level to be modified by both Database and Grid Oracle Home Users.
- Refer to the Oracle Database Upgrade Guide for the supported upgrade paths.
- Starting with Enterprise Manager 13c Release 5 Update 5 (13.5.0.5) you can now update a non Container Database (CDB) into a new Pluggable Database using Fleet Maintenance, select Upgrade Non-CDB as Pluggable Database under Operation to filter databases that are available for this type of upgrade. This type of upgrade, takes a non CDB database and performs an upgrade into a new PDB. For a step by step tutorial see: Patching an Oracle Pluggable Database Using Oracle Enterprise Manager's Fleet Maintenance Graphic User Interface.
- Starting with Enterprise Manager 13c Release 5 Update 17 (13.5.0.17) logical standby database upgrade must be performed using the Fleet Maintenance UI. Non-rolling updates are no longer be available through the main database menu.
- Starting with Enterprise Manager 13.5 Release update 19 (13.5.0.19), before and after a UPDATE/ROLLBACK operation on single instance or RAC databases, the deployment procedure activity logs include the number of invalid objects found.

UPDATE DB

This command moves an Oracle Database from its Oracle Home to the newly created Gold Image standard Oracle Home.

```
emcli db_software_maintenance
-performOperation -purpose="UPDATE_DB"
-target_list="db1221"
-target_type="oracle_db"
-name="Operation Name"
-description="Operation description"
-normal_credential="NORMAL:SYSMAN"
-privilege_credential="PRIV:SYSMAN"
-database credential="DB SYS CREDS:SYSMAN"
```



Note:

For Windows Single Instance database update operations use windows_sec_credential="Named Credential: Credential Owner" instead of privilege_credential. Starting with Enterprise Manager 13. Release Update 2 you can also use either -is_virtual_account="true" or -is_builtin_account="true" accounts.

Examples:

```
emcli db software maintenance -performOperation -name="Update SIDB" -
purpose=UPDATE DB
-target type=oracle database -target list="ORCL1910.abc.sample.com" -
rolling=false
-normal credential="LOCALADMIN:SYSMAN" -
windows sec credential="LOCALUSERCRED:SYSMAN"
-database credential="SYSDBA:SYSMAN"
emcli db software maintenance -performOperation -name="Update SIDB" -
purpose=UPDATE DB
-target_type=oracle_database -target_list="ORCL1910.abc.sample.com" -
rolling=false
-normal credential="LOCALADMIN:SYSMAN"
-is virtual account="true" -database credential="SYSDBA:SYSMAN"
emcli db software maintenance -performOperation -name="Update SIDB" -
purpose=UPDATE DB
-target_type=oracle_database -target_list="ORCL1910.abc.sample.com" -
rolling=false
-normal credential="LOCALADMIN:SYSMAN" -is builtin account="true"
-database credential="SYSDBA:SYSMAN"
```

UPDATE PDB

Moves a Pluggable Database 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="pdb1331"
-target_type="oracle_pdb"
-name="Operation Name"
-description="Operation description"
-normal_credential="NORMAL:SYSMAN"
-privilege_credential="PRIV:SYSMAN"
-database_credential="DB_SYS_CREDS:SYSMAN"
```



UPDATE_GI

```
emcli db_software_maintenance
-performOperation
-name="Update Cluster"
-purpose=UPDATE_GI
-target_type=cluster
-target_list= CLUSTER1
-normal_credential="NC_HOST_CREDS:SYSMAN"
-privilege_credential="HOST_PRIV:SYSMAN"
-node_list="host1.example.com"
-startupDatabase=<true/false>
-SKIP_CVU_CHECK=<true/false>
-drain timeout=<seconds>
```

For a use case on <code>UPDATE_GI</code> and a detailed explanation see: Node Wise RAC Database / Cluster Update

SKIP_CVU_CHECK: This flag allows skipping the CVU check if set to true. This is an optional flag only used for update gi and update siha operations.

When performing <code>UPDATE_GI</code> operations and you need to pass multiple clusters in the -target list, -node list is not supported.

To avoid service interruptions, patching of Oracle Grid Infrastructure is done most of the time in a rolling manner. While the database instance is brought down on the node that is getting patched. However, RAC databases remain up on the remaining nodes.

Oracle Grid Infrastructure patches in a rolling manner to avoid service interruption. The database instance is brought down only on the patched node, while RAC databases continue running on the remaining nodes. Starting with Oracle Enterprise Manager 13.5 Release Update 17 (13.5.0.17) Fleet Maintenance you can use the DRAIN_TIMEOUT parameter to reduce the overall downtime for the node being patched. For the UPDATE_GI and UPDATE_RACDB cases, you can leverage the Grid Infrastructure capabilities to shutdown instances on one node in parallel. Define the DRAIN_TIMEOUT value (in seconds) when using the db_software_maintenance EM CLI verb. If the DRAIN_TIMEOUT parameter is not specified, the default drain time specified at service creation will be used.

When specifying the drain time in seconds, to allow for resource draining to complete and services to stop, the following are accepted values:

- DRAIN_TIMEOUT=0: Drain occurs immediately, database instances are shutdown immediately.
- DRAIN_TIMEOUT=<time>: Time in seconds allowing for the drain to occur before database instances are shutdown.
- DRAIN_TIMEOUT is not specified in the input file, the default drain timeout specified at database service creation will be used. This is the default option.



UPDATE_GIMR



From Oracle Database 21c onward, the GIMR update is now separated from GI update or patching. Performing an update of both GI and GIMR will happen in 2 separate operations.

Before starting a GIMR update, make sure that you have completed the

DEPLOY GIMR SOFTWARE instructions outlined in: Deploy an Oracle Grid Infrastructure Database.

```
emcli db_software_maintenance
-performOperation
-name="Update GIMR 21c"
-purpose=UPDATE_GIMR
-target_type=cluster
-target_list= <cluster target name>
-normal_credential="NC_HOST_CREDS:SYSMAN"
-privilege_credential="HOST_PRIV:SYSMAN"
```

UPDATE RACDB

The EM CLI command format for UPDATE RACDB:

```
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_credential="HOST_PRIV:SYSMAN"
-database_credential=<SYSDBA credentials of the database>
-rolling=<true/false>
-node_list="host1.example.com"
-drain timeout = <time in seconds>
```

- 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 only 2 hosts at a time; the value of this parameter needs to be specified as node_list="host1, host2"
- drain_timeout: This optional parameter can be used for both rolling and non rolling use
 cases and can be passed as an EM CLI argument or as an input_file parameter. This drain
 parameter is a time measure of seconds allowed in order to complete the resource
 draining action. By default, this parameter is not set.



This is an optional parameter. If no date is provided, the fleet operation will start immediately.

For primary and standby databases, upgrades need to be done through the Rolling Upgrade user interface in Enterprise Manager. The wizard can be launched from the target database home page (either primary or standby database). On the Database home page, select **Oracle Database**, **Provisioning**, then select **Upgrade Database**.

UPGRADE DB

This purpose is only to be used for non-CDB to PDB migration:

```
emcli db_software_maintenance
-performOperation -purpose="UPGRADE_DB"
-target_list="db1916"
-target_type="oracle_db"
-name="Operation Name"
-description="Operation description"
-normal_credential="NORMAL:SYSMAN"
-privilege_credential="PRIV:SYSMAN"
-database_credential="DB_SYS_CREDS:SYSMAN"
-migrate_non_cdb pdb="true"
```

Note:

migrate_non_cdb_pdb is only used when you want to migrate a database to a pluggable database where the source version is lower than 20c.

UPGRADE RACDB

This purpose is only to be used for non-CDB to RAC migration:

```
emcli db_software_maintenance
-performOperation -purpose="UPGRADE_RACDB"
-target_list="db1221"
-target_type="oracle_db"
-name="Operation Name"
-description="Operation description"
-normal_credential="NORMAL:SYSMAN"
-privilege_credential="PRIV:SYSMAN"
-database_credential="DB_SYS_CREDS:SYSMAN"
-migrate_non_cdb_pdb="true"
```

Note:

migrate_non_cdb_pdb is only used when you want to migrate a database to a pluggable database where the source version is lower than 20c.

Transparent Data Encryption (TDE) Support in Fleet Maintenance

Fleet Maintenance now supports updating Databases, Container Databases and Pluggable Databases that have tables encripted using TDE.



For Oracle Databases and Container Databases, using oracle_database Target_type, the TDE wallet password must be provided in the input file -input_file="data:/<file pathc>/input.properties" during UPDATE_DB

Example: srcWalletPassword=welcome123.

Note:

- If the -database_credential is not provided, TDE configuration will be verified by connecting to the database using SQL*Plus.
- If the wallet status is CLOSED, TDE configuration will be skipped.
- After the migration completes, review the wallet status. If the status is Closed manually open the Wallet File to access the encrypted data in the database.

For Oracle Pluggable Databases, using oracle pdb Target_type:

- When executing UPDATE_PDB, the Database credentials should be provided using the
 existing flag -database_credential="DB_SYS_CREDS:SYSMAN". If the Database credentials
 are not provided, the preferred credentials, if set, will be used.
- TDE must be already configured in the destination CDB
- The TDE wallet source and destination passwords must be provided in the input file input_file="data:/scratch/input.properties" during UPDATE_PDB

 Example: srcWalletPassword=welcome123, destWalletPassword=welcome123

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. A cluster update, followed by a RAC database update results in a database instance getting restarted twice.

There may be several situations where you as the administrator may require more control over the switch process. For example, you may choose to perform a node-wise update of the cluster and RAC databases in order to avoid multiple restarts of the database instances. You may also need to perform node specific pre and post steps.

This option provides this control by enabling you to perform the following tasks for each node:

- Switch the cluster instance
- Leave the RAC database instances shutdown
 For example, 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
-normal_credential="NC_HOST_CREDS:SYSMAN"
-privilege_credential="HOST_PRIV:SYSMAN"
-rolling=true -node list="host1.example.com"
```



- -startupDatabase=false
 -skip_cvu_check=true
- 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.
- **SKIP_CVU_CHECK**: This flag allows skipping the CVU check if set to *true*. This is an optional flag in UPDATE GI operations.
- When performing UPDATE_GI operations and you need to pass multiple clusters in the target list, -node list is not supported.

Update Oracle Homes Using Patched Oracle Homes

The following section highlights how Oracle Homes can be switched to other patched 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 patched Home.
	 oratab file entries are updated
	 Instance is brought down and brought up from the patched home.
	 Target is brought of Blackout mode.
	 For pre-12c databases SQLs are run using SQL*Plus. For 12c and higher databases, the Datapatch utility is called to make data changes.
Grid Infrastructure	 All relevant targets are put in Blackout mode.
	 Backup CRS configuration files from the source home directory: \$SRC_HOME/crs/utl/crsconfig_dirs.
	• Configuration files like: spfile, password, listener.ora, sqlnet.ora tnsnames.ora are copied over to the patched 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 Patched Home as new Oracle Homes.
	 Configuration scripts are run.
	 All targets are brought out of Blackout mode.
	 If the management database is used, pre-12c databases SQLs are run using SQL*Plus. For 12c and higher databases, the 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 patched Home.
	• oratab file entries are updated
	 Instance is brought down and brought up with Patched Home as new Orac Home.
	Target is brought out of Blackout mode.
	 For pre-12c databases SQLs are run using SQL*Plus. For 12c and higher databases, the Datapatch utility is called to make data changes.



Database Type	Steps Performed
RAC Non-Rolling Node	All RAC nodes are put in Blackout mode.
Node	 Configuration files like spfile, password, listener.ora, sqlnet.ora, tnsnames.ora are copied over to the patched Home.
	oratab file entries are updated
	 Instance is brought down and brought up with Patched Home as new Oracle Home concurrently.
	 For pre-12c databases SQLs are run using SQL*Plus. For 12c and higher databases, the Datapatch utility is called to make data changes.

Patch History

Starting with Oracle Enterprise Manager 13.5 Release Update 21 (13.5.0.21) you can now access database patching history for update and upgrade operations in Fleet Maintenance.



Your environment's data retention policy dictates how far back Fleet Maintenance can show update and upgrade history.

Currently the following target types are supported for Patch History:

- Single Instance databases
- RAC databases
- Grid Infrastructure
- Single Instance High Availability (SIHA)

Obtain Patch History Via EM CLI

Read Patch History:

```
emcli lcm_operations
-getLcmHistory
-input_file="<path of the JSON file>"
```

Where the input file contains the following:

```
"targetNames" : [<List of targets>],
   "targetType" : "oracle_database",
   "targetIds" : [<List of ID's>],
   "tasks" : ["update", "upgrade"],
   "timeStartedGreaterThanOrEqualTo" : "<Start date format YYYY-MM-DD
HH:MM:SS>",
   "timeStartedLessThanOrEqualTo" : "<End date format YYYY-MM-DD
HH:MM:SS>",
   "status" : "completed",
   "sort" : "targetType",
   "limit" : 100,
```



```
"page" : "<Page ID>" }
```

· Export Patch History:

```
emcli lcm_operations
-exportLcmHistory
-input file="<path of the JSON file>"
```

Where the input file contains:

```
"destinationHostName" : "<host name>",
  "destinationHostCredential" : "<host_cred>",
  "csvFileLocation" : "<CSV file path>",
  "csvFileName" : "<file_name.csv>"
  "targetNames" : [<list of targets>],
  "targetType" : "oracle_database",
  "targetIds" : [<list of ids>],
  "tasks" : ["update", "upgrade"],
  "timeStartedGreaterThanOrEqualTo" : "<Start date format YYYY-MM-DD
HH:MM:SS",
  "timeStartedLessThanOrEqualTo" : "<End date format YYYY-MM-DD HH:MM:SS",
  "status" : "completed",
  "sort" : "targetType"
}</pre>
```

Use the following EM CLI commands to read and export your Fleet Maintenance update and upgrade history: For a complete list of definitions and examples of the $lcm_operations$ verb see: $lcm_operations$

Obtain Patch History Via REST API

For REST API see: Database Maintenance- Updates (Patching) And Upgrades REST Endpoints.



Upgrade databases and Grid Infrastructure Using Fleet Maintenance

Upgrading transforms an existing Oracle Database environment (including installed components and associated applications) into a new release Oracle Database environment. The data dictionary for the database is upgraded to the new release.

Upgrading does not directly affect user data; no data is touched, changed, or moved during an upgrade. Review the supported options for direct upgrades to the latest Oracle Database release. The path that you must take to upgrade to the latest Oracle Database release depends on the release number of your current database.

Note:

The process of upgrading databases and Grid Infrastructure follows the same model as discussed for patching in the previous section. Unless otherwise explicitly specified, the process documented below for databases applies for Grid Infrastructure as well.

Upgrade is broken down into the following steps:

- 1. Plan Your Upgrades
- 2. Prepare Your Upgrade
- 3. Upgrading Databases using AutoUpgrade
- 4. Data Guard Rolling Upgrade
- 5. Downgrade Databases

Plan Your Upgrades

The upgrade process entails additional planning steps compared to updates, to properly plan an upgrade ensure the following points are considered before starting an upgrade.

Before starting to plan for an upgrade the following things must be considered:

- Become familiar with the features of the new release of Oracle Database
- Determine the upgrade path to the new release
- Select an upgrade method
- Select an Oracle home directory for the new release
- Develop a testing plan
- Prepare a backup strategy
- Follow pre-upgrade recommendations
- Run a pre-upgrade fixup scripts, or perform manual pre-upgrade system updates

Planning your upgrade is broken down into four steps:

- 1. Define your upgrade goals: There are many reasons for you to embark on an upgrade journey, for example: you might be running on an older database version and have no support from Oracle to fix the issues, or your compliance policy suggests that you should be on newer version. Whichever be the reason, you must ensure that to call your upgrade successful, in terms of upgrade process, performance acceptance, outage duration, business impact due to non-availability and schedule of upgrade are clearly stated and called out.
- Define your upgrade policies: Your upgrade policy must take care of your rollback scenarios. These could be related to failed upgrades, poor database performance after upgrade. You must also take into consideration that rollback will add to the non-availability of database.
 - Similar to rollback, you should also consider cleanup of unused Oracle Homes. Once upgrade is successfully complete and all participating teams have signed off the results post upgrade, based on your cleanup policy, you should remove the unused homes.

3. Organize:

- Amount of images: It is recommended that you have the fewest number of images as
 possible, maintaining large numbers of images can add greatly to the maintenance
 overhead. Images can be based on you different application types, data center name
 or location, application usage (test, development, production, or master).
- Targets to subscribe per image: Based on the criteria for image creation, a similar criteria can be followed for image subscription for targets. You can have a gold image for a particular application type such as ERP and subscribe all the ERP supporting databases to this image.
- Consolidation plans: An upgrade is an opportunity to re-organize your database estate.
 Aim to consolidate wherever possible and reduce the number of databases, Multi-tenant is the way forward. Try to consolidate multiple unique databases into one or more Container databases (CDB).
- Prepare for SQL Performance Analytics: Its important to have a benchmark for SQL performance so that you can review, analyze and fix any SQL queries that are performing poorly after database upgrade.

Prepare Your Upgrade

To begin preparing for your upgrade there is a list of actions to take place. You will need to create the image, subscribe targets to the image, deploy the Oracle Homes and check upgrade readiness.

The following is a set of steps that need to be done to properly prepare for your upgrade:

- Create image for new release: Based upon the upgrade route (direct or indirect), create gold image of that particular release. In case of direct upgrade route, simple create the gold image of the upgrade release. For more information see: Image Maintenance.
- 2. Subscribe Targets: Use Fleet Maintenance Hub, EM CLI or REST APIs to subscribe the targets to the new image. For more information see: Subscription Tasks
- 3. Deploy new Oracle Home: Deploy the higher version binaries of Oracle Home before your upgrade outage. Deploying a new Oracle Home will not interfere with your existing, in-use Oracle Database. Deploying the Oracle Home in advance will allow you to utilize the maximum outage window for the upgrade process. For more information see: Update Databases Using Fleet Maintenance



4. Check for upgrade readiness: Ensure you have all the approvals to begin the upgrade,. Its always advisable to perform few iterations on non-production setup first and resolve any issues. During the production cut over, ensure that you document and apply all the fixes from non-production setup to production.

Upgrading Databases using AutoUpgrade

Starting with Enterprise Manger 13.5 Release Update 18 (13.5.0.18) AutoUpgrade will be the default tool used for upgrading databases using Fleet Maintenance. With AutoUpgrade, Enterprise Manager automatically upgrades Software Library subscribed single instance or RAC databases version 12.2 and above images.

Should you prefer to use the Database Upgrade Assistant (DBUA) for upgrades, you will need to set the oms property oracle.sysman.db.upgrade_using_autoupgrade_tool to false. Use the following sample code:

emctl set property -name oracle.sysman.db.upgrade_using_autoupgrade_tool sysman pwd <EM SYSMAN PASSWORD> -value false

Prerequisites for using AutoUpgrade

To use AutoUpgrade with Enterprise Manger Fleet Maintenance managed databases the following prerequisites must be met:

- Data Collections must be added and running. For more information see: Extending Configuration Data Collections.
- Upload the latest Autoupgrade.jar to the software library (either in the default location, or a custom location: to use during upgrades. To download the most recent version of AutoUpgrade see: My Oracle Support Document 2485457.1.
- The Fast Recovery Area (FRA) should be set for all databases that will update automatically, AutoUpgrade will create the Guaranteed Restore Point (GRP) by default. If FRA is not setup AutoUpgrade will not proceed.
 For more information on FRA see: About the Fast Recovery Area and the Fast Recovery Area Disk Group.
- The database must be in ARCHIVELOG mode for AutoUpgrade to be able to create the GRP by default. If the database is in NOARCHIVELOG mode a warning will be shown during the prerequisite checks. You can pass the flag isEnableArchiveLog as true to enable it as part of an upgrade deployment procedure.
 When using the Fleet Maintenance UI this flag will be always passed as true.
- For TDE (Transparent Data Encryption) configured database, the wallet should be in AUTOLOGIN mode. If the wallet is a PASSWORD type you will see a warning during the prerequisite checks. Pass the flag isConvertToAutoLoginWallet as true to convert it to AUTOLOGIN as part of an upgrade deployment procedure.

 When using the Fleet Maintenance UI this flag will be always passed as true.

For more information on Oracle AutoUpgrade see: Using AutoUpgrade for Oracle Database Upgrades .

AutoUpgrade JAR Location

The AutoUpgrade jar will be used from the default component location in the software library, located at Fleet Maintenance, then Tools, and click on **AutoUpgrade**. It is recommended to always use the latest version of the autoupgrade.jar file, the file must be maintained and managed by you.



In the Fleet Maintenance UI the source of AutoUpgrade utility will always be auto populated and modifiable. You can provide a custom component location from the software library for the autoupgrade.jar with the parameter autoUpgradeToolPath.

You can run the AutoUpgrade jar from a Destination Oracle home with the parameter isAutoUpgradeToolFromOracleHome as true.

If you have not passed the jar location, and the default component location for the AutoUpgrade jar doesn't exist in the software library, then AutoUpgrade jar will be used from the destination Oracle home.

Users can pass all options supported by AutoUpgrade using the parameter autoUpgradeAdditionalOptions, for example:

```
autoUpgradeAdditionalOptions= {"restoration":"no","catctl options":"-n 64"}
```

For a complete list of AutoUpgrade utility configuration files, parameters and options see: AutoUpgrade Utility Configuration Files Parameters and Options.

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 shutdown 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:

 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



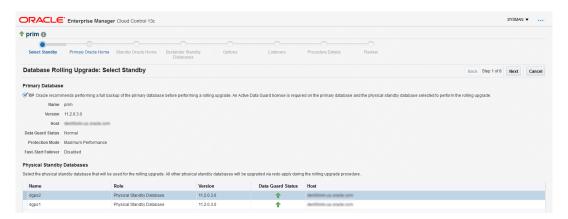
Upgrade wizard, the Select Standby page, which displays the list of all the standby databases associated with the primary database that can be selected for the rolling upgrade process. The Select Standby page also displays an overall summary of the Data Guard configuration such as the name of the primary 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.

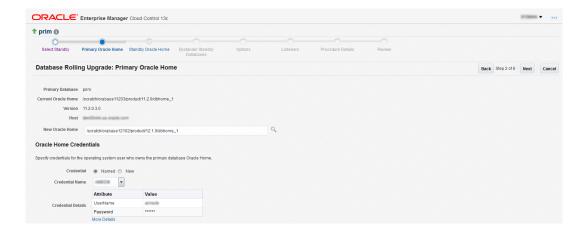
2. On the Select Standby page, select the physical standby database to use for the rolling upgrade from the physical standby database table, then click **Next**. All other physical standby databases will be upgraded individually after the rolling upgrade is complete.

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.



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.

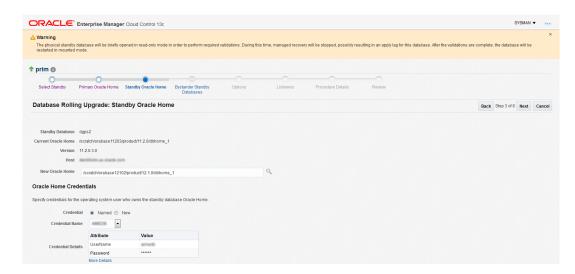




On the Standby Oracle Home page, enter the new Oracle Home location for the selected physical standby database. Click Next.

Enterprise Manager performs the pre-requisite checks and storage checks for the physical standby database and then opens the Options page.

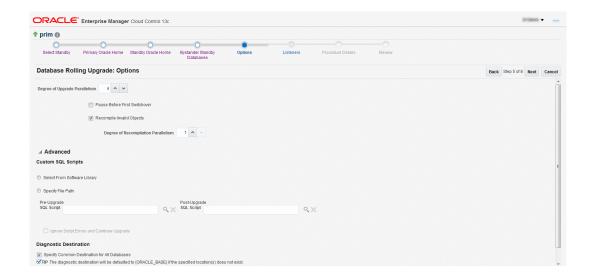
If you are submitting a Rolling Upgrade Procedure for a primary database with multiple physical standby databases, in this step Enterprise Manager performs the pre-requisite checks and storage checks for the physical standby database and then opens the Bystander Physical Standby Databases page. On the 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.



5. 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 bystander 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.





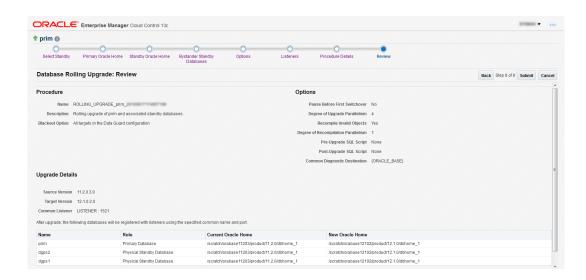
6. On the Listeners page, choose the default Common Listener listed or select from a list of common listeners on the primary and the standby database hosts. The upgraded databases will be registered with the selected listener. You can also specify a new listener name and port to create and register the 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.
- 8. 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.

Downgrade Databases

In the event of an error or misconfiguration Grid Infrastructure and Single Instance High Availability (SIHA) databases can be retroactively downgraded to the previous version.

Downgrade Using EM CLI

To downgrade a database using EM CLI, use the following commands:

Grid Infrastructure:

```
emcli db_software_maintenance
-performOperation -purpose=DOWNGRADE_GI
-name="Downgrade GI"
-description="Downgrade GI"
-target_list=<cluster target list>
-target_type=cluster
-normal_credential=NORMAL:SYSMAN
-privilege credential=ROOT:SYSMAN
```

SIHA:

```
emcli db_software_maintenance
-performOperation
-name="Downgrade SIHA"
-description="Downgrade SIHA"
-purpose=DOWNGRADE_SIHA
-target_list=<databse target list>
-target_type=has
-normal_credential=NORMAL:SYSMAN
-privilege_credential=ROOT:SYSMAN
```



Use -force="true" in case of a downgrade, due to upgrade failure. This will force the downgrade operation to execute.

Downgrade Using REST API

To downgrade using REST API use the following:

Table 19-1 Downgrade Grid Infrastructure

Feature	Description	
Request Method	POST	



Table 19-1 (Cont.) Downgrade Grid Infrastructure

Feature	Description	
URI	<oms_console_host>:OMS_CONSOLE_PO RT>/em/websvcs/restful/emws/db/ fleetmaintenance/performOperation/rollback</oms_console_host>	
Headers	Authorization: Basic Accept: application/json	
Payload		
	{	
	"name": "Downgrade GI",	
	"targetName" :	
	"example-cluster",	
	"targetType" :	
	"cluster",	
	"normalCredString" :	
	"USER:SYSMAN",	
	"privCredString" :	
	"ROOT:SYSMAN",	
	"downgradeOperation":"true"	
	}	

Table 19-2 Downgrade SIHA

- eature	Description	
Request Method	POST	
URI	<pre><oms_console_host>:OMS_CONSOLE_PO RT>/em/websvcs/restful/emws/db/ fleetmaintenance/performOperation/rollback</oms_console_host></pre>	
Headers	Authorization: Basic Accept: application/json	
Payload	<pre>"name": "Downgrade SIHA",</pre>	

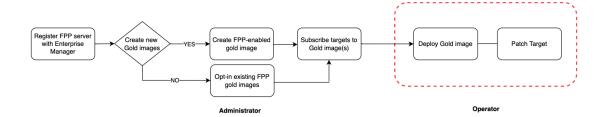
Oracle Fleet Patching and Provisioning and Enterprise Manager Integration

Starting with Oracle Enterprise Manager 13.5 Release Update 14 (13.5.0.14), Fleet Maintenance integrates Oracle Fleet Patching and Provisioning (FPP) solution. This enables FPP customers to leverage native capabilities of Enterprise Manager automated risk assessment, security patch recommendations, and proactively patch databases and Grid Infrastructure using Enterprise Manager Fleet Maintenance smart User Interface, and automation via REST APIs and EMCLI.

Oracle Fleet Patching and Provisioning (FPP) enables mass deployment and maintenance of standard operating environments for databases, clusters, and user-defined software types. With Oracle Fleet Patching and Provisioning, you can also install clusters and provision, patch, scale, and upgrade Oracle Grid Infrastructure and Oracle Database 11g release 2 (11.2), or later.

A Oracle Database Lifecycle Management Pack license is required in order to use these options.

Figure 20-1 FPP Workflow



Terminology

- **Gold Image**: Are Enterprise Manager provisioned base software version plus the additional patches. Gold is the qualifier for the ideal level for a software configuration.
- **Images**: Are Enterprise Manager provisioned software sets of specific version, for a particular database or Grid Infrastructure.
- **Image Series**: Are FPP provisioned sets base software version plus the additional patches. Gold is the qualifier for the ideal level for a software configuration.
- Versions: Are Enterprise Manager provisioned software sets of specific version within an image.

Enterprise Manager Fleet Maintenance manages Gold Images using **Images** and **Versions**, while FPP manages Gold Images using **Image Series** and **Images**. Within Enterprise Manager, FPP Image Series will appear as an Image, and a FPP Images will appear as a Version.

To use an existing image series for patching and upgrading your Grid Infrastructure or Oracle Databases, you will need to **opt-in** that image series into Enterprise Manager.

Exadata Patching Using Fleet Maintenance FPP

Starting with Oracle Enterprise Manager 13.5 Release Update 23 (13.5.0.23), Fleet Maintenance integrates Oracle Fleet Patching and Provisioning (FPP) solution for Exadata Patching. For a complete guide to patching your Exadata targets and the different supported Exadata types for patching see: Exadata Patching in *Advanced Management and Monitoring for Engineered Systems*.

FPP for Enterprise Manager Prerequisites

In order to integrate Fleet Patching & Provisioning (FPP) services into Enterprise Manager you must perform the following prerequisites tasks. An <code>EM_ALL_ADMINISTRATOR</code> account is required for these tasks.

Environment Prerequisites

For full FPP support with Enterprise Manager the following configurations are required:

- Oracle Enterprise Manager 13c Release 5 Update 14 (13.5.0.14) or above
- FPP Server version 19.18 or above with patch 35159977
- FPP Client version 19.18 or above is required before patching databases on clusters

For more details for setting up and upgrading your FPP server, see Oracle Fleet Patching and Provisioning Administrator's Guide.

Deployment Prerequisites

- Existing FPP Customers:
 - Discover the FPP server and client, running on Grid Infrastructure, as cluster targets in Enterprise Manager. If the Grid Infrastructure where the FPP Server or FPP Client are running are already monitored by Enterprise Manager, no additional discovery of targets is required. You can discover targets in EM by configuring autodiscovery, or by using either a guided or a declarative discovery process. For more details, see Discovering and Adding Host Targets.
 - 2. Register existing FPP server in Enterprise Manager using REST API or EM CLI:
 - REST API: Register FPP Server
 - EM CLI: Use command emcli db_software_maintenance -setup registerFPPServer -input_file="<fileLocation>" with the input file containing:

```
"name":"host1",
"restUrl":"https://host1.example.com:24500/rhp-restapi/rhp",
"agentCredential":"CUSER:SYSMAN"
"username":"restuser",
"password":"<password>"
"recurring":"frequency:weekly;days:1"
```

- New FPP Customers:
 - It is recommend that you setup and configure an FPP server outside of Enterprise Manager, for more details see: Configure an Oracle Fleet Patching and Provisioning server.

- Alternatively you can also set up the FPP server within Enterprise Manager via FPP Setup Automation, for more information see: FPP Setup Automation.
- 2. Discover the FPP server and client, running on Grid Infrastructure, as cluster targets in Enterprise Manager. If the Grid Infrastructure where the FPP Server or FPP Client are running are already monitored by Enterprise Manager, no additional discovery of targets is required. You can discover targets in EM by configuring autodiscovery, or by using either a guided or a declarative discovery process. For more details, see Discovering and Adding Host Targets.
- 3. Register the new FPP server in Enterprise Manager using REST API or EM CLI:
 - REST API: Register FPP Server
 - EM CLI: Use command emcli db_software_maintenance -setup registerFPPServer -input_file="<fileLocation>" with the input file containing:

```
"name": "string",
    "restUrl": "https://xyz.com:4473/em/api/fppServers",
    "agentCredential": "USER:SYSMAN"
    "timeToStart": "<TIME>",
    "timeToEnd": "<TIME>",
    "isTargetZone": true,
    "scheduleType": "string",
    "recurring": "Weekly",
    "startGracePeriodInMinutes": 60
```

Once the prerequisites have been met, proceed to FPP Setup Automation.

FPP Setup Automation

With FPP Setup Automation you can setup and provision an FPP server of client while provisioning new Grid Infrastructure or on existing Grid Infrastructure.

You will need to provision both an FPP server and client, in these examples you can see sample commands to provision both server and client in new Grid Infrastructure or into existing Grid Infrastructure through Enterprise Manager



This method is currently supported for 19c releases only.

List of Variables Used for Provisioning FPP Server and Client

List of supported variables for FPP server provisioning:

Variable	Required	Description
TARGET_NAME	Yes	Specify the cluster target name in EM to set up the FPP Client
NORMAL_USER	Yes	Specify the named credentials for the user who owns the Grid Infrastructure home (Credential Name:Owner).



Variable	Required	Description
ROOT_USER	Yes	Specify the named credentials for the privileged user (Credential Name:Owner). In the event that you do not have root access, the procedure will require manual intervention to run the root scripts.
FPP_DISK_GROUP	Yes	Specify the disk group that FPP Server can use to store the software images. Oracle recommends a minimum of 100 GB for this disk group.
FPP_IMAGES_MOUNTPOINT	Yes	Provide a mount path that exists on all nodes of the cluster. The FPP Server uses this path to mount gold images.
REST_USER_NAME	Yes	Specify the REST user name to be created on the FPP server. This REST user will be registered and used by Enterprise Manager to communicate with the FPP Server.
REST_USER_PASSWORD	Yes	Specify the password for the REST user to be created.
REST_USER_EMAIL	Yes	Specify the email for the REST user.
CONFIGURE_GIMR	Optional	For 19c database releases, Fleet Patching and Provisioning server setup requires Grid Infrastructure Management Repository (GIMR) as a prerequisite. This procedure will setup GIMR if it is not already set up. Options: true/false.
GIMR_DISK_GROUP	Optional	Specify the disk group for the Grid Infrastructure Management Repository. If it is not provided, it will be configured with the same disk group as the Oracle Clusterware files.
CONFIGURE_GNS	Optional	Grid Naming Service is required as a prerequisite to setting up the Fleet Patching and Provisioning server. Options: true/false.
GNS_VIP_ADDRESS	Optional	The Grid Naming Service (GNS) virtual IP address is required as a prerequisite to setting up the Fleet Patching and Provisioning Server. GNS will be created if it does not already exist.
REFERENCE_HOST	Optional	Specify the node on which to perform the FPP setup. If not specified, any one of the nodes in the cluster will be automatically chosen.



Variable	Required	Description
WORKING_DIRECTORY	Optional	Specify the working directory path to stage the required binaries for execution. By default, it will use the agent stage directory.
MIN_HEAP_SIZE	Optional	Specify the minimum heap size to be configured for the FPP server. Supported formats in gigabytes. (Ex: 3G)
MAX_HEAP_SIZE	Optional	Specify the maximum heap size to be configured for the FPP server. Supported formats in gigabytes. (Ex: 5G)
PREREQS_ONLY	Optional	Only run the prerequisite checks to verify readiness to setup the FPP server. Options: true/false (default: false).

List of supported variables for FPP client provisioning:

Variable	Required	Description
TARGET_NAME	Yes	Specify the cluster target name in EM to set up the FPP server.
NORMAL_USER	Yes	Specify the named credentials for the user who owns the Grid Infrastructure home (Credential Name:Owner).
ROOT_USER	Yes	Specify the named credentials for the privileged user (Credential Name:Owner). In the event that you do not have root access, the procedure will require manual intervention to run the root scripts.
FPP_SERVER_TARGET_NAME	Yes	Specify the cluster target name in EM which is already registered as an FPP Server.
FPP_SERVER_HOST_USER	Yes	Specify the named credentials for the user who owns the Grid Infrastructure home in FPP Server (Credential Name:Owner).
REFERENCE_HOST	Optional	Specify the node on which to setup the FPP Client. If not specified, any one of the nodes in the cluster will be automatically chosen.
WORKING_DIRECTORY	Optional	Specify the working directory path to stage the required binaries for execution. By default, it will use the agent stage directory.



Variable	Required	Description
FPP_SERVER_WORKING_DIRE CTORY	Optional	Specify the working directory path from FPP Server to stage required binaries for execution. By default, it will use the agent stage directory.
PREREQS_ONLY	Optional	Only run the prerequisite checks to verify readiness to setup the FPP client. Options: true/false (default value: false)

Provision an FPP Server on Existing Grid Infrastructure

```
emcli submit_procedure -input_file=data:"/home/user_name/FPP/
input.properties" -name="SETUP FPP SERVER"
```



The parameter -name="SETUP_FPP_SERVER" is required when provisioning an FPP server. SETUP_FPP_SERVER is the command action name and not a user defined name for the operation.

Sample input file:

TARGET_NAME=emracn0607

FPP_DISK_GROUP=MGMTDB

FPP_IMAGES_MOUNTPOINT=/scratch/cuser/fppImages

GNS_VIP_ADDRESS=emracn0607-gns.subnet1rg.emsubnet.companyvcn.com

NORMAL_USER=CUSER:SYSMAN

ROOT_USER=CUSER_ROOT:SYSMAN

MIN_HEAP_SIZE=3G

MAX_HEAP_SIZE=5G

REST_USER_NAME=restuser

REST_USER_NAME=restuser

REST_USER_PASSWORD=<password>

REST_USER_EMAIL=user@company.com

Provision an FPP Client on Existing Grid Infrastructure

```
emcli submit_procedure -input_file=data:"/home/user_name/FPP/
input.properties" -name="SETUP_FPP_CLIENT"
```

Note:

The parameter -name="SETUP_FPP_CLIENT" is required when provisioning an FPP server. SETUP_FPP_CLIENT is the command action name and not a user defined name for the operation.



Sample input file:

```
TARGET_NAME=emracp1011
NORMAL_USER=CUSER:SYSMAN
ROOT_USER=CUSER_ROOT:SYSMAN
FPP_SERVER_TARGET_NAME=emracn0607
FPP_SERVER_HOST_USER=CUSER:SYSMAN
WORKING DIRECTORY=/scratch/cuser/tmp
```

Provision an FPP Server and Grid Infrastructure

```
emcli submit_procedure -input_file=data:"/home/user_name/FPP/
gi input.properties" -name="SETUP FPP SERVER"
```



The parameter -name="SETUP_FPP_SERVER" is required when provisioning an FPP server. SETUP_FPP_SERVER is the command action name and not a user defined name for the operation.

Sample input file:

```
ASMDBA GROUP=oinstall
ASMOPER GROUP=oinstall
ASM DISK LISTS=/mnt/emracdev/emracn0607/ASMDISK/disk1,/mnt/emracdev/
emracn0607/ASMDISK/disk2,/mnt/emracdev/emracn0607/ASMDISK/disk3
ASM DISK STRINGS=/mnt/emracdev/emracn0607/ASMDISK/disk*;/mnt/emracdev/
emracn0607/MGMTDB/disk*
AUTOVIP ENABLED=false
CLUSTER NAME=emracn0607
CONFIGURE ASMFD=false
CONFIGURE GIMR=true
CONFIGURE MODE=CONFIGURE ASM
CVU CHECK ERRORMODE=stop
DB COMPONENT=Fleet/OH19/linux64/DB19 18 - OH19 RU18
DB HOST NORMAL CREDNAMES=CUSER:SYSMAN
DB HOST ROOT CREDNAMES=CUSER ROOT:SYSMAN
DB INS CS GUID=0
DB ORACLE BASE LOC=/scratch/cuser/19c/racbase
DB ORACLE HOME LOC=/scratch/cuser/19c/rachome
DB PRODUCT VERSION=19.0.0.0.0
DEPLOY MODE=DEPLOY GI
DISKGROUP AU COMPATIBLE=DATA:4:::;MGMTDB:4:::
DISK GROUP NAMES=DATA, MGMTDB
FLEX ENABLED=false
GIMR DB ORACLE BASE LOC=/scratch/cuser/19c/racbase
GIMR DB ORACLE HOME LOC=/scratch/cuser/19c/rachome
GIMR DISK GROUP=MGMTDB
GIMR TYPE=NOGIMR
GI AU SIZE=4
GI CLUSTER TYPE=1
GI COMPONENT=Fleet/GI1919/linux64/GI 19 19 - GI 19 RU19
GI HOME SHARED=false
```



```
GI HOST NORMAL CREDNAMES=CUSER:SYSMAN
GI HOST ROOT CREDNAMES=CUSER ROOT:SYSMAN
GI ORACLE BASE LOC=/scratch/gridbase
GI ORACLE HOME LOC=/scratch/gridhome
GI PRODUCT VERSION=19.0.0.0.0
GI STORAGE TYPE=1
GNS ENABLED=true
GNS VIP ADDRESS=emracp1011-gns.subnet1rg.emsubnet.companyvcn.com
GRID DISK FLAG=true
GROUP VERSION=19.0.0.0.0
INTERFACE PARAM VALUE=ens3:100.75.0.0:1,ens4:192.168.1.0:5
NOT RELINK DB HOME SELECTED=false
NOT RELINK GIMR DB HOME SELECTED=false
NOT_RELINK_GI HOME SELECTED=false
OCR VDSK DISK GROUP=DATA
OINSTALL GROUP=oinstall
OSASM GROUP=oinstall
OSBACKUPDBA GROUP=oinstall
OSDBA GROUP=oinstall
OSDGDBA GROUP=oinstall
OSKMDBA GROUP=oinstall
OSOPER GROUP=oinstall
OSRACDBA GROUP=oinstall
PAUSE AFTER PREREQ=false
PLATFORM ARUID=226
PREREQ CHECK ERRORMODE=continue
RAC DBCA MODE=CreateRACDatabase
RAC HOME SHARED=false
RAC STORAGE TYPE=1
REDUNDANCYS=NORMAL, NORMAL
ROLE NODENAMES=HUB, HUB
RUN CVU FROM OH=true
SCAN NAME=emracn0607-r.racdb-scan.emsubnet.companyvcn.com
SCAN PORT=1521
SKIP CVU CHECK=false
SKIP PREREQUISITE CHECKS=false
SOURCE TYPE=SOFTWARE LIBRARY
SSH SETUP=false
TARGET HOST LIST=emracn07.subnet1rg.emsubnet.companyvcn.com,emracn06.subnet
1rg2.emsubnet.companyvcn.com
VERSION GREATER 12102=true
VIP NODENAMES=emracn07-vip.subnet1rg.emsubnet.companyvcn.com,emracn06-
vip.subnet1rg.emsubnet.companyvcn.com
WORK DIR LOC=/tmp/
ASM SYS PASSWORD=oracle
FPP CONFIGURATION TYPE=server
FPP DISK GROUP=MGMTDB
FPP IMAGES MOUNTPOINT=/scratch/cuser/mount
REST USER NAME=restuser
REST USER PASSWORD=<password>
REST USER EMAIL=user@company.com
```

Provision an FPP Client and Grid Infrastructure

emcli submit_procedure -input_file=data:"/home/user_name/FPP/
gi client input.properties" -name="RACPROV2"



The parameter <code>-name="RAC_PROV"</code> is required when provisioning an FPP Client and Grid Infrastructure. <code>RAC_PROV</code> is the command action name and not a user defined name for the operation.

Sample input file:

```
ASMDBA GROUP=oinstall
ASMOPER GROUP=oinstall
ASM DISK LISTS=/mnt/emracdev/emracp1011/ASMDISK/disk1,/mnt/emracdev/
emracp1011/ASMDISK/disk2,/mnt/emracdev/emracp1011/ASMDISK/disk3
ASM DISK STRINGS=/mnt/emracdev/emracp1011/ASMDISK/disk*
ASM FAILURE AND QUORUM GROUPS=/mnt/emracdev/emracp1011/ASMDISK/disk1||
false,/mnt/emracdev/emracp1011/ASMDISK/disk2||false,/mnt/emracdev/
emracp1011/ASMDISK/disk3||false
ASM OCRVDSK FAILURE AND QUORUM GROUPS=/mnt/emracdev/emracp1011/ASMDISK/
disk1||false,/mnt/emracdev/emracp1011/ASMDISK/disk2||false,/mnt/emracdev/
emracp1011/ASMDISK/disk3||false
ASM SYS PASSWORD=sample1
AUTOVIP ENABLED=false
CLUSTER NAME=emracp1011
CONFIGURE ASMFD=false
CONFIGURE GIMR=false
CONFIGURE MODE=CONFIGURE ASM
DB ORACLE BASE LOC=/u02/app/oracle
DB ORACLE HOME LOC=/u02/app/oracle/app/product/19.0.0.0.0/db
DB PRODUCT VERSION=19.0.0.0.0
DEPLOY MODE=DEPLOY GI
DISKGROUP AU COMPATIBLE=DATA:4:::
DISK GROUP NAMES=DATA
FLEX ENABLED=false
GIMR DB ORACLE BASE LOC=/u03/app/oracle
GIMR DB ORACLE HOME LOC=/u03/app/oracle/app/product/21.0.0.0/gimrdb
GIMR DISK GROUP=DATA
GIMR TYPE=NOGIMR
GI AU SIZE=4
GI CLUSTER TYPE=1
GI COMPONENT=Fleet/OH19/linux64/DB19 - RU19
GI HOME SHARED=false
GI HOST NORMAL CREDNAMES=CUSER:SYSMAN
GI HOST ROOT CREDNAMES=CUSER ROOT:SYSMAN
GI ORACLE BASE LOC=/scratch/gridbase
GI ORACLE HOME LOC=/scratch/gridhome
GI PRODUCT VERSION=19.0.0.0.0
GI_STORAGE_TYPE=1
GNS ENABLED=false
GRID DISK FLAG=true
```



```
GROUP VERSION=19.0.0.0.0
INTERFACE PARAM VALUE=ens3:100.75.0.0:1,ens4:192.168.1.0:5
NOT RELINK DB HOME SELECTED=false
NOT RELINK GIMR DB HOME SELECTED=false
NOT RELINK GI HOME SELECTED=false
OCR VDSK DISK GROUP=DATA
OINSTALL GROUP=oinstall
OSASM GROUP=oinstall
PAUSE AFTER PREREQ=false
PLATFORM ARUID=226
PREREQ CHECK ERRORMODE=continue
RAC DBCA MODE=CreateRACDatabase
RAC HOME SHARED=false
RAC STORAGE TYPE=1
REDUNDANCYS=NORMAL
ROLE NODENAMES=HUB, HUB
RUN CVU FROM OH=true
SCAN NAME=emracp1011-r.racdb-scan.emsubnet.companyvcn.com
SCAN PORT=1521
SKIP CVU CHECK=true
SKIP PREREQUISITE CHECKS=false
SOURCE TYPE=SOFTWARE LIBRARY
SSH SETUP=false
TARGET HOST LIST=emracp10.subnet1rg2.emsubnet.companyvcn.com,emracp11.subne
t1rg2.emsubnet.company.com
VERSION GREATER 12102=true
VIP NODENAMES=emracp10-vip.subnet1rg2.emsubnet.companyvcn.com,emracp11-
vip.subnet1rg2.emsubnet.companyvcn.com
FPP SERVER TARGET NAME=emracn0607
FPP SERVER HOST USER=CUSER:SYSMAN
FPP CONFIGURATION TYPE=client
#PREREQS ONLY=true
WORK DIR LOC=/tmp/
```

Optional FPP Client Registration

FPP Client registration occurs automatically using Enterprise Manager, should you optionally need to register the FPP client manually use the following EM CLI or REST API operations.

- REST API: Register FPP Client
- EM CLI: Use command emcli db_software_maintenance -setup -registerFPPClient-input file="data:<fileLocation>"w with the following input file:

```
"serverName": "<Server Name>",
"clientName": "<Client Name>"
```

Opt in the FPP Image Series

After completing the prerequisites you have opt in your FPP Image Series to be able to create the Gold Image.

 Review available FPP Server series, by querying the FPP server and listing the image series. EM CLI, use the following command:

```
emcli db_software_maintenance -setup -getFPPSeries
-input file="data:<fileLocation>"
```

Where the input file contains the following:

```
"targetName": "<GI_DATABASE>",
"targetType": "oracle database"
```

- REST API, List Images in an FPP Image Series
- List all available FPP images.
 - EM CLI, use the following command:

```
emcli db_software_maintenance -setup -getFPPImages
-input file="data:<fileLocation>"
```

Where the input file contains the following:

```
"seriesId": "<Series ID>",
"seriesName": "<Series Name>",
"serverName": "<Server Name>"
```

- REST API, List Images in an FPP Image Series
- Opt in the FPP series, this step stores the image series information in Enterprise Manager.
 - EM CLI, use the following command:

```
emcli db_software_maintenance -setup -optInFPPSeries -
input_file="data:<fileLocation>"
```

Where the input file contains the following:

```
"seriesId": "<Image series ID>",
"currentVersionImageId":"<Image ID>"
```

REST API, Opt-In FPP Image Series



"currentVersionImageId": "<Image ID>" is an optional parameter, if it is not specified the latest image created in the series is marked as *current*.

Create an FPP Gold Image

After registration of your FPP server you need to create the Gold Image. For FPP patching the Gold Image will reside in the FPP server and not in Enterprise Manager.

Follow these steps to create and FPP Gold Image:

- Verify the Oracle Home name and location. From Enterprise Manager menu bar, navigate
 to the **Targets** drop-down menu and then select **All Targets**. In the All Targets page, in the
 upper left search field, type or copy the FPP Oracle Home name.
- 2. Run the following EM CLI command to create the Gold Image: emcli db_software_maintenance -createSoftwareImage -input_file="data:/<location>/ fpp_image". The input file must contain the information obtained in Step 1, the following is a sample input file:

```
IMAGE_NAME=19CGI_Generic
VERSION_NAME=19.16GIRU
IMAGE_DESCRIPTION=GI gold image for release 19C
REF_TARGET_NAME=OraGI19Home1.<subnet>.<subdomain>.<domain>.com_4079
REF_HOST_CREDENTIALS=CUSER:SYSMAN
fppRestUser=RestUser1
WORKING_DIRECTORY=/u01/tmp
```

3. Once the deployment procedure completes, check the status of the Gold Image with the emcli db_software_maintenance -getImages command. A table with the images is displayed, the following is a sample result:

______ ______ IMAGE ID IMAGE NAME IMAGE TYPE DESCRIPTION VERSION CREATION DATE OWNER ______ ______ E79A3E2EAE871EBFE0533081 <SERVER>:DB gold image FPP 2022-09-01 13:57:15.0 gold image for 19.0.0.0.0 SYSMAN 4564041B 19.16 19.16 _____ E7B5422D9914364EE0533081 <SERVER>:GI gold image FPP image for vesrio 19.0.0.0.0 2022-09-02 22:01:13.0 SYSMAN _19.16 4564CC9F _____



Tip:

The Image Type column will show Software Library for Enterprise Manager native images and FPP for FPP server images.

4. Subscribe the target Grid Infrastructure home to the Gold Image. This is a one time step and is not required when adding new versions of future releases to this Gold Image. To perform this step run the following EM CLI command:

```
emcli db_software_maintenance -subscribeTarget -target_name=<CLUSTERNAME>-
cluster -target_type=cluster -image_id=<IMAGE ID>
```

.

Once all steps have been completed successfully you can now deploy the Gold Image and complete Grid Infrastructure patching. For more information on Gold Image creations see: Create a Gold Image Using EMCLI.

Deploy a Gold Image and Patch Using FPP

You can deploy and patch using FPP and Fleet Maintenance by using the Fleet Maintenance UI, EM CLI, and REST API.

- FPP Patching Using Fleet Maintenance UI
- FPP Patching Using EM CLI
- FPP Patching Using REST API

FPP Patching Using Fleet Maintenance UI

The following are a series of use cases for patching using FPP and the Fleet Maintenance UI

- Single Instance Database Patching Using FPP
- · Deploy and Patch Independently, Starting with a Standby Database
- Deploy and Patch Together, Starting with a Standby Database

Single Instance Database Patching Using FPP

Follow these instructions to deploy the software and patch single instance databases using FPP.

- In Enterprise Manager click on Targets and select Databases. In the Databases page click on Administration and from the drop down menu select Fleet Maintenance.
- Select the FPP image to which the database is subscribed to, filter by Image selecting the Gold Image to be used, Target Type select Database Instance, and select Update under Operation. This will show only the databases subscribed to the image.
- Select all databases requiring an update and click Next.
- 4. In the Operation page Software Deployment will be automatically selected, ensure the information presented is correct and select a Start time. Enter the Home Name Prefix if your environment requires a custom Oracle Home name prefix, this feature is available starting with Oracle Enterprise Manager 13.5 Release update 20 (13.5.0.20). Enable the Update Database task by check marking it and select a Start time.



Note:

The scheduled time for the update can be either Immediately/Auto or Later. If you select **Immediately** the patching operation will begin immediately after the deployment operation is successfully completed. When selecting Later make sure there is ample time between operations in case of issues.

In the Options section, select the options most convenient for your environment. Specify **Credentials**, including the FPP Rest User credential and **Pre/Post Script** inputs and click **Next**.

5. In the Review click on Validate and select Quick to perform a quick validation or Full to perform a full validation (full validations submit a deployment procedure). Once the validation completes successfully, submit the operation by clicking Submit. Enter a name for the procedure and click Submit once again.

Note:

If you encounter a *Destination home not found* error when performing a Quick validation you will need to perform a Runbook Session. Navigate to **Enterprise**, then **Monitoring**, then **Runbook Session**. Select **Start Universal Runbook Session**, and initiate the session: **Fleet Maintenance Operation Failed: Home Deployment not recognized**. This will help identify the root cause of this issue.

6. Click on **Monitor Progress** to view the operation in the Procedure Activity Page. Once the procedure completes successfully have updated the primary databases.

Deploy and Patch Independently, Starting with a Standby Database

This is the recommended use case when patching Grid Infrastructure databases, first deploying the software to the primary and standby databases. Then patching the standby databases and finally the primary databases.

- 1. Deploy the software to the hosts of primary and standby databases.
 - **a.** In Enterprise Manager click on **Targets** and select **Databases**. In the Databases page click on **Administration** and from the drop down menu select **Fleet Maintenance**.
 - b. Select the FPP image to which the databases are subscribed to, filter by Image selecting the Gold Image to be used, Target Type use Cluster Database (Database Instance for non Grid Infrastructure databases), and select Update under Operation. This will show only the databases subscribed to the image.
 - Select all primary databases requiring an update and click Next.

Note:

You cannot select primary and standby databases in the same Fleet Maintenance session.

d. In the Operation page Software Deployment will be automatically selected ensure the information presented is correct and select a Start time. Enter the Home Name Prefix if your environment requires a custom Oracle Home name prefix, this feature is available starting with Oracle Enterprise Manager 13.5 Release update 20 (13.5.0.20).



Make sure the **Update Cluster** task is NOT enabled, as this is a deploy only step. In the Options section, under **Software Deployment** the Data Guard Options select **Standby Auto Deploy**, this will allow the software deployment to occur for both primary and standby databases. Specify **Credentials**, including the FPP REST User credential and **Pre/Post Script** inputs and click **Next**.

- e. In the Review click on Validate and select Quick to perform a quick validation or Full to perform a full validation (full validations submit a deployment procedure). Once the validation completes successfully, submit the operation by clicking Submit. Enter a name for the procedure and click Submit once again.
- f. Click on **Monitor Progress** to view the operation in the Procedure Activity Page. Once the procedure completes successfully you can update the standby databases.
- 2. Update the standby databases
 - a. In Enterprise Manager click on **Targets** and select **Databases**. In the Databases page click on **Administration** and from the drop down menu select **Fleet Maintenance**.
 - b. Select the FPP image to which the databases are subscribed to, filter by Image selecting the Gold Image to be used, Target Type use Cluster Database (Database Instance for non Grid Infrastructure databases), and select Update under Operation. This will show only the databases subscribed to the image.
 - The **Deploy** status should show **Complete** for both primary and standby databases.
 - c. Select all standby databases requiring an update and click **Next**.
 - d. In the Operation page **Software Deployment** will be grayed out. Enable the **Update Cluster** task by check marking it and select a **Start** time.



Set the schedule time for the update at a time that starts after the deployment completes. This could be the time in your maintenance window when the standby can be patched. Allow enough time for the software deployment to complete and time for corrective actions. If the deployment is already successfully complete, you may schedule this **Immediately**.

In the Options section, under **Software Deployment** the Data Guard Options select the option most convenient for your environment:

- Startup After Switch: Selected by default, this will be restore the standby to the same status as before the update.
- Disable Data Guard: Choose this option if the standby is managed by a Data Guard Broker.

Specify Credentials, including the FPP REST User credential and $Pre/Post\ Script$ inputs and click Next.

- e. In the Review click on **Validate** and select **Quick** to perform a quick validation or **Full** to perform a full validation (full validations submit a deployment procedure). Once the validation completes successfully, submit the operation by clicking **Submit**. Enter a name for the procedure and click **Submit** once again.
- f. Click on **Monitor Progress** to view the operation in the Procedure Activity Page. Once the procedure completes successfully you can update the primary databases.
- 3. Update the primary databases



- a. In Enterprise Manager click on **Targets** and select **Databases**. In the Databases page click on **Administration** and from the drop down menu select **Fleet Maintenance**.
- b. Select the FPP image to which the databases are subscribed to, filter by Image selecting the Gold Image to be used, Target Type use Cluster Database (Database Instance for non Grid Infrastructure databases), and select Update under Operation. This will show only the databases subscribed to the image.

The **Deploy** status should show **Complete** for both primary and standby databases.

The **Update** status should show **Complete** for standby databases.

- c. Select all primary databases requiring an update and click Next.
- d. In the Operation page the **Software Deployment** area will be grayed out. Enable the **Update Cluster** task by check marking it and select a **Start** time.



Set the schedule time for the update at a time that starts after the deployment completes. This could be the time in your maintenance window when the primary can be patched. Allow enough time for the software deployment to complete and time for corrective actions. If the deployment is already successfully complete, you may schedule this **Immediately**.

In the Options section, under **Software Deployment** the Data Guard Options select the option most convenient for your environment:

- Enable Data Guard: Specify if this database system is managed by Data Guard Broker and of you disabled the broker in step 2.
- Enable FSFO (Fast-Start Failover): FSFO can be enabled when DG broker is on.
- Startup Standby: Specify to startup the Standby databases after primary is patched.

Specify **Credentials**, including the FPP REST User credential and **Pre/Post Script** inputs and click **Next**.

e. In the Review click on Validate and select Quick to perform a quick validation or Full to perform a full validation (full validations submit a deployment procedure). Once the validation completes successfully, submit the operation by clicking Submit. Enter a name for the procedure and click Submit once again.

Note:

If you encounter a *Destination home not found* error when performing a Quick validation you will need to perform a Runbook Session. Navigate to **Enterprise**, then **Monitoring**, then **Runbook Session**. Select **Start Universal Runbook Session**, and initiate the session: **Fleet Maintenance Operation Failed: Home Deployment not recognized**. This will help identify the root cause of this issue.

f. Click on **Monitor Progress** to view the operation in the Procedure Activity Page. Once the procedure completes successfully you can update the primary databases.



Deploy and Patch Together, Starting with a Standby Database

This is an alternative to patching Grid Infrastructure databases, first deploying the software and patching to the standby databases. Then deploying and patching the primary databases.

- 1. Deploy the software and Update the standby database.
 - a. In Enterprise Manager click on **Targets** and select **Databases**. In the Databases page click on **Administration** and from the drop down menu select **Fleet Maintenance**.
 - b. Select the FPP image to which the databases are subscribed to, filter by Image selecting the Gold Image to be used, Target Type use Cluster (Database Instance for non Grid Infrastructure databases), and select Update under Operation. This will show only the databases subscribed to the image.
 - c. Select all standby databases requiring an update and click Next.



You cannot select primary and standby databases in the same Fleet Maintenance session.

d. In the Operation page Software Deployment will be automatically selected ensure the information presented is correct and select a Start time. Enable the Update Cluster task by check marking it and select a Start time. Enter the Home Name Prefix if your environment requires a custom Oracle Home name prefix, this feature is available starting with Oracle Enterprise Manager 13.5 Release update 20 (13.5.0.20).



The scheduled time for the update can be either Immediately/Auto or Later. If you select **Immediately** the patching operation will begin immediately after the deployment operation is successfully completed. When selecting Later make sure there is ample time between operations in case of issues.

In the Options section, under **Software Deployment** the Data Guard Options select the option most convenient for your environment:

- Startup After Switch: Selected by default, this will be restore the standby to the same status as before the update.
- Disable Data Guard: Choose this option if the standby is managed by a Data Guard Broker.

Specify **Credentials**, including the FPP REST User credential and **Pre/Post Script** inputs and click **Next**.

- e. In the Review click on Validate and select Quick to perform a quick validation or Full to perform a full validation (full validations submit a deployment procedure). Once the validation completes successfully, submit the operation by clicking Submit. Enter a name for the procedure and click Submit once again.
- f. Click on **Monitor Progress** to view the operation in the Procedure Activity Page. Once the procedure completes you have successfully updated the standby databases and can move to updating the primary databases.



- 2. Deploy the software and update the primary database.
 - a. In Enterprise Manager click on **Targets** and select **Databases**. In the Databases page click on **Administration** and from the drop down menu select **Fleet Maintenance**.
 - b. Select the FPP image to which the databases are subscribed to, filter by Image selecting the Gold Image to be used, Target Type use Cluster (Database Instance for non Grid Infrastructuredatabases), and select Update under Operation. This will show only the databases subscribed to the image.
 - The **Deploy** status should show **Complete** for both primary and standby databases.
 - c. Select all primary databases requiring an update and click Next.
 - d. In the Operation page Software Deployment will be automatically selected, ensure the information presented is correct and select a Start time. Enable the Update Cluster task by check marking it and select a Start time.



The scheduled time for the update can be either Immediately/Auto or Later. If you select **Immediately** the patching operation will begin immediately after the deployment operation is successfully completed. When selecting Later make sure there is ample time between operations in case of issues.

In the Options section, under **Software Deployment** the Data Guard Options select the option most convenient for your environment:

- Enable Data Guard: Specify if this database system is managed by Data Guard Broker and of you disabled the broker in step 2.
- Enable FSFO (Fast-Start Failover): FSFO can be enabled when DG broker is on.
- Startup Standby: Specify to startup the Standby databases after primary is patched.

Specify **Credentials**, including the FPP REST User credential and **Pre/Post Script** inputs and click **Next**.

e. In the Review click on **Validate** and select **Quick** to perform a quick validation or **Full** to perform a full validation (full validations submit a deployment procedure). Once the validation completes successfully, submit the operation by clicking **Submit**. Enter a name for the procedure and click **Submit** once again.

Note:

If you encounter a *Destination home not found* error when performing a Quick validation you will need to perform a Runbook Session. Navigate to **Enterprise**, then **Monitoring**, then **Runbook Session**. Select **Start Universal Runbook Session**, and initiate the session: **Fleet Maintenance Operation Failed: Home Deployment not recognized**. This will help identify the root cause of this issue.

f. Click on **Monitor Progress** to view the operation in the Procedure Activity Page. Once the procedure completes successfully have updated the primary databases.



FPP Patching Using EM CLI

To deploy software images and patch using EM CLI use the following examples.

Deploy a Gold Image for a Database

To deploy a software image use the EM CLI command

```
emcli db_software_maintenance -performOperation
-name="Deploy 19.14 RAC Home"
-purpose=DEPLOY_RAC_SOFTWARE
-target_type=rac_database
-target_list="demodbfpp"
-normal_credential="CUSER:SYSMAN"
-input_file="data:/<location>/deploy_input.prop"
```

The input file must contain the following parameters:

```
NEW_ORACLE_HOME_LIST=/scratch/<user>/<location>/swDeploy
dispatchLoc=/scratch/cuser/workDir
workingDir=/scratch/cuser/workDir
SKIP_PREREQUISITE_CHECKS=false
workingCopy=EM_WORKING_COPY_1
*workingCopyPrefix=EM_WORKING_LIST
homeNamePrefix=<User defined Oracle Home name>
```

Note:

- If there is more than one target in the target_list user will have to specify workingCopyPrefix instead of workingCopy.
- homeNamePrefix=<User defined Oracle Home name>: Allows you to define a
 custom Oracle Home name prefix, allowing for the ability to discover Oracle
 Homes with custom name in deployment command. If the home name prefix is
 not specified, the default value will be used. This variable is available starting
 with Oracle Enterprise Manager 13.5 Release update 20 (13.5.0.20).

Patch a Database

To patch a database use the EM CLI command

```
emcli db_software_maintenance
-performOperation -name="Update dbfpp"
-purpose=UPDATE_RACDB
-target_type=rac_database
-target_list="demodbfpp"
-normal_credential="HHKAUR:SYSMAN"
```



First deploy the software on all standby databases, then patch all standby databases. Once the standby process is complete, repeat on primary databases by first deploying the software and then patching.

FPP Patching Using REST API

To deploy software images and patch using REST API use the following examples.

To deploy and update a database using FPP and REST API use the following sample files: Database Maintenance- Using Oracle Fleet Patching And Provisioning (FPP) REST Endpoints.

Opt Out of the FPP Series

If you need to opt out of the FPP series, use the following commands:

To opt out using EM CLI run the following command emcli db_software_maintenance setup -optOutFPPSeries -input_file="data:<fileLocation>" with the input file
containing either the SeriesID or the SeriesName and ServerName. Use one of the two
options listed below:

To opt out using REST API, use the following link: Opt-Out Image Series

Patch Exadata PDBs Using Fleet Maintenance

With Enterprise Manager's Fleet Maintenance you can easily patch and upgrade PDBs running on an Exadata Cloud@Cutomer (Exa C@C). This feature is available with Oracle Enterprise Manager 13.5 Release Update 21 (13.5.0.21) on up, including Oracle Enterprise Manager 24.1ai (24.1.0.0).

Follow these steps in order to successfully patch and upgrade your Exadata Cloud PDB's

- 1. Prerequisites
- 2. Obtain Patch Recommendations
- 3. Create a CDB With Target Patch Version
- 4. Gold Image Operations
- 5. Validate Patching Operations

Prerequisites

- The source CDB to be used should be discovered in EM
- The source CDB to be used must be subscribed to an image
- The source CDB must contain all patch versions to be applied
- Ensure all Fleet Maintenance prerequisites are met, for more information see: Before you begin: Meeting the Prerequisites

Obtain Patch Recommendations

Before you begin patching, you will need to review and identify the list of vulnerable targets, as well as the list of patches to be applied to these vulnerable targets. Fleet Maintenance Hub provides patch recommendation insights, assisting you in identifying the list of images, their subscribed vulnerable targets, and the patches that need to be applied. For detailed information on Fleet Maintenance Hub and patch recommendations see: About

For detailed information on Fleet Maintenance Hub and patch recommendations see: About Database Fleet Maintenance.

Create a CDB With Target Patch Version

Fleet Maintenance uses out-of-place patching, which requires that you already have a destination CDB pre-created and discovered in Enterprise Manager, with the target patches applied to it. There are two different methods to create a CDB with the target patches applied, use the method that best works for your environment:

- Create a CDB with the default Oracle Home on Exadata Cloud@Customer: For detailed information see Using the Console to Create a Database in Oracle Exadata Database Service on Cloud@Customer. For a video tutorial see Create Oracle Database on Exadata Cloud@Customer Video
- Creating a CDB using a custom Oracle Home on Exadata Cloud@Customer: To create a CDB using this method, follow the three steps below.
 - Create a custom software image: For detailed information see Creation and Storage of Software Images in Manage Software Images. For a video tutorial see: Create a Custom Database Software Image on Exadata Cloud@Customer.

- 2. Creating a custom Oracle Home using the software image: For detailed information see Create a Custom Database Software Image on Exadata Cloud@Customer. For a video tutorial see: Create Oracle Database Home on Exadata Cloud@Customer in Oracle Exadata Database Service on Cloud@Customer.
- 3. Create the CDB with custom Oracle Home: For detailed information see Creating First Database on an Oracle Exadata Database Service on Cloud@Customer System in Oracle Exadata Database Service on Cloud@Customer.

Discover the CDB

Once you have created the CDB it needs to be discovered in Enterprise Manager. This will allow it to be visible to Enterprise Manager and allow Fleet Maintenance to use it for patching. Follow the two steps outlined bellow:

- Discover the new Oracle Home: For detailed instructions see Discovering and Promoting Oracle Homes.
- Discover and add new CDB targets to the Oracle Home: Discovering and Adding Container Database and Pluggable Database Targets.

Gold Image Operations

A Gold Image represents a database release. A Gold Image consists of versions which are mapped to the Database version. When you create a Gold Image (sometimes referred as image), you create it with a version. Its advisable that you create only one Gold Image for a particular database release.

- 1. Create a Gold Image: If you have not created a Gold Image, or wish to create a new one follow the instructions outlined in Image Maintenance.
- 2. Update a Gold Image: In these steps you will create a new version of an existing Gold Image and mark it as current, see Create a new version of the image.

Deploy and Update Patches

- 1. To perform Fleet Maintenance operations, navigate to **Targets**, and select the **Databases** you wish to update.
- Once the databases have been selected. Click on Administration, then Fleet Maintenance. This will open the Fleet Maintenance home page,
 - a. Select the Gold Image to be used from the **Image** drop down menu.
 - b. Select target type Pluggable Database.
 - In the Operation drop down menu, select Update or Upgrade based on your requirement.
 - d. Select the PDB targets that you want to patch or upgrade and click **Next**.
- In the Operations screen, check mark the operations to be performed under Maintenance Tasks:
 - Under Maintenance Tasks, select Attach an existing CDB as Destination CDB.
 - b. In the Attach an Existing CDB section, select your Destination CDB.
 - c. Update PDB: Select a start time or set to start automatically after the previous step.
- 4. Enter the **Credentials** for Normal Host, Privileged Host and Database Credentials. You can select whether to use a Named or Preferred credentials. For Exadata Cloud RAC use the OPC user for Normal Host credentials.



- 5. Under **Options**, you can set **Work Directory**, and optionally upload **Custom Scripts** which can be executed either before or after update. Click **Next**.
- 6. In the **Review** screen, verify all options entered in the previous screen. Click on **Validate** to perform a **Quick** or **Full** validation of your Fleet Maintenance Operation before submitting.
- 7. To start your Fleet Maintenance operation, click **Submit**.

Validate Patching Operations

The DB Sync job introduced with the second Generation Exadata cloud synchronizes certain actions performed outside of the OCI framework with the Cloud UI. This is a fully automated process managed by a daemon that executes every 10 minutes and synchronizes the metadata with Cloud UI. After the patching operation is complete you can verify the status of the PDB on both EM and OCI console.

For more information see Differences Between Managing Resources with dbaascli and the Database API, and Gen 2 Exadata Cloud@Customer New Features: DB Sync Part-1.



Patching End-to-End Use Cases

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.

For specific Fleet Maintenance patching use cases see: Fleet Maintenance End to End Use Cases.

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

This use case 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 the data center plays a critical role in keeping the IT functions alive and business going.

Administrators in a data center carry out several data maintenance, data backup, and lifecycle management operations every day. The challenges they face in carrying out these system management activities are can be immeasurable. These activities 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

Enterprise Manager uses 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. Ensuring 100% compliance with your policies and standards.

Figure 22-1 Patching Your Datacenter Use Case



Identify the Patches Relevant to Your Data Center

Use Enterprise Manager's patch recommendations to identify the patches 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 patches based on your enterprise configuration.

Download, Prepare, and Test Patches

Analyze your environment and verify that the targets in your data center can be patched. Once you are sure they can be patched, download the recommended patches, test the patches, diagnose and resolve all patch conflicts.

Patch Using Database Fleet Maintenance

Follow the easy to use Fleet Maintenance UI that guides you through setting up and using Fleet Maintenance to Migrate, Deploy and Update your environments databases. You can also use the EM CLI verb <code>db_software_maintenance</code> to perform Fleet Maintenance activities in the command line.

Monitor Progress and Report the Status of the Change Activities

Monitor the progress of the various change activities, track the status of the patching activity, identify drifts, and report the overall status.

Executing the Example Scenario

The following table lists the tasks that will be performed in this example scenario.

- 1. Create Administrators with the Required Roles
- 2. Set Up the Infrastructure
- 3. Analyze the Environment and Identify Patches to be Applied
- 4. Identify the Relevant Patches
- 5. Execute Database Patching with Fleet Maintenance
- 6. Verify If the Targets Have Been Patched

Create Administrators with the Required Roles

You will need to create administrators with the follwing roles:

Table 22-1 Creating Administrators with the Required Roles

Enterprise Manager Role	Privileges
EM_CAP_ADMINISTRATOR	CREATE_JOB, CREATE_CAP_PLAN, BASIC_CAP_ACCESS
EM_CAP_USER	BASIC_CAP_ACCESS

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 online patching mode, this is the default patching mode in Enterprise Manager, you do not have to manually set this up the first time. If you set it to offline mode and want to reset it to online mode, or if you want to verify that the online mode is set, see: Setting up MOS Connection for Patch Recommendations and Patch Downloads.

If you have a proxy server set up in your environment, register the proxy details. To register the proxy server details with Enterprise Manager, see Setting up MOS Connection for Patch Recommendations and Patch Downloads.

Analyze the Environment and Identify Patches to be Applied

Before patching your databases, Oracle recommends you use the Software Standardization Advisor that scans the database configuration across all targets and generates a report (ondemand) 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.

To use the Software Standardization Advisor see Get Organized With Standardization Advisor.

Identify the Relevant Patches

View the Patch Recommendations region to identify the recommended and the relevant patches to be rolled out in your data center. These patches 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, and download them.

To view the recommended patches, see Downloading Patch Recommendations and Patches

Execute Database Patching with Fleet Maintenance

Once the patches have been identified and downloaded into the software library use Fleet Maintenance to select the databases to be patched by following the steps outlined in; Image Maintenance .

Check the Status of the Patching Activities

Track the status of the Fleet Maintenance activity and reports. For more information see *Track Activity* in Track, Fix and Resume Update Procedures.

Verify If the Targets Have Been Patched

Verify that the targets have been successfully patched.

Run the Oracle-supplied configuration search *Search Patches Applied on Oracle Products* from the **Configuration Search Library**. Search for the **Patch ID** that were applied to the targets, this will list all targets with that patch ID. For more information, see Managing Configuration Searches.

Fleet Maintenance End to End Use Cases

The following end to end use case Migration Workbench tutorials are available:

 Patching an Oracle Database RAC Container Database Using Oracle Enterprise Manager's Fleet Maintenance Graphic User Interface Patching and Oracle Pluggable Database Using Oracle Enterprise Manager's Fleet Maintenance Graphic User Interface



Part IV

Manage Configuration, Compliance and Change Management

This part contains the following chapters:

- Manage Configuration Information
- Manage Compliance
- Manage Database Schema Changes



Manage Configuration Information

This chapter explains how Oracle Enterprise Manager simplifies the monitoring and management of the deployments in your enterprise.

This chapter covers the following:

- Configuration Management
- Configuration Searches
- · Configuration Browser
- Configuration History
- Comparisons
- Configuration Extensions and Collections
- Parsers
- Relationships
- Configuration Topology Viewer

Configuration Management

Enterprise Manager 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 Enterprise Manager.

Enterprise Manager 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 23-1 provides a snippet of configuration information collected for a small sampling of target types as an example.

Table 23-1 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 23-1 (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 VM Company and to details
	VM Server member details
Client	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.
vCPU	vCPU is a Engineered Systems Management Framework having Virtualization
	enabled on it to gather data about maximum VCPU used by products or license editions deployed on the system. It reports maximum-vcpu consumed by individual
	engineered-systems and all engineered-systems of a type for 3 time periods hourly,
	daily, and monthly. Engineered Systems currently registered with ECM VCPU are:
	Exalogic system
	Exalytics system
	To enable collection of VCPU data for an Engineered System integrator for the system
	needs to register with ECM VCPU. Registration happens via MRS using xml.

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 Enterprise Manager 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

Comparisons

This section outlines the procedure for creating templates and utilizing rules, as well as details on configuring and administering Comparison Templates and Comparisons.

This section covers the following topics:

- About Comparisons
- Get Started Using Comparisons
- Get Started Using Rules
- Work with Comparison Results
- Comparisons and Drift Management with Oracle Enterprise Manager

About Comparisons

Enterprise Configuration Management involves the gathering, storing, and monitoring of configuration data associated with managed entities within your enterprise. For instance, a host has configuration items that relate to hardware and software components such as: the number of CPUs, memory, IO devices, operating system platform and version, installed software products, and so on.

Modifications to configuration data are inevitable often due to routine events such as patches and upgrades. A change to a single component can negatively impact the entire system, making it crucial to identify the root cause.

Enterprise Manager provides two types of comparisons:

- Configuration Drift: This allows you to compare the configurations of one target with those
 of another target of the same type.
- Configuration Consistency: Reflects the modifications made to target members within a system. For instance, you can use configuration consistency to guarantee that the configuration parameters for all databases in a cluster database are uniform.

Comparisons can be done on the current configuration or on previously saved configurations (before applying a patch or upgrade).

Comparisons allow you to:

- Select the particular attributes which need to be compared throughout targets
- · Notify key personnel when differences are detected
- Design and share comparison templates with other administrators



- Conduct comparisons of entire target systems; by automatically or manually matching target system members
- 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 re-comparison. If you have setup notifications, you will be notified of the changes.

About Comparison Templates

A comparison template is a blueprint for fine-tuning comparisons between similar configurations. The template is linked to a specific target type that defines the configuration item types, items, and properties to be compared. Oracle Enterprise Manager offers a default set of templates for specific target types. The template allows you to specify settings to consider during comparisons of the target type, such as which property differences to ignore or trigger an alert. You can also use constraints to set acceptable values for specific properties, any configuration that does not comply with the constraint is considered a difference.

The template applies rules or expressions to determine when a match for comparison purposes has been found and when differences detected in the comparison should be ignored.

Templates can be used as is or as a guide. For example, you may find that an existing comparison template with some modifications that meet your requirements. If the template disregards property differences that you consider important, you can use the **Create Like** feature to make changes to the existing template and save it under a new name.

For systems, create a system template that references member templates based on the target types that make up the system. We recommended you create the member templates first before creating the system template.

About One-Time Comparisons

One-time comparisons enable you to quickly view the differences between target configurations. Unlike drift and consistency comparisons one-time comparisons are only evaluated once. The results of a one-time comparison can be saved for future reference.

Before creating a drift or consistency comparison, Oracle recommends conducting a one-time comparison to verify that the initial drift or consistency evaluation will not reveal unexpected differences. With the results, you can correct configuration values or refine your comparison template to ensure that only relevant configuration items are compared.

To perform a one-time comparison:

- 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 helps maintain consistency across multiple targets, such as ensuring that the configuration of a database is uniform across all databases in the organization. This is especially important in today's business environment with frequent company acquisitions.

To perform a Drift Management Comparison:

- From the Enterprise menu, select Configuration, then select Comparison & Drift Management.
- 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 linked to a comparison, automatic comparisons will be run for it when necessary (for example: when the target's configuration changes). If differences are detected during an automatic comparison with the reference target notifications will be sent, and the differences will be displayed in drift comparison reports.

Before associating targets to a comparison, you can run a test comparison to verify that the comparison template is correctly configured and comparing relevant items. The **Save and Test Targets** option allows you to perform tests without sending notifications or affecting reports.

To test targets, go to the Test Targets page and click **Add** a target to compare. Run the test comparison by clicking **Run Test**. The number of differences found will be displayed in the **Differences** column of the table. To view the test comparison results, click on the **Difference Count**.

If the results are satisfactory, you can associate the target with the drift comparison by clicking **Associate**. This will enable 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 ensures that all members within a system or group have consistent configuration settings. For example, this can be applied to check that all database instances in a cluster database have uniform configuration parameters.

To perform a Consistency Management Comparison:

- From the Enterprise menu, select Configuration, then select Comparison & Drift Management.
- 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
- 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.
- 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 were

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 serves as the repository for all 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

Get Started Using Comparisons

This section describes how to create, edit, and otherwise manage comparisons and 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.

- From the Enterprise menu, select Configuration, click on Comparison & Drift Management.
- 2. In the Comparison & Drift Management page click the **Templates** tab icon on the left.
- Search for a template, select the target types that you want to search for; specify the Template Name, Owner name, and whether it is a default template or an Oracle provided template. Click Search.



Tip:

Templates have a lock by the template name.

- Closed lock: Oracle provided template that cannot be edited.
- Open lock: User defined templates that can be edited.
- 4. For a new template, click Create and enter a name, target type, and description; or click Create Like to create a template from an existing one and enter a name. These options create a new template row.
- 5. To edit a template select the template row and click on **Edit**, this takes you to the Template Details page.
- To create or edit the comparison template for each member, select the Member Settings tab. Select how
 - to view the Mapping Display in Hierarchical or Flat format.

You can edit the following:

- Select the member template to use for each system member type (optional).
- For any given member type, you can compare configurations by selecting the check box.

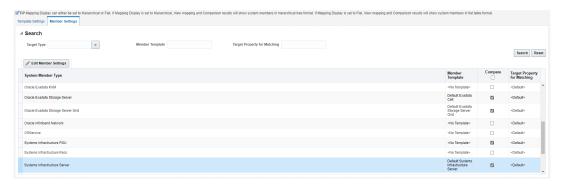




When you clear a check box for a system member, any associated child instance will be automatically ignored during that template's comparison.

Select a target property to use as a matching key for member types being compared.
 The default is target name, we recommend you use a distinguishable property to align comparison entities, like Department or Location.

Figure 23-1 Member Settings



In the Template Settings tab, select a Configuration Item to view its properties. A key icon
denotes a property that is defined as a key column in the configuration item type's
metadata.



Tip:

The **Compare** check box in the Template Settings tab allows you to streamline the comparison by selecting only the items you want to compare. When you select the check box, the comparison engine includes the item type and its children.

You can also compare individual columns and rows on the **Property Settings** tab. Settings are stored as part of the comparison results, you can view the compared properties in the Results page.

For example, when comparing host configurations you determine that CPU properties differences are not required. Expand the Hardware Configuration Item type and deselect the CPU check box excluding all its associated properties.

Figure 23-2 Template Settings



8. Click the Property Settings tab and check boxes for property differences to be compared and alerted if different. When comparing with property settings you are comparing for all differences detected in the property value for the Configuration Item type.
Use Value Constraint rules to filter property values, this compares the property value in the configurations being compared (second through n configurations) to the constrained value.

Any property value that fulfills the constraint constitutes a difference. See Get Started Using Rules for more information.

- 9. Repeat steps 7 and 8 to set property settings on other configuration items.
- 10. Optional, select an item and click the Rules for Matching Configuration Items tab. For a given property, specify a rule expression to be evaluated to determine when a match exists between configuration instances. See Get Started Using Rules for more information. 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. Optional, select an item 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 and use an AND logical operator within a subset of rules and an OR logical operator between rule subsets. 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.

Set Up a Comparison Template

To set up a comparison template, follow these steps:

- 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. Enter the Name, Target Type, and provide an optional description. When you provide the target type, configuration items for that target are automatically added to the template definition.
- 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.

- 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 of a target, when the system members change, or when the template changes.
- 7. If you choose **Table for the Mapping Display**, override the default flat map when defining the comparison. If not the members of the systems will be matched for comparison without



regard to their level in the system hierarchy. When viewing the results the original system hierarchy cannot be retrieved.



Chose **Table** option only if you are not concerned about how the members of the system are related to each other.

Managing Comparison Templates

You can manage Comparison Templates by doing the following:

- View a template's settings and composition (read-only)
- Delete a template (requires permission)
- 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.

- 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.

- Select a template in the Comparison Templates page, select the Actions menu and click Export.
- Select the save radio button and click OK.
- 3. Browse to the desired location in the file system, save the file and optional rename the file.



You cannot change the name of an Oracle provided template when exporting.



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, if the template owner is not the same as the login ID of the person importing the template it retains its original ownership. To change ownership of the imported template edit the owner attribute in the template XML file prior to import. You can also remove the attribute and it will set by default to the ID of the importer.



You cannot import an Oracle provided template of the same name. Change the name attribute in the template XML file prior to import

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

A follow-on step would be to review the results and drill down to differences details.

Set Up a Drift or Consistency Comparison

To create a drift or consistency comparison, follow these steps:

- 1. Determine what configuration items you want to compare.
- Create a template or use an existing template.
- Create a definition of either a Configuration Drift or a Configuration Consistency.
- 4. Perform the comparison.
- Set up email notification by setting up mail servers and then creating incident rules. See Create Notifications for Comparisons.



Create Notifications for Comparisons

Get notified every time there is a comparison change by creating notifications.

There are two parts to set up notifications: setting up the mail servers and creating an incident rule.

Follow these steps to setup notifications:

- From the Setup menu, select Notifications, then select Mail Servers.
- 2. On the Mail Servers page, provide the Server Identity information and 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.
- 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 by selecting one of the following:
 - a. Select:Incoming events and updates to events.
 - b. Select: Specific events of type Compliance Standard Rule Violation.
 - c. Select either:
 - Compliance Standard: Select a standard of type Drift or Consistency
 - 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 (comma separated) to notify if there is a drift alert.
 - f. Click **Next** and specify rule name and description. Click **Continue**.
- Click Save.

Exclude Target To Be Compared In a Template

Specific target types can be excluded from a comparison within a template when required.



The following SQL views are used to obtain a compliance score:

- MGMT\$COMPLIANT_TARGETS
- MGMT\$COMPLIANCE_SUMMARY
- MGMT\$COMPLIANCE TREND

Follow these steps to exclude targets from comparison templates:



- To start, create a new Generic System this will allow multiple types of targets to be compared. The Generic System must contain Host, Database Instance and Oracle Home Targets.
 - a. From the Targets menu select **Systems**, then click **Add**, and select **Generic System**. Enter a Name and enter optional information into Comment and System Properties.
 - b. Under Members click Add, a new pop up to Select Targets appears. Under Target Type select Host, Database Instance, and Oracle Home targets. This loads a list with these targets.
 - Select the targets to be added from the list and click **Select**.
 - c. Define Associations for the Generic System, click Add to add any associations. Once all associations are added click Next.
 - Setup the Availability Criteria for Key Members in the Generic Systems. Once complete click Next.
 - Specify the charts that will be shown in the System Charts page. Click on Add to add required charts. Once complete click Next.
 - f. Review information entered and click Finish to create the Generic System.
- Once the Generic System has been created, you will need to create a Generic System Template.
 - a. To create a Generic System Template, go to Enterprise, select Configuration, then select Comparison and Drift Management. In the Templates tab click on Create.
 - b. Enter a Name, and under Target Type select Generic System. Click OK.
 - c. Select the newly created template and click the **Edit** button.
 - d. In the Member Settings tab, select Host and Oracle Home targets for comparing and **Exclude** the Database Instance. Click **Save**.
- 3. Create a Consistency Definition for the Generic System with the template created in previous step.
 - a. Go to the **Definition Library** tab and click **Create**.
 - b. In the pop up select Consistency Definition, for Applicable To select Generic System, and for Comparision Template select the newly created template from the previous step. Click Continue.
 - c. Click Save and Associate Targets, click Add to add the Generic System created in the first Step. Once complete click Ok.
- 4. The Comparison Results contains all the Targets: Oracle Home, Database Instance, and Host.

Exclude Target To Be Compared From Reference Target

Targets can be excluded from comparison from a reference target, to do this follow these steps:

- 1. Create a Generic System with Host, Database Instance and Oracle Home Targets.
 - a. From the Targets menu select Systems, then click Add and select Generic System. Enter a Name and enter optional information into Comment and System Properties.
 - b. Under Members click Add, a new pop up to Select Targets appears. Under Target Type select Host, Database Instance, and Oracle Home targets. This loads a list with these targets.
 - Select the targets to be added from the list and click **Select**.



- c. Define Associations for the Generic System, click Add to add any associations. Once all associations are added click Next.
- d. Setup the Availability Criteria for Key Members in the Generic Systems. Once complete click **Next**.
- Specify the charts that will be shown in the System Charts page. Click on Add to add required charts. Once complete click Next.
- f. Review information entered and click **Finish** to create the Generic System.
- 2. Create a Generic System Template, select all the targets for Host, Oracle Home and Database Instance for comparison.
 - a. To create a Generic System Template, go to **Enterprise**, select **Configuration**, then select **Comparison and Drift Management**. In the Templates tab click on **Create**.
 - b. Enter a Name, and under Target Type select Generic System. Click OK.
 - c. Select the newly created template and click the **Edit** button.
 - **d.** In the Member Settings tab, select Host, Database Instance and Oracle Home targets for comparing. Click **Save**.
- Create Consistency Definition for the Generic System with the template created in previous step.
 - a. Go to the **Definition Library** tab and click **Create**.
 - b. In the pop up select Consistency Definition, for Applicable To select Generic System, and for Comparison Template select the newly created template from the previous step. Click Continue.
 - c. Click Save and Associate Targets, click Add to add the Generic System created in the first Step. Once complete click Ok.
- 4. Edit the Reference Target and Select the Host targets to ignore.
- 5. Save the changes to verify the result.
- 6. In the Comparison Results, all targets except Host will be visible.

Oracle Home Consistency With Target Properties

You can review an Oracle Home consistency through target properties. To do so follow these steps:

- 1. Create a Generic System with Oracle Home Targets.
 - a. From the Targets menu select **Systems**, then click **Add** and select **Generic System**. Enter a Name and enter optional information into Comment and System Properties.
 - b. Under Members click **Add**, a new pop up to Select Targets appears. Under Target Type select **Host**. Select the targets to be added from the list and click **Select**.
 - c. Define Associations for the Generic System, click Add to add any associations. Once all associations are added click Next.
 - Setup the Availability Criteria for Key Members in the Generic Systems. Once complete click Next.
 - Specify the charts that will be shown in the System Charts page. Click on Add to add required charts. Once complete click Next.
 - f. Review information entered and click **Finish** to create the Generic System.
- Create an Oracle Home Template.



- **a.** To create an Oracle Home Template, go to **Enterprise**, select **Configuration**, then select **Comparison and Drift Management**. In the Templates tab click on **Create**.
- Enter a Name, and under Target Type select Oracle Home. Click OK.
- **c.** Select the newly created template and click the **Edit** button.
- In the Member Settings tab, select the Oracle Home targets for consistency. Click Save.
- 3. In the Target Properties go to Edit Rule and set Property Name to Location.
 - a. To edit the rule go to **Enterprise**, select **Configuration**, then select **Comparison and Drift Management**. In the Templates tab, select the template click on **Edit**.
 - **b.** Under the Template Settings tab select **Target Properties**. Next move over to the Property Settings tab and select **Property Name**. Click on **Edit Rule**.
 - In the dialog box set the Operator to is equal to, and set Operands to Location. Click
 Ok.
 - d. Click Save.
- Create a Generic System Template.
 - a. To create a Generic System Template, go to **Enterprise**, select **Configuration**, then select **Comparison and Drift Management**. In the Templates tab click on **Create**.
 - b. Enter a Name, and under Target Type select Generic System. Click OK.
 - c. Select the newly created template and click the **Edit** button.
 - d. In the Member Settings tab, select Oracle Home targets. Click Save.
- Select the newly created Oracle Home Template in the Generic System and go to Member Template.
- Create Consistency Definition for the Generic System with the template created in previous sten.
 - a. Go to the Definition Library tab and click Create. In the pop up select Consistency Definition, for Applicable To select Generic System, and for Comparison Template select the newly created template from the previous step. Click Continue.
 - b. Click Save and Associate Targets, click Add to add the Generic System created in the first Step. Once complete click Ok.

Get Started Using 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**.



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]
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.
	['match-literal']
is greater than or equal†	A literal value to match; required
	match-literal
is greater than †	A literal value to match; required
	match-literal
is less than or equal to†	A literal value to match; required
	match-literal
is less than†	A literal value to match; required
	match-literal
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
	<pre>match-literal-1[,match-literal-n,]</pre>
is between†	A range specified as start and end literal values; both must be specified; range is inclusive
	start-range-literal , end-range-literal

Operator	Operands
contains†	A string literal on which to perform pattern matching; required
	[FALSE TRUE,] 'pattern-literal'
	FALSE (default) means string must comply with Oracle LIKE operator syntax; TRUE means string must comply with Posix regular expression syntax
replace‡	A string literal to match and replace with a second string literal
	[FALSE TRUE,]'pattern-literal'[,'replacement-literal'][,position-integer][,occurrence-integer]
	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
substring‡	Extract specified segment of string value
	<pre>[FALSE TRUE,]position-integer[,length-integer][,'pattern- literal'[,occurrence-integer]]</pre>
	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

Create 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:

- 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.

Create a Matching Rule

Specify matching rules as follows:

- 1. Select a configuration item in the left pane.
- 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.

Create a Rule for Including and Excluding Configuration Items

Specify ignore rules as follows:

- 1. Select a configuration item in the left pane.
- 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 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.

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.

Rule Examples

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.
- 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.
- 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

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.
- 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.
- 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.

- In the Rules for Including or Excluding Configuration Items tab, click New.
- 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.
- 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.
- 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.

Work 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 Configuration Extensions and Collections, for information on configuration extensions.



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.



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.

- 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 timestamped 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 effects 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.

Comparisons and Drift Management with Oracle Enterprise Manager

Comparisons and Drift Management offers new dashboard features for Oracle Enterprise Manager.

To access Comparison and Drift Management in Enterprise Manager, navigate to **Enterprise**, then select **Configuration**, and click on: **Comparison & Drift Management**.

This section contains the following:

- About the Dashboard
- Create a Drift Definition
- Create a Consistency Definition
- Create a One-Time Comparison
- Manage the Definition Library
- Manage Comparison Templates

About the Dashboard

The Comparison and Drift Management Dashboard lets you access your drifted and inconsistent targets from a single page.

The dashboard contains the following sections:



Dashboard

Dashboard

Consistency Overview

Comparison Overview

Compari

Figure 23-3 Comparison and Drift Management Dashboard

Drift Overview

Compare the latest or a saved target configuration to one or more targets. Some of the uses are:

- Ensure a target's configuration remains unchanged compared to a reference or a saved target configuration.
- Get notified when a target configuration change results in undesired differences.

Typical Scope: Large number of targets

Consistency Overview

Compare targets within a container for consistency. Some of the uses are:

- Ensure targets of a similar target type remain the same within a system.
- Ensure targets within a static or dynamic group remain consistent.

Typical Scope: Large number of targets

Comparison Overview

Compare the latest or a saved target configuration to one or more targets. Some of the uses are:

- Investigate problems where a change is suspected.
- Tune a comparison template before using it in Drift or Consistency Management.

Typical Scope: Small number of targets

Drift Results and Consistency Results by Target Type

A bar graph demonstrating the different target's drift and consistency results sorted by target type.



Library

The Library offers a way to see created Drift Definitions, Consistency Definitions, and Comparison Templates.

Create a Drift Definition

Create a Drift Definition to select a target and see when it drifts from the selected source.

To perform a Drift Management Comparison:

- From the Enterprise menu, select Configuration, then click Comparison & Drift Management.
- 2. Click Drift Results. In Drift Results, click Create Definition.
- In the Applicable To dialog box, select the Target Type and Template. For example, select
 Database Instance and the template of choice, in this case Database Instance Template.
 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.

Severity

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.

Create a Consistency Definition

Create a Consistency Definition to select a target type within a system and verify all targets have the same configuration.

To perform a Consistency Comparison:

- From the Enterprise menu, select Configuration, then click Comparison & Drift Management.
- 2. Click Comparison Results. In Drift Results, click Create Definition.
- 3. Select a location.



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
- 4. 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**.
- 5. On the Consistency Definition Details page, change the Definition Name as needed. The Compare Template and Applicable To information are replicated.
- 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.

7. 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.

Create a One-Time Comparison

Create a one-time comparisons 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.

To perform a One-Time Comparison:

- From the Enterprise menu, select Configuration, then select Comparison & Drift Management.
- Click One-Time Comparison Results.
- 3. Click Create Comparison.
- 4. 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.

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.

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.

Manage the Definition Library

Run management tasks from the Definition Library, the repository for all drift and consistency definitions.

To access the Definition Library:

- From the Enterprise menu, select Configuration, then select Comparison & Drift Management.
- 2. Click Definition Library.

From the Definition Library, you have access to the following tasks:

- Create a new definition.
- Edit an existing definition.



- Run tests against a definition.
- Search for 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.
- Import and export a definition.

Manage Comparison Templates

Run management tasks from the Comparison Templates screen.

To access the Comparison Templates screen:

- From the Enterprise menu, select Configuration, then select Comparison & Drift Management.
- Click Comparison Templates.

From the Comparison Templates screen, you have access to the following tasks:

- Create a new comparison template.
- Create Like an existing comparison template. This will import the configuration of an existing comparison template.
- Edit an existing comparison template.
- Search for a comparison template.
- Compare comparison templates.
- Delete a comparison template.
- Import and export comparison templates.

Compare Configurations with Scripts

You can compare configuration setups with EM CLI code and Perl scripts.

The following is sample Perl script that compares configurations by entering the databases to be compared, and their CDB (container database):

```
#! /bin/perl
#
#define map reference target => list of comparing targets
my %compare_map = ('<target.subnet.domain.com_CDBROOT>' =>
'<target.subnet.domain.com_PDB1,<target.subnet.domain.com_PDB2>',
'<target.subnet.domain.com_DB1' =>
'orcl.subnet1rg2phxsu.emdevinfraphx1.oraclevcn.com_CDBROOT,o<target.subnet.dom
ain.com_PDB2');

my %emcli_cmd = "./emcli_work_dir/emcli";
# target type is not the display name, please do the query in the db
# to get the correct target type name
# select target_type from em_ecm_config_template where
template_name='<template name>';
my %target_type = 'oracle_pdb';
my $template_name = 'pdb cp';
my $compare name root = 'comp test';
```



```
my $count = 1;
my $compare name;
my $cmd;
my $login cmd = "$emcli cmd login -username=\"sysman\" -
password=\"<password>\"";
#login to emcli session
system($login cmd);
print("build up one time compare..\n");
# Loop to create the one time comparation.
for(keys %compare map) {
  print("target type: $target type\n");
  print("template name: $template name\n");
  $compare name = $compare name root . $count;
  print("compare name: $compare name\n");
  count = count + 1;
  print("source target $_ , refer targets $compare map{$ }\n");
  $cmd = "$emcli cmd create config onetimecompare -name=\"$compare name\" -
template name=\"$template name\" -template owner=\"<SYSMAN>\" -
target_type=\"$target_type\" -compare_type=\"L\" -ref_target_name=\"$_\" -
dest target list=\"$compare map{$ }\"";
 print("execute : $cmd\n");
 system($cmd);
# logout from emcli session.
system("$emcli cmd logout");
```

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.
- Specify the owner of the configuration search.



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.

- Select the target type.
- Select Latest or History search type.
- Select the mode that was used to create the search, such as All, Modeler, or SQL.
- 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:

- 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.
- 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

Search for and import or export existing configuration searches using the: Configuration Search Library, Comparison Templates, Metric Extensions, Compliance Library pages.

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.
- 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.



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.**

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.
- 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 compare a saved configuration:

- From the Enterprise menu, select Configuration, then select Saved.
- 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:

- 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 export a configuration:

1. From the Enterprise menu, select Configuration, then select Saved.



2. Click **Export** and follow the instructions in the export dialog to save the results in a ZIP file. For a single target, the export dialog will determine the target's configuration items and immediate relationship details and will package the data into a zip file.

Or if the configuration already exists:

- From the Enterprise menu, select Configuration, then select Saved.
- 2. Select the configuration, then click **View**.
- Select Configuration Report and click Export.

To import a previously exported configuration:

- From the Enterprise menu, select Configuration, then select Saved.
- Click the Import button.
- 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:

- 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.
- 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.
- 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 *oralnst.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.

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:

- 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).



- 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.**

- 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.
- 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 using SQL, select Configuration Changes or Relationship Changes.
- In the Configuration History Modeler page, click Search to trigger the operation. A
 progress indicator verifies ongoing search activity. Results appear in the table at the
 bottom.



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 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 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 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 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 view annotiation history:

- 1. Select a row in the change details table and click **Details**.
- 2. In the window that pops up, click **Annotation Details**.

To annotate configuration changes:

- 1. Select the change row. To add the same annotation to multiple lines, use the multiselect feature (Ctrl+click or Shift+click).
- 2. Click Add Annotation.
- 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 popup 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.

- 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

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.

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

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.

- 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.
- 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**.
- 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**.
- 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.

- 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:

- 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.
- 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.
- 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 gueries.

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**.
- 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**.



- Enter the username and password for the credential set and click Save (or Test and Save for database credentials).
- 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 <code>SKIP</code> 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:

- 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.

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 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.

- In the Configuration Extensions library, select the specification table row and click View Details.
- 2. Peruse the settings and rules on the various tabs.

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.
- 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 autosynchronized monitoring facet	EM_COMPLIANCE_DESIGNER
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 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:

- 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.
- 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.
- Click Remove. A new column appears denoting a pending action of Undeploy; status remains Deployed.
- 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.
- 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:

- 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**.
- 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.
- 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.
- 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.
- 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**.
- After a while, a message confirms target type creation. Click Yes to add a sample target instance.
- 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.
- 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 httpdl.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.

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:

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:

- Access the Configuration Topology Viewer.
- 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:

- Access the Configuration Topology Viewer.
- 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.
- 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.
- 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.
- 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.
- 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:

- Access the Configuration Topology Viewer.
- 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:

- Access the Configuration Topology Viewer.
- 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 at 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...
- 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:

- 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.

- 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:

- Access the Configuration Topology Viewer.
- 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.
- 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:

- 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:

- Access the Configuration Topology Viewer.
- 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.
- To turn the coloring of the links on and off, on the Customize menu, select Highlight "Down" Root Cause.
- 6. To group targets:
 - 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.

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 parers, 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.

- 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.
- 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.
- 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.
- 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

Enterprise Manager 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 (~)
STORE_CONTENT_AS	Name to give to parsed attributes derived from element text content, where the element contains XML attributes; default is text_value
CONTAINER_NAME	A list of XML attribute names delimited by the value of the MULTIKEY_DELIMITER 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 Stooges that has attributes Moe, Larry, and Curly. As Moe appears first in the delimited list, its value, leader, becomes the parsed container name; Larry and Curly become parsed attributes. The tag name Stooges is discarded. The original XML fragment might be as follows:
	<pre><?xml version="1.0" encoding="UTF-8"?> <comedy></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:

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 <code>CONTAINER_NAME</code>, the container takes the value of the attribute (ajax), obviating the name=ajax parsed attribute. Remember that the <code>CONTAINER_NAME</code> 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).
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 IFTHENELSE 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.



Parser	Description
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.
Unix Recursive Directory List	Parser for output of Unix recursive directory listing ($1s -1 -R$). 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.
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 .ini type sections; default is true. $ \label{eq:continuous} % \begin{subarray}{ll} \end{subarray} % \b$

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 :.
COMMENT	Character that tells the parser to ignore the line that follows; default is #.

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.
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.



Parameter	Description
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 cron.allow and cron.deny files.
Cron Directory	Parser for Unix etc and cron.d 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 hosts.allow and hosts.deny files.
Kernel Modules	Parser for kernel modules files.
Linux Directory List	Parser for Linux directory listing data format (for example, output of a ls -l command).
PAM Configuration	Parser for pam.conf files.
PAM Directory	Parser for Unix etc/pam.d files.
Process Local	Parser for process.local files.
Secure TTY	Parser for Unix etc/securetty files.
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 ls -l command).
Unix Groups	Parser for Unix etc/group files. The parser ignores group name and password information.
Unix GShadow	Parser for Unix etc/gshadow files.
Unix Hosts	Parser for Unix etc/hosts files.
Unix INETD	Parser for Unix etc/inetd.conf files.
Unix Passwd	Parser for Unix etc/passwd files. The parser ignores password values.



Parser	Description
Unix Protocols	Parser for Unix etc/hosts files.
Unix Services	Parser for Unix etc/services.conf files.
Unix Shadow	Parser for Unix etc/shadow 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 $PATH=/a/b$.

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 <code>QUOTE_DELIMITER</code> 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, # [^\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.
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.



Parameter	Description
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.

Description
-
Parser for AIX installed packages files.
Parser for Apache HTTPD.conf files.
Parser for Autosys.jil files.
Parser for custom .cfg files. This syntax defines an element with $\mathbb{E}=\{\}$ syntax, where the brackets may contain name value pairs, nested elements, or both.
Parser for java.policy files.
Parser for java.properties files.
Parser for LDAP .cfg files.
Parser for mime.types files.
Parser for Radia .cfg files.
Parser for files containing name value pairs organized into sections, such as a Windows .ini file.
Parser for SiteMinder agent files.
Parser for SiteMinder .registry files.
Parser for SiteMinder SmReport.txt files.
Parser for SiteMinder SmWalker.dat files.
Parser for Sun ONE magnus.conf files.
Parser for Sun ONE obj.conf files.
Parser for Tuxedo files.
Parser for Unix etc/config files.
Parser for Unix etc/login.defs files.
Parser for Unix etc/proftpd.conf files.



Parser	Description
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 <code>java.properties</code> 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, $\# [^\r]^*$ 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_PROPE RTIES	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.
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.



Parameter	Description
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_PROP ERTIES	A flag that indicates whether sections support explicit properties. Default: false.
STRUCTURE_EXPLICIT_PR OPERTIES	A flag that indicates whether structures support explicit properties. Default: false.
NEWLINE_CONTINUE_LIN	A flag that indicates whether newlines can be line continuation sequences. Default: false.
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 newline 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 <code>KEYWORD_FIELD</code> 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 <code>DIRECTIVE_PROPERTIES</code> 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_. SiteMinder 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 = \setminus
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 (1) to indicate line continuation. The advanced properties parser parameters PROPERTY_DELIMITER and CONTINUE LINE define the respective format characters.

Example 2:

```
EC = \  EC2 = \setminus A = B, \setminus 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:



```
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:

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 <code>name/text()</code> 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
                                     host1
                    HOST
                                     1630
                    PORT
            ADDRESS LIST
                    FAILOVER
                                     on
                    LOAD BALANCE off
                    ADDRESS
                            PROTOCOL
                                              tcp
                            HOST
                                              host2a
                                              1630
                            PORT
                    ADDRESS
                            PROTOCOL
                                              tcp
                            HOST
                                              host2b
                            PORT
                                              1630
            ADDRESS
                    PROTOCOL
                    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>
```

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 $/{\tt HOST/text}$ () 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
                  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:



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 <code>SERVER_NAME/text()</code> 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.

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:

- Click Add.
- 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 dropdown 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.



Manage 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 Management
- Evaluate Compliance
- Configure Compliance Management
- Examples

Overview of Compliance Management

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 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
- Access the Compliance Features
- Roles and Privileges Needed for 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 pyramid 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 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.

Agent-Side Rule

Used to perform configuration checks on the agent and upload violations into the Management Repository.

Configuration Consistency Rule

Reflects changes of target members within a system. For example, configuration consistency rules ensures that the configuration parameters for all databases within a cluster are the same.

Configuration Drift Rule

Compares the configuration of a target with the configuration of another target of the same type.

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."

Missing Patches 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."

Monitoring Rule

Used to monitor actions to files, processes, and database entities as the changes occur. Also captures users logging in and logging out, and SU and SUDO activities.

Repository Rule

Used to perform a check against any metric collection data in the Management Repository

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.

Facets

The 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.

Observations

Observations are the actions that were seen on a host or target that were configured to be monitored through 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 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.



Access 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 and compliance standard rules.

Observations

Observations are the actions that were seen on a host or target that were configured to be monitored through 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 monitoring rule. Multiple UI-based reports are provided to allow users to analyze the actions that are being observed.

Roles and Privileges Needed for Compliance Features

To use the compliance standard features, you need to have access to the following roles and privileges.

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, and compliance standard rules.	 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 mange 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
			This privilege is part of the Compliance Framework resource privilege. This privilege is granted by default for EM_COMPLIANCE_OFFI CER role but it is not granted by default for the EM_COMPLIANCE_DESI GNER role.
Create Compliance Entity	Resource	EM_COMPLIANCE_ DESIGNER	Allows you to create compliance frameworks, compliance standards, and compliance standard rules. 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, and compliance standard rules. This privilege is part of the Compliance Framework resource privilege.
Compliance Framework	Resource	DESIGNER	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)
Configuration Future in a	Doca	EM COMPLIANCE	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



Privilege	Type of Privilege	Included in Role	Description
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 monitoring rule	EM_COMPLIANCE_DESIGNER role		



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.

Configure 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 Dashboard, then select Library.
- 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 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 and 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 for 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 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.
- 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.
- Customize the fields as needed.

Ensure that the Compliance Framework name is different from the original compliance framework and any other existing compliance frameworks.

- Click Save.
- 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:

From the Enterprise menu, select Compliance, then select Library.

- 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.

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 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:

From the Enterprise menu, select Compliance, then select Library.



- Click the Compliance Frameworks tab.
- 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.
- 5. Click OK.

Browsing Compliance Frameworks

To browse a compliance framework, follow these steps:

- 1. From the Enterprise menu, select Compliance, then select Library.
- 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.
- 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:

- 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 Basic Installation Guide
- Oracle Enterprise Manager 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.
- 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 Enterprise Manager structure needs to be defined. The steps to provide this structure include the following:

- Super Administrator creates three Enterprise Manager users: Compliance Author, IT Administrator, and Compliance Auditor.
- 2. Super Administrator assigns the appropriate roles and privileges to the Compliance Author and IT Administrator.
- Super Administrator assigns the same target privileges to IT Administrator and Compliance Auditor.



- Compliance Author logs in to Enterprise Manager 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 Enterprise Manager and associates the targets for which he has target privileges with the appropriate compliance standards.
- 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.
- Compliance Auditor logs in to Enterprise Manager 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.

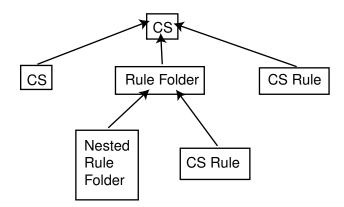
Compliance standards are made up of the following in a hierarchical structure:

- 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 24-1 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.
- 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**.

For information on available Enterprise Manager Compliance Standards see: Oracle Enterprise Manager Cloud Control Oracle Compliance Standards Reference.

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 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.

Enterprise Manager Release Updates may contain Compliance fixes and changes. When there are fixes to Compliance standards (rules fixed, new rules, or removal of outdated rules) a reassociation to the Compliance standard is required. This is because an Enterprise Manager environment has associated these standards to targets, and the corresponding results are invalid the moment the upgraded Enterprise Manager updates the modified standard or rule. Agent side standards require the updated rules to propagate from the OMS to the target(s), which happens as part of association. If the associations are not removed, the agent continues to run old version of rule checks. There are the following times you do not need to re-associate a Compliance standard:

- First time Enterprise Manager installation.
- If the release update has new standards, there is no impact.



 If the release update has compliance standards that are NOT currenlty being used in your environment.

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 < 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:

- Super Administrator creates three Cloud Control users: Compliance Author, IT Administrator, and Compliance Auditor.
- Super Administrator assigns the appropriate roles and privileges to the Compliance Author and IT Administrator.
- 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. The author 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. The administrator 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. The auditor 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
- Create Like for Compliance Standards
- 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 for 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 prompted for the Name, Author, target type to which the standard is applicable.. and the standard type. The standard types are:
 - Repository
 - Monitoring
 - · Agent-side

Click Continue.

- 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 included 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.

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:

- 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

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.

refer to production compliance standards and production compliance standard rules.

Associating a Compliance Standard

Associate the standard with a target and define the target type-specific settings.

- On the Compliance Standards Library page, ensure the correct compliance standard is highlighted.
- 2. Click the Associate Target button.
- 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.
- Click Select.

After you associate the targets with the compliance standard, you can edit the parameters associated with the target.

- While on the Target Association for Compliance Standard page, click Edit.
- On the Customize Compliance Standard Parameters page, change the parameters as needed.



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**.

Self Update for Compliance Standards

The Self Update for Compliance Standards feature allows you to update Compliance standards through Enterprise Manager components whenever new or updated standards become available.

This feature is accessed via the Self Update home page, a common dashboard used to obtain information about new updates and a common workflow to review, download and apply updates. With Self Update Compliance standards are released outside of Enterprise Manager Cloud Control major and minor release cycles, allowing for faster and up to date standards to be applied in your environment.

Before you begin, make sure your Enterprise Manager is already set up for Self Updates, for more information see: Setting Up Self Update in *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

Once you have successfully set up Self Update, there are two methods for applying updates:

- Online Mode: Enterprise Manager requires an active internet connection and access to the Enterprise Manager store to download the latest Compliance updates
- Offline mode: Is used in high security environments, where an active Internet connection between Enterprise Manager and the Enterprise Manager Update Store may not be available. This method utilizes a third party computer that is internet connect to the Enterprise Manager Update Store, update files are then transferred behind the firewall to Enterprise Manager.

The following is a list of self update enabled compliance standards:

- CIS Compliance Standards
- Security Technical Implementation Guide



For more information on applying updates under both methods see: Applying an Update in *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

Create Like for Compliance Standards

To create a Compliance Standard *like* another Compliance Standard using the existing standard as a reference source. Follow these steps:

- 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.

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.
- Update the parameters as needed.
- 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:

- 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.
- 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:

- From the Enterprise menu, select Compliance, then select Library.
- 2. Click the Compliance Standards tab.
- 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.
- 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:

- 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:

- From the Enterprise menu, select Compliance, then select Library.
- 2. Click the Compliance Standards tab.
- Highlight the compliance standard you want to associate with various targets. Click the Associate Target button.



- Select the targets you want to associate with this compliance standard. Click OK.
- With the compliance standard still highlighted, click the Override Target Type Settings button.
- 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.

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 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 you have privileges to access the group target you want to associate the compliance standards to. For more information see: Roles and Privileges Needed for Compliance Features

Perform the following steps:

- From the Enterprise menu, select Compliance, then select Library.
- Click the Compliance Standards tab.
- Highlight the compliance standard you want to associate with the group target. Click the Associate Groups... button.
- 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.

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:

- From the Enterprise menu, select Monitoring, then select Monitoring Templates.
- In the Search area, select Display Oracle provided templates and Oracle Certified templates and click Go.
- Select Oracle Certified-Enable Database Security Configuration Metrics and click Apply.
- In the Destination Targets region on the Apply Monitoring Template Oracle Certified-Enable Database Security Configuration Metrics: General page, click Add.
- 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 requirement. One may check that configuration settings are set properly. Another may be used to check in 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:

- From the Enterprise menu, select Compliance, then select Library.
- 2. Click the Compliance Standards tab.
- 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.
- Type the name of the folder and click OK.
- 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.

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 noncompliant 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.

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 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. For more information see: Roles and Privileges Needed for 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.
- 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.

- 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 < 100</WhereClause>

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.
- 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.

- 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.
- Click Continue.
- 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:

- 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.

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 compliance="" of="" rule="" standard=""> is compliant.</name>
	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 compliance="" of="" rule="" standard=""> is not compliant.</name>
	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.
- Click the Create button.
- 4. In the Create Rule popup, select **Configuration Consistency Rule** as the type.
- 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.

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.
- 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:

- 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.
- Customize the fields as needed.
- Click Save.

Editing a Compliance Standard Rule

To edit a compliance standard rule, follow these steps:

- From the Enterprise menu, select Compliance, then select Library.
- 2. Click the Compliance Standard Rules tab.



- Highlight the rule you want to edit and click the Edit button.
- 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 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 monitoring rules, 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:

- 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 reimport the generated compliance standard rule definitions into another Management Repository.

To export a compliance standard rule, follow these steps:

From the Enterprise menu, select Compliance, then select Library.

- Click the Compliance Standard Rules tab.
- Highlight the rule you want to export.
- 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 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:

- From the Setup menu, select Incidents, then select Incident Rules.
- 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.
- 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.
- Select either All events of type Compliance Standard Rule Violation or Specific events of type Compliance Standard Rule Violation.
- 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.

Evaluate 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 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 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 the 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 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:

- Access Compliance Statistics
- View Compliance Summary Information
- View Target Compliance Evaluation Results
- View Compliance Framework Evaluation Results
- Manage Violations
- Investigate Compliance Violations and Evaluation Results
- Investigate Evaluation Errors
- Analyze Compliance Reports
- Overview of Compliance Score and Importance
- About the Compliance Dashboard

Access 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.

Use 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 be 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.

View Compliance Summary Information

Compliance summary information is available from the Enterprise Manager Compliance Results page and individual target home pages.

To view compliance summary information from the Enterprise Manager 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.
- 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.

View 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 monitoring observations.			
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 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 zeroInvoke a function with incorrect parameter values			

To view results using Cloud Control home page, follow these steps:

- From the Enterprise menu, select Compliance, then select Results.
- 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.

View 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**.
- 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.



If you notice a mismatch in dates in the Results page and in the Configuration Refresh page, runAll can be used to sync the collection date correctly. runAll fixes a stale or stuck collected date, the command does not force data collect. If there is no previously collected data before executing runAll, evaluations and scores will not be refreshed.

Manage 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:

- 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.

- Select one or more violations.
- 2. Click Suppress Violations.
- 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.
- Click Unsuppress Violations.
- 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:

- Select one or more manual rule violations.
- Click Clear Violations.
- 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.

Investigate 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.
- Navigate to the Rule Details tab under the Compliance Results page to see a summary of the Compliance rule, including its Type, Severity, State, Description, Rationale and Recommendation for resolution.

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.

- 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.
- 5. In the navigational list, click the **Secure Ports** compliance standard rule. In the resulting Secure Ports Summary, navigate to the Rule Detail tab. Here you will see a detailed summary of the rule, its Severity, Rule State (Production, Test), Description, Rationale and Recommendation for solutions.

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:

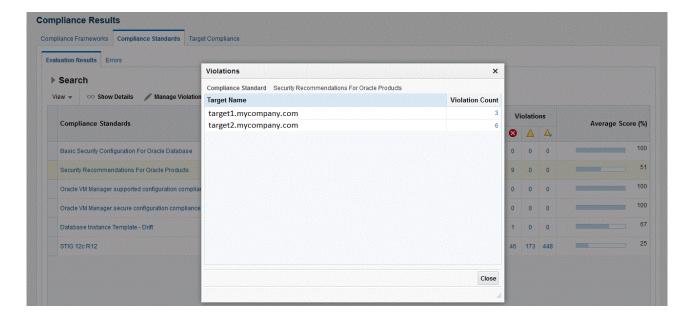
Example 1 - Accessing Violation Details of a Compliance Framework: To see the
violations of a compliance framework, click the Compliance Frameworks tab then the
Evaluation Results tab. The Violations columns list how many violations exist for each
framework. When you click the number in a Violations 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 24-2 Violations for a Compliance Standard

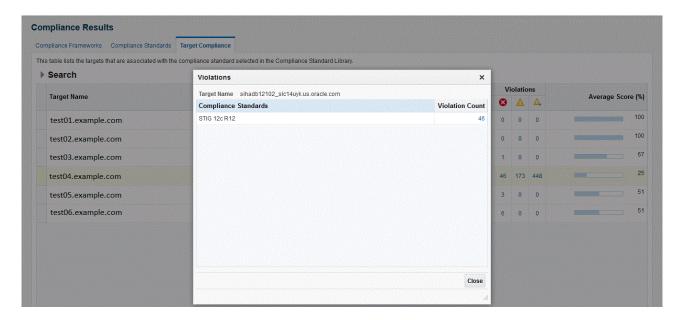


Again, click the number in the Violation Count column and the Violations pop-up appears. All the Compliance Standard Rules, for example Security Recommendations, are listed.

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.

Figure 24-3 Violations Using the Target Compliance Tab



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.

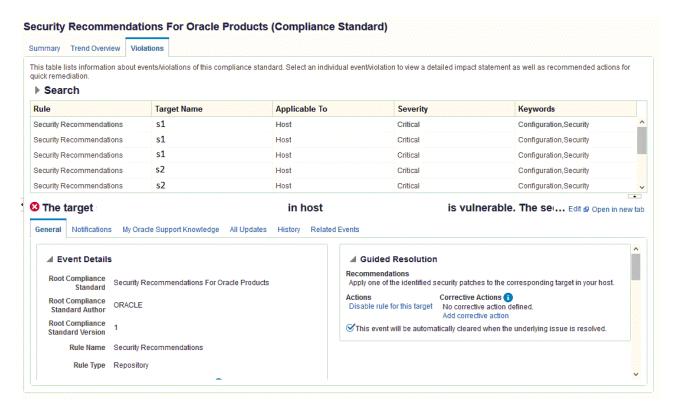
Figure 24-4 Show Details Page



On the resulting page, you have the option of seeing violations by target or by compliance standard rule.

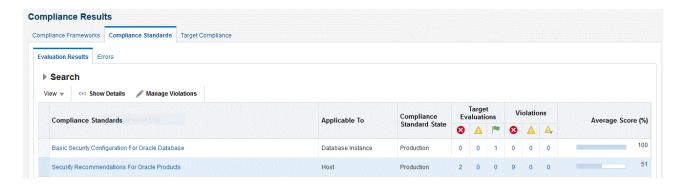
When you click the **Violations** tab, details regarding the compliance standard are listed including Event Details and Guided Resolution.

Figure 24-5 Event Details and Guided Resolution



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.

Figure 24-6 Compliance Summary Region on Enterprise Summary Page



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.
- Target evaluations are only one level while Violations are multi-level.

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.

Investigate 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.

Analyze 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.
- Suppressed Compliance Violations
 Displays to identify targets with disabled compliance rules details. It aims to determine which user disabled the rules the specific time and reason for suppressing the rules and explanation of violations.
- Configuration Extensions uploading large amount of data
 Configuration extensions provide a way to identify and upload configuration data
 changes more than 14 times per target last 7 days. These customized configuration
 extension metrics collections can be collected on well-known target types or a target
 type introduced as part of the configuration extension definition.

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 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:

```
hirange - (hirange - lorange) * (number of violations / number of rows evaluated)
```

The following table provides the combination of the severity and importance values used to calculate a compliance score.

Table 24-1 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

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.



^{(2) 0} is the lorange; 25 is the hirange

Figure 24-7 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

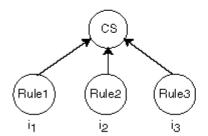
it: importance for Rulet

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

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 calculated based on weighted average of all the compliance score for each standard and target association within the compliance framework hierarchy. The weight is based on the importance of a compliance standard.

Parent Node Compliance Score

The compliance score of a hierarchy node/parent node is calculated as shown below. Compliance standards are hierarchical, thus the top node in the tree is known as the parent node.

Figure 24-8 Compliance Score of Parent Node

Compliance Score of Parent =
$$\frac{\sum_{i=1}^{|S_i| \times |S_i|}}{\sum_{i=1}^{|S_i|}}$$

i represents the number of children

- S is the score of the child node
- I is the importance of the child node

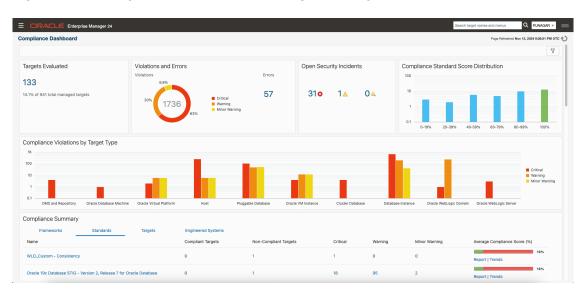
About the Compliance Dashboard

The Compliance Dashboard lets you access and manage your infrastructure's compliance from a single page.

To access the Compliance Dashboard, navigate to **Enterprise**, then select **Compliance**, and click on **Dashboard**.

The dashboard contains the following sections:

Figure 24-9 Compliance Dashboard with Engineered Systems Tab



Starting with Enterprise Manager 13c Release 5 Update 3 (13.5.0.3) the Compliance Dashboard features a new Engineered Systems tab, detailing information for AHF EXAchk compliance management of Exadata Engineered Systems targets, violations and reports. For more information on AHF EXAchk Compliance Standards see: AHF EXACHK Compliance Standards in *Oracle Enterprise Manager Cloud Control Oracle Compliance Standards Reference*.

Filters

The filter ribbon located at the top of the Compliance Dashboard allows you to filter by:

- Target Type: Filters by the type of monitored target in Enterprise Manager.
- Severity: Filters by the severity of the warning. Critical, Warning, Minor Warning.
- Properties: A set of values that are defined within the monitored target metadata, for example:
 - Contact
 - Downtime Contact
 - Cost Center
 - Department



- Line of Business
- Location
- Site
- Life Cycle

The filter ribbon will show all currently applied filters, you can close them by clicking on each filter. To apply new filters click on the **Filter** icon on the right hand side of the ribbon.



All widgets are filtered in real-time when filtering by target type, severity and properties.

Targets Evaluated

This widget shows all targets evaluated with the selected filters. Clicking the number will change the Compliance Summary to filter by the selected target type.

Violations and Errors

The Violations and Errors widget shows how your infrastructure compares to a compliance standard or framework regarding violations. Clicking a violation severity will filter the dashboard by severity. Clicking the error number will take you to the Compliance Standards Errors page.

To learn more about Compliance Standard Errors, see Manage Violations.

Open Security Incidents

The Open Security Incidents widget shows pending security incidents that require your attention. The incidents are sorted by severity and clicking the number will take you to Incident Manager where you can monitor and resolve service disruptions.

To learn more abut Incident Manager, see Using Incident Manager in *Enterprise Manager Cloud Control Administrator's Guide*.

Compliance Standard Score Distribution

The Compliance Standard Score Distribution graph filters your infrastructure by the selected target type and severity and sorts them by Compliance Score.

Compliance Violations By Target Type

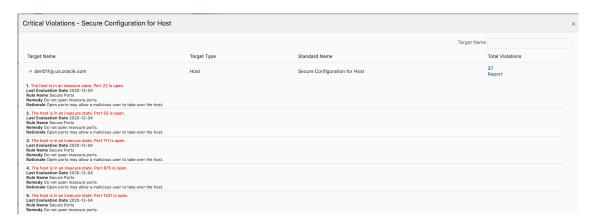
This graph presents the number of compliance violations sorted by target type and severity. Clicking any bar will filter the Dashboard by the selected target type and severity.

Compliance Summary

The Compliance Summary widget shows how your infrastructure compares to compliance frameworks and standards selected while evaluating compliance. Clicking a Violation number for any target will open a modal with all violations of that type (critical, warning, minor warning) in a single page. These violations can be expanded to display the description of the violation, last evaluation date, rule name, remedy, and rationale.



Figure 24-10 Compliance Summary



Compliance Evaluation Report

With the Compliance Evaluation Report you can access a full report of all violations for each target, for a particular standard, or for multiple targets or for multiple standards directly from the Compliance Dashboard. Follow these steps to obtain the reports:

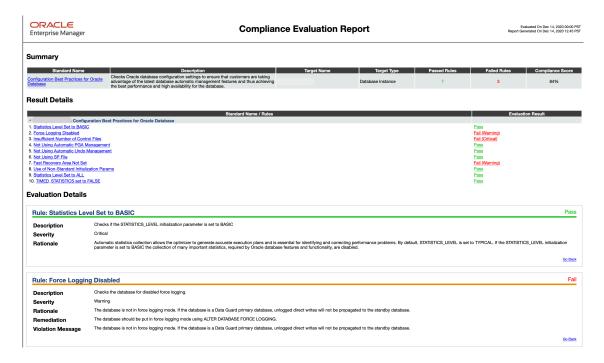
- On the Compliance Dashboard, navigate to the Standards or Targets tabs.
- Drill down into any of the Evaluations or Violations rolled up information (indicated by number) that you wish to review.
- 3. Click the Report link under the Average Compliance Score (%) column.
- 4. A new pop up appears click **Report** to view the Compliance Evaluation Report. This report provides a summary of the compliance findings as well as details on the standard(s) and rule(s) that are in violation and suggestions to resolve these violations.



This report can also be saved as a single HTML file, with embedded images, for easy sharing within your organization. Depending on the browser you use, to save the report, right click the Report link and select **Save Link As**.



Figure 24-11 Compliance Evaluation Report



DBSAT Reports

To generate these reports Enterprise Manager uses additional enhanced Oracle Database Security Assessment (DSAT) rules. There are two available DBSAT reports Security Assessment Report and Sensitive Data Assessment eport. For more information on the Oracle DBSAT Compliance Standard see: Oracle DBSAT Compliance Standard in *Oracle Enterprise Manager Cloud Control Oracle Compliance Standards Reference*.



Oracle Enterprise Manager 13 Release 5 Update 11 (13.5.0.11) is required to use the Sensitive Data Assessment Report.

To access these reports On the Compliance Dashboard

- Navigate to the Standards tab and locate the evaluation you wish to view the DBSAT report on.
- 2. In the Compliance Dashboard scroll down to the Compliance Summary and select either the **Standards** or **Targets** tab.
- Click on the number below Compliant Targets or Non-Compliant Targets.
- 4. A new pop up appears click **DBSAT Report**
- 5. In the DBSAT Report pop up there are two report options Security Assessment Report and Sensitive Data Assessment Report, click on the report of your choosing. These reports provide a summary of the DBSAT compliance findings as well as details on the



standard(s) and rule(s) that are in violation and suggestions to resolve these violations.





This report can also be saved as a single HTML file, with embedded images, for easy sharing within your organization. Depending on the browser you use, to save the report, right click the Report link and select **Save Link As**.

Figure 24-12 Oracle Database Security Assessment Report

Oracle Database Security Assessment

Highly Confidential

Assessment Date & Time

Date of Data Collection	Date of Report	Reporter Version
Fri Jan 28 2022 20:12:29 UTC+00:00	Fri Jan 28 2022 20:13:12 UTC+00:00	2.2.2 (June 2021) - 6003

Database Identity

A	Name	Container (Type:ID)	Platform	Database Role	Log Mode	Created
	NORMALDB	CDB\$ROOT (Root:1)	Linux x86 64-bit	PRIMARY	NOARCHIVELOG	Wed Sep 08 2021 21:44:42 UTC+00:00

<u>Summary</u>

Section	Pass	Evaluate	Advisory	Low Risk	Medium Risk	High Risk	Total Findings
Basic Information	0	0	0	0	0	1	1
User Accounts	7	1	2	2	1	0	13
Privileges and Roles	4	17	1	0	0	0	22
Authorization Control	0	0	2	0	0	0	2
Fine-Grained Access Control	0	0	5	0	0	0	5
Auditing	0	8	5	0	0	0	13
<u>Encryption</u>	0	2	1	0	0	0	3
Database Configuration	8	3	0	1	0	1	13



Figure 24-13 Oracle Database Sensitive Data Assessment

Oracle Database Sensitive Data Assessment

Highly Confidential

Assessment Date & Time Date of DBSAT Report Generation DBSAT Discoverer Version Thu Nov 03 2022 18:06:07 2.2.2 (June 2021) **Database Identity** Platform Database Role Log Mode **Date Created** S028NOV Linux x86 64-bit PRIMARY NOARCHIVELOG Wed Nov 02 2022 20:38:24 **Database Version** Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production **Discovery Parameters** Parameter Schema Scope ALL **Exclusion List File** myexclusion_list Minimum Rows Count 1 Pattern File(s) sensitive_en.ini

Summary

Sensitive Category	# Sensitive Tables	# Sensitive Columns	# Sensitive Rows
BIOGRAPHIC INFO - ADDRESS	7	13	1022
BIOGRAPHIC INFO - FAMILY DATA	1	1	630
BIOGRAPHIC INFO - RESTRICTED DATA	2	2	634
FINANCIAL INFO - CARD DATA	2	2	949

Examples

This section provides examples of using compliance. Examples include:

- Create Repository Rule Based on Custom Configuration Collections
- Create Compliance Standard Agent-side and Manual Rules
- Suppress Violations
- Clear Violations

Create 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:

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_css 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".

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**.
- 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.
- 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.
- On the Create Rule: Repository Rule: Check Definition (Query) page, click Model Query. New Search Criteria page appears.
 - 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.
- 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.



- On the Include Rule Reference popup, select compliance_css_rule. Click OK. Click Save to save the compliance_css_cs.
- 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.
- 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:

- 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.
- Click the Violations tab associated with the compliance_css_rule. The target is associated with one violation.
- 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.

Create 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
- Create a compliance standard
- 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:

- 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.
- Click the SQL tab.



- Click Add to add the first SQL statement.
 - In the SQL field, type:

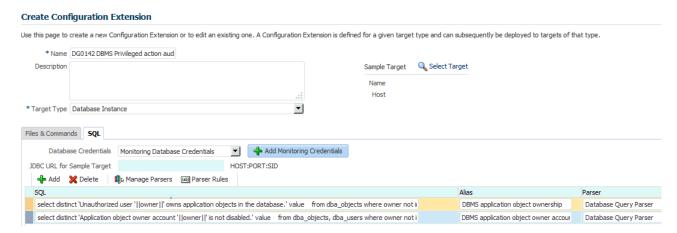
```
select distinct 'Unauthorized user '||owner||' owns application objects
in the database.' value
    from dba_objects where
owner not in ('ANONYMOUS','AURORA$JIS$UTILITY$',
'AURORA$ORB$UNAUTHENTICATED',
'CTXSYS','DBSNMP','DIP','DVF','DVSYS','EXFSYS','LBACSYS','MDDATA',
'MDSYS','MGMT_VIEW','ODM','ODM_MTR', 'OLAPSYS','ORDPLUGINS', 'ORDSYS',
'OSE$HTTP$ADMIN','OUTLN','PERFSTAT',
'PUBLIC','REPADMIN','RMAN','SI_INFORMTN_SCHEMA',
'SYS','SYSMAN','SYSTEM','TRACESVR', 'TSMSYSWK_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.

```
select distinct 'Application object owner account '||owner||' is not
disabled.' 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', 'TSMSYS',
'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 24-14 Completed Create Configuration Extension Page



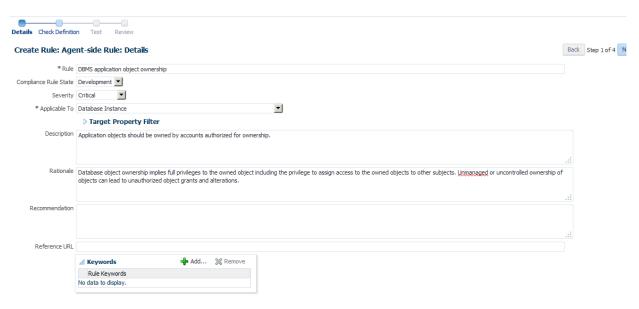
Creating an Agent-Side Compliance Standard Rule

To create an agent-side compliance rule:

- From the Enterprise menu, select Compliance, then select Library.
- 2. On the Compliance Library, click Compliance Standard Rules.
- Click Create. On the Create Rule pop-up, choose Agent-side Rule.
- Click Continue.
- 5. On the Create Rule, Agent-side Rule Details page provide the following information:
 - a. Name: DBMS application object ownership
 - b. Compliance Rule State: Development
 - c. Severity: Critical
 - d. Applicable To: Database Instance
 - 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.
 - g. Click Next.



Figure 24-15 Completed Compliance Standard Rule Details Page



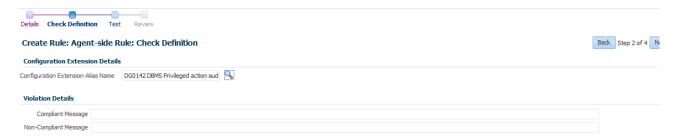
On the Create Rule: Agent-side Rule: Check Definition Page search for the configuration extension and alias you defined earlier.



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 24-16 Completed Compliance Standard Rule Check Definition Page



 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.



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 24-17 Completed Compliance Standard Rule Test Page



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**.

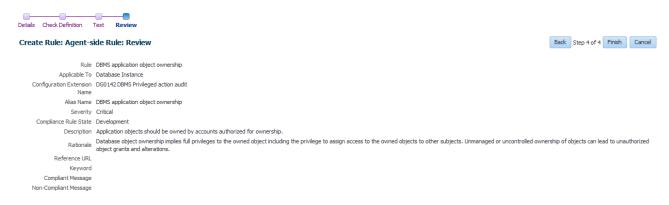


The compliance standard rule is not defined until you click **Finish**.

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 24-18 Completed Compliance Standard Rule Review Page



9. Repeat these steps for the second rule.



The compliance standard rule is not defined until you click Finish.



Creating a Manual Rule

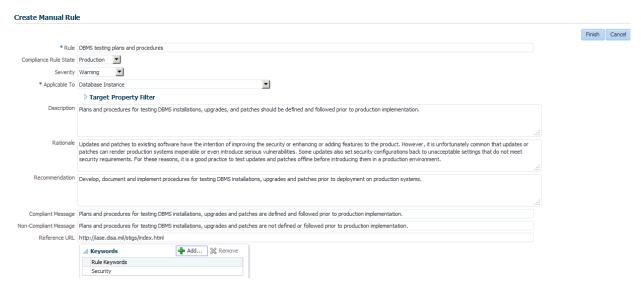
The purpose of creating this manual rule is to keep track of the checks that cannot be automated: ensuring that test plans and procedures have been followed prior to production.

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.
 - 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.
 - 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
 - I. Click Finish.



Figure 24-19 Completed Manual Rule Page



Creating a Compliance Standard

To create a compliance standard, perform the following steps:

- 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:
 - Name: CS1 DB Check
 - Applicable To: Select Database Instance
 - Author: SYSMAN
 - Standard Type: Agent-side
 - Click Continue

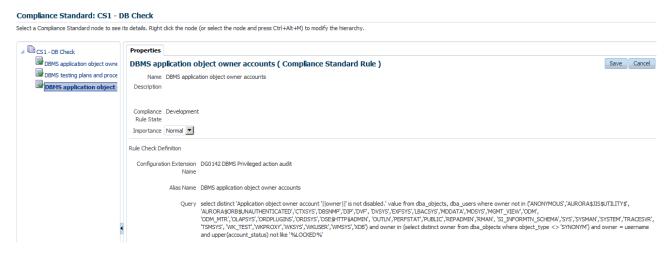
Figure 24-20 Completed Create Compliance Standard Pop-Up



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.
Click OK.



Figure 24-21 Compliance Standard Rules



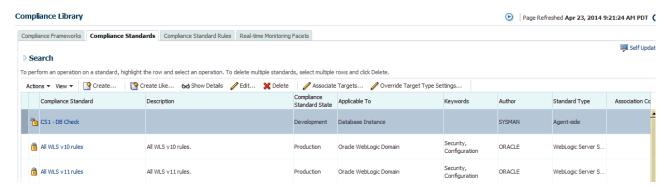
Click Save.

Associating the Compliance Standard to a Target

To associate the compliance standard to a target, perform the following steps:

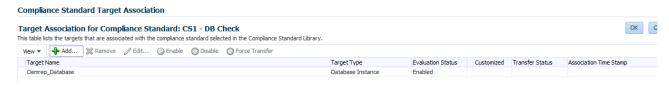
- From the Enterprise menu, select Compliance, then select Library.
- 2. On the Compliance Library page, click **Compliance Standards.**

Figure 24-22 Compliance Standards Library Page



- 3. 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.
- 5. Choose one or more targets, for example, Oemrep_Database.

Figure 24-23 Completed Target Association Page



6. Click Select. Click OK.



Click Yes to Save the Association.

Suppress Violations

The purpose of this example is to suppress violations. We will suppress the violation that arose due to the manual rule defined in Create Compliance Standard Agent-side and Manual Rules.

Follow these steps:

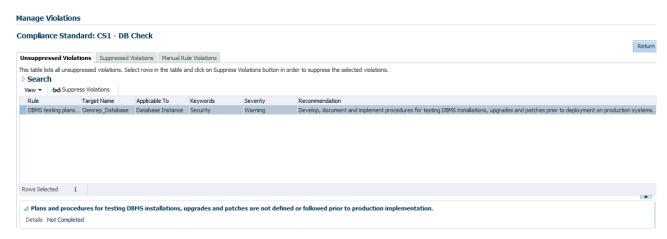
- 1. From the Compliance menu, select Results.
- In the Evaluation Results tab, locate the compliance standard named CS1 DB Check. Notice that there is a violation against the standard.

Figure 24-24 CS1 - DB Check Compliance Standard in Evaluation Results Tab



- 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.

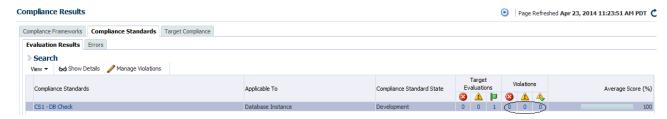
Figure 24-25 Manage Violations Page - Unsuppressed Violations



- 6. To suppress the violation, click the **Suppress Violations** tab.
- 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.



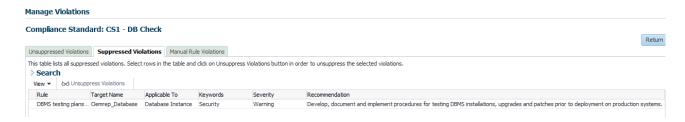
Figure 24-26 Evaluation Results Page After Violation Is Suppressed



To unsuppress the violation, use the Suppressed Violations tab. 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 24-27 Manage Violations Page Showing the Suppressed Violations Tab

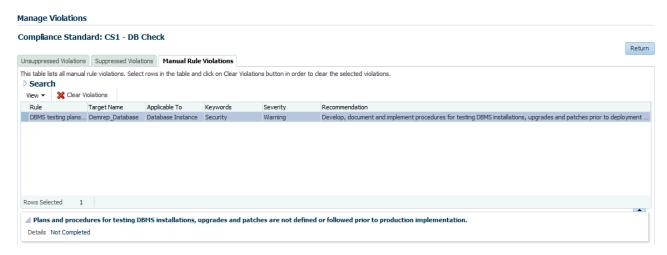


Clear Violations

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:

- 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.
- On the Manage Violations page, highlight the rule and click the Manual Rule Violations tab.

Figure 24-28 Clearing Manual Rule Violations



6. Select the rows and then click Clear Violations. On the Clear Violations Confirmation popup, select either Clear Violations Indefinitely or Clear Violations Until and specify a date. For completeness, provide a reason for clearing the violation.



Manage 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 Tablespaces) 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 versions 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 Enterprise Manager 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 Enterprise Manager 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.



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 recreating 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 restart 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.

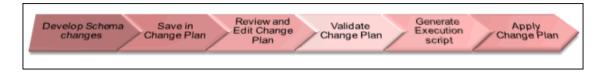


 Change Plans complement Change Management Schema Synchronizations, allowing for finer control of changes and "change-only" change requests.

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 25-1 Steps in a Change Plan



The image 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 Enterprise Manager

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

- 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



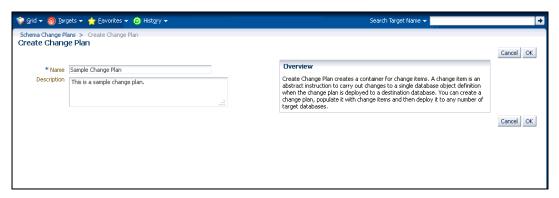
- Create Privileges for Job System (Resource Privilege)
- Create new Named Credential (Resource Privilege)
- Edit Resource Privilege on the change plans
- Execute Command Anywhere (Target Privilege)
- EM_ALL_OPERATOR privilege
- Ensure that the Database Administrator (DBA) is an Cloud Control user who 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 Enterprise Manager.

Creating a Change Plan

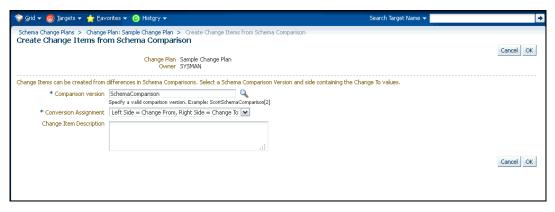
Follow these steps to create a change plan:

- Log in to Cloud Control as a database administrator (DBA).
- 2. Identify the schemas that contain application objects.
- 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.
- Save the baseline.
- Use Schema Comparisons wizard to define a comparison between the baseline version and the development database.
- Schedule a job to create the first version of the comparison and save the comparison.
- 7. In the Schema Change Plans page, click **Create**.
- 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:

- 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.
- 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 Enterprise Manager

Enterprise Manager provides support for external clients such as SQL Developer to create and access change plans. You can use these applications to connect to the Enterprise Manager 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:

- 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 nontrusted clients.
- 2. Set up an Enterprise Manager administrator for use by an external client.

The following section describes how to set up administrators for change plans, for access from Enterprise Manager and from external clients.

Setting Up Enterprise Manager Administrator For Change Plans

Follow these steps:

- 1. Log in to Enterprise Manager as a super administrator.
- 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 Enterprise Manager administrator. Click **Next**.
- In the Create Administrator: Roles page, click Next.
- In the Create Administrator: Target Privileges page, click Next.
- 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.

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, then, 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:

- 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>;
```

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 <code>\$ORACLE_HOME/network/admin directory</code>). The following wallet location must be specified at the reference database:

WALLET_LOCATION = (SOURCE=(METHOD=FILE) (METHOD_DATA=(DIRECTORY=/net/
slc05puy/scratch/dbwallets/s wallets)))

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.

Select your comparison name from the list, open the Actions menu, then select View Results. The Data Comparison Results page appears.

- 8. 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.
- 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.



Part V

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



26

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.

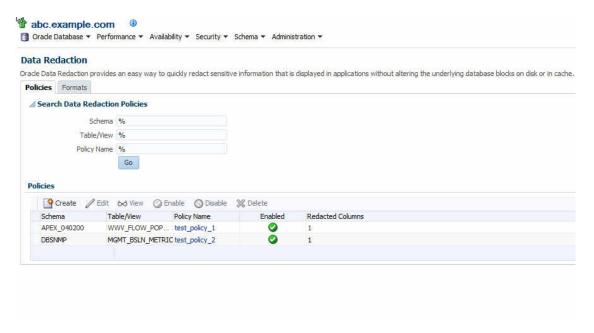
You can manage and monitor Oracle AVDF components in Enterprise Manager using the Oracle AVDF plug-in. For information on how to install this plug-in and manage Oracle AVDF using Enterprise Manager, see *Oracle Enterprise Manager System Monitoring Plug-in User's Guide for Audit Vault and Database Firewall.*



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 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 27-1 Data Redaction Page



For detailed information on using Oracle Data Redaction, see Using Oracle Data Redaction in Oracle Enterprise Manager in Oracle Database Advanced Security Guide.

28

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.*

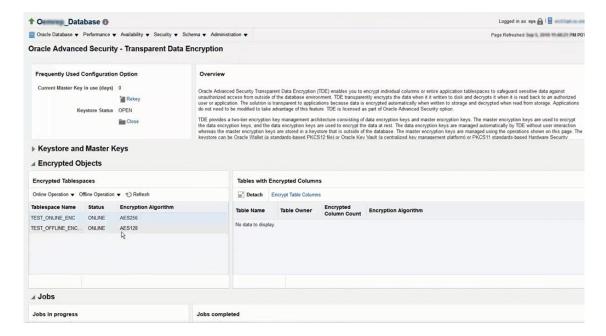


Figure 29-1 Transparent Data Encryption Home Page

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

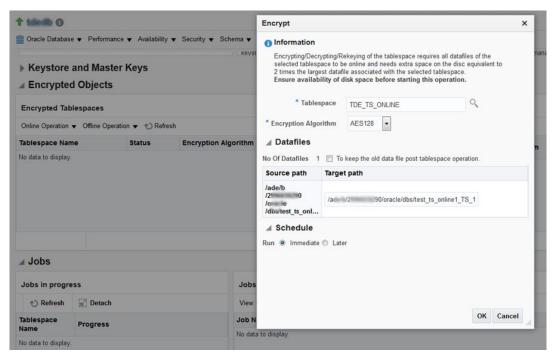
To encrypt or decrypt a tablespace using TDE, meet the following prerequisites:

- 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:

- 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.



- 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.
- 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.

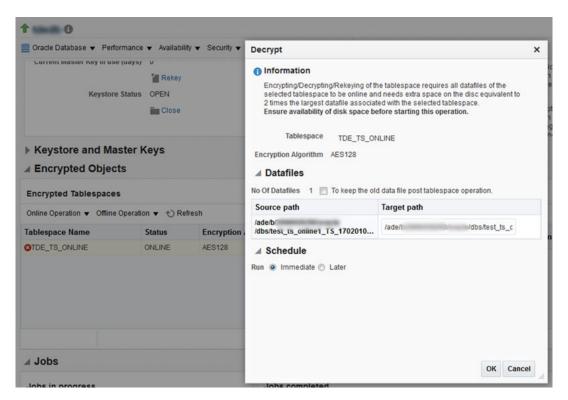


You may need to refresh the view to see the latest status of tablespaces in the Encrypted Objects section and the jobs.

Decrypting a Tablespace in Online Mode

To decrypt an encrypted tablespace, follow the steps below:

- From the Database home page, click Security and then select Transparent Data Encryption.
- 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.

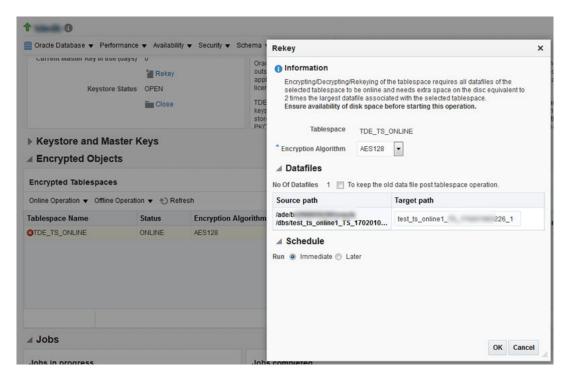


A job is scheduled to complete the encryption process. You can view the job status in the **Jobs** section on the Transparent Data Encryption home page.

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:

- 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.



- 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.

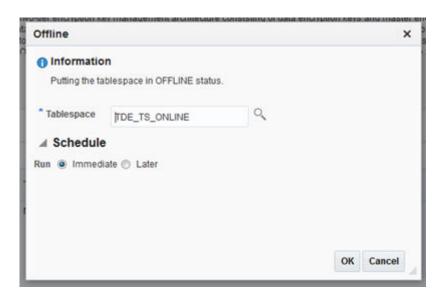
A job is scheduled to complete the encryption process. You can view the job status in the **Jobs** section on the Transparent Data Encryption home page.

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:



- 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:

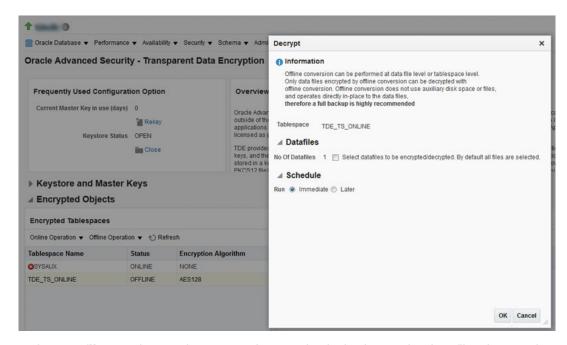
- From the Database home page, click Security and then select Transparent Data Encryption.
- 2. On the Transparent Data Encryption home page, expand the **Encrypted Objects** section.

- 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.

30

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:

- 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:

- 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.
- 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:

- Click Create Data Discovery Job.
- 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.
- 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.
- 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.

- 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:

- Open the Application Signature link from the Enterprise Data Governance dashboard.
 Click Create. The editor page opens.
- 2. Specify a name and optional description.
- 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.



A

Appendixes

This part contains the following appendixes:

- Using Enterprise Manager Command Line Interface
- Using emctl partool Utility
- End-to-End Use Case: Patching Your Data Center
- Troubleshooting Issues

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 Enterprise Manager 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





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.



To change the current kernel parameters, run the following command with root user privileges:

/sbin/sysctl -p

Parameter	Command
semmsl, semmns, semopm,	# /sbin/sysctl -a grep sem
and semmni	This command displays the value of the semaphore parameters in the order listed.
shmall, shmmax, and	# /sbin/sysctl -a grep shm
shmmni	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	<pre># /sbin/sysctl -a grep ip_local_port_range</pre>
	This command displays a range of port numbers.
rmem_default	<pre># /sbin/sysctl -a grep rmem_default</pre>
rmem_max	<pre># /sbin/sysctl -a grep rmem_max</pre>
wmem_default	<pre># /sbin/sysctl -a grep wmem_default</pre>
wmem_max	<pre># /sbin/sysctl -a grep wmem_max</pre>

Note:

For more information about the Kernel requirements, see the *Oracle Database**Installation Guide available in the following location: https://docs.oracle.com/en/database/oracle/oracle-database/index.html



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 you 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 glibcdevel gcc gcc-c++ libaio libaio-devel libstdc++ libstdc++-devel make compatlibstdc++ 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



3. 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 adaptors 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:

- 1. 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.
- 3. 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	Туре	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	
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)	or ASMCA, but	No	No	Not supported by OUI or ASMCA, but supported by the	
		Supported Storage Options for Oracle Clusterware and Oracle RAC / can be added			
	software. They can be added or removed after installation.			or removed after installation.	

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
Todalida ici providod by ordelo obravaror		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.



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

Using Enterprise Manager Command Line Interface

For information about Enterprise Manager Command Line Interface verb usage, syntax, and examples, see *Oracle Enterprise Manager Command Line Interface*

Overview

Enterprise Manager Command Line Interface (EM CLI) is a command line utility available for power users in Oracle Enterprise Manager that enables you to perform most of the console-based operations. It enables you to access Enterprise Manager 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 Enterprise Manager console.

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.

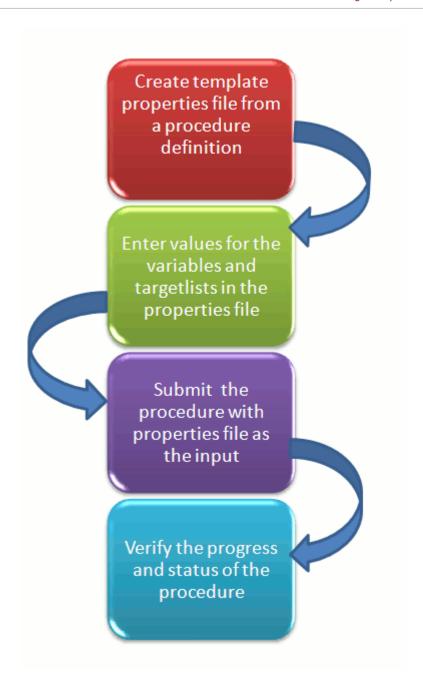
This section covers the following scenarios:

- 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:

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



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
instance}] [-notification={procedure status}]
[-grants={users and their corresponding accessing levels}] [-
schedule=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 -
procedure=B35E10B1F427B4EEE040578CD78179DC -schedule="start_time:2006/6/21 21:23;
tz:America/New_York" -grants="user1:VIEW_JOB; user2:FULL_JOB" -
notification="scheduled, action required, running"
```

Submitting the properties file using the procedure name/owner pair:

```
emcli submit_procedure -input_file=data:procConfiguration.properties -
name=SIHA_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, action
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:

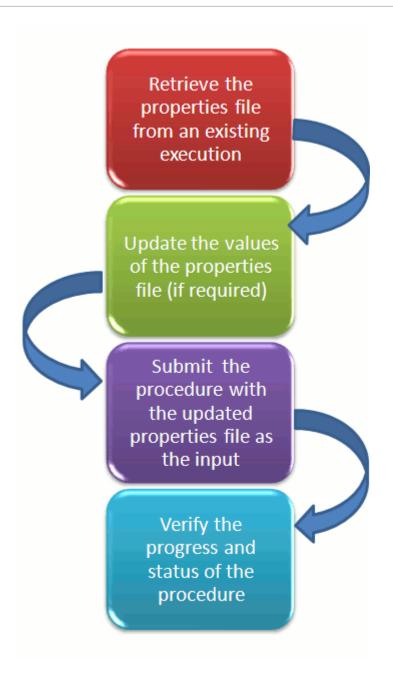
```
emcli get_instance_status
[-instance={instance guid}]
[-exec=execution guid]
[-xml
[-details]
[-showJobOutput]
[-tailLength={last N characters}]]]

Example:
emcli get_instance_status -instance=B35E10B1F427B4EEE040578CD78179F1 -details -showJobOutput

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 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_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
```

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_ssubbural, Failed
```

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 >
instanceData.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 <code>emcli get_procedures</code> fetches all the procedures which also include the Procedure Configurations.

To launch a procedure using an Existing Procedure Configuration File, do the following:

- Run the verb emcli get procedures to fetch an existing Procedure Configuration file.
- 2. Update the properties file if required.
- 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.



You can update the Procedure Configuration using the update_procedure_input verb as described in Updating the Procedure Configuration of a ProcedureAfter Updating the Procedure Configuration, follow Step 4 to resubmit the procedure.

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 location is SWLIB or not.>]
```

```
[-pdbTemplateInSWLIB=<If -sourceFromSWLIB, specify the URN of pdb template
component 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
qualified archive location>]
[-sourcePDBMetadataFile=<If -unpluggedPDBType=RMAN or XML, this is fully
qualified 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
container 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
specified, 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, - sourceCDBHostCreds, -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.>]
[-numOfPdbs=<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.>]
[-numOfPdbs=<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.>1
[-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 exported
files will be stored>]
[-importDir=<Temporary file system location on the CDB host used to stage the
migration metadata and/or datafiles>]
[-useOMF=<UseOMF for datafile location if CDB is OMF enabled <math>(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 encrypted
tablespaces>1
[-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 -
noncdbTargetName=NON_CDB_NAME -noncdbTargetType=oracle_database -
noncdbHostCreds=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 -
pdbAdminPassword=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 Prerequisites for Unplugging and Dropping a Pluggable Database.

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>
[-unplug {\tt PDBToSWLIB=<Specifies}\ that\ unplugged\ {\tt PDB}\ should\ be\ uploaded\ to\ {\tt 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 directory
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 qualified
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 - cdbTargetType=oracle_database -cdbHostCreds=HOST_CREDS -cdbTargetCreds=CDB_CREDS - pdbName=db_pdb -unplugPDBTemplateType=ARCHIVE -pdbArchiveLocation=/u01/app/unplugged/db pdb.tar.gz

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
- 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_without_env_shift_username1,
Prov_112_db_using_SH_locked_acc_without_env_shift_ssubbura11, 1.0, USERNAME1,
SIHA_SIDB_PROC

B35E10B1F42AB4EEE040578CD78179DC, DB_PROV_UPGRADE, DbProvUpgradeDP, Upgrade Oracle
Database, 1.0, ORACLE
B35E10B1F427B4EEE040578CD78179DC, DBPROV, SIHA_SIDB_PROC, Provision Oracle Database,
1.0, ORACLE

Select the GUID corresponding to the SIHA_SIDB_PROC, which is
B35E10B1F427B4EEE040578CD78179DC
```

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 Enterprise Manager.

Enterprise Manager supports the following use cases 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>, <Dsplay name>, <version>, <Parent procedure name>

For example:
    ./emcli get_procedures | grep FMWPROV_
B35E10B1F154B4EEE040578CD78179DC, FMW Provisioning, FMWPROV_DP, Provision Middleware, 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-1.

```
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-1 describes the parameters used in the aforementioned example.

Table A-1 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_PASSWOR D	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
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:

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-2.

```
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.O.name=example.com
ADMIN HOST.O.type=host
ADMIN HOST.O.normalHostCreds=NAME:<password>:user1
ADMIN HOST.O.SERVER ADDRESS=example.com
ADMIN HOST.O.IS SERVER MIGRATABLE=false
ADMIN HOST. 0. ORACLE HOME=n/a
ADMIN HOST.O.SERVER PORT=19030
ADMIN HOST.O.DOMAIN HOME ADMIN HOST=/scratch/oracle wls/mwBase/domains/base domain
ADMIN HOST.O.MIDDLEWARE HOME=/scratch/oracle wls/mwBase/middleware
ADMIN HOST.O.TLOG DIR=/scratch/oracle wls/mwBase/domains/base domain/servers/
Managed Server1/tlogs
ADMIN HOST.O.CLUSTER NAME ADMIN HOST=Cluster 1
ADMIN HOST.O.JDK LOC=/ade autofs/gd29 3rdparty/nfsdo generic/JDK7/MAIN/
LINUX.X64/150608.1.7.0.85.0B015/jdk7/jre
ADMIN HOST.O.SSL PORT=0
ADMIN HOST.O.EX PATTERN=*.mem
ADMIN HOST.O.MS OVERRIDE PORT=0
ADMIN HOST.O.CLONE MODE=true
ADMIN HOST.O.WORK DIR LOC ADMIN HOST=/tmp/scaleUpSrc
ADMIN HOST.O.MACHINE NAME ADMIN HOST=Machine 1
ADMIN HOST.O.WLS HOME ADMIN HOST=/scratch/oracle wls/mwBase/middleware/wlserver 10.3
ADMIN HOST.O.SOURCE SERVER NAME=Managed Server 1
ADMIN HOST.O.SERVER NAME ADMIN HOST=Managed Server1
DEST MANAGED SERVERS.O.name=example.com
DEST_MANAGED_SERVERS.O.type=host
DEST MANAGED SERVERS.O.normalHostCreds=NAME:<password>:user1
DEST MANAGED SERVERS.O.START SERVER REQUIRED=true
DEST MANAGED SERVERS.O.JOC PORT=9988
DEST MANAGED SERVERS.O.WLS HOME DEST MANAGED SERVERS=/scratch/oracle wls/mwBase/
```

```
middleware/wlserver 10.3
DEST MANAGED SERVERS.O.MACHINE NAME=Machine 1
DEST_MANAGED_SERVERS.O.SERVER_NAME_DEST_MANAGED_SERVERS=Managed_Server1
DEST_MANAGED_SERVERS.O.JRE_LOC=/usr/local/packages/jdk7
DEST_MANAGED_SERVERS.O.START_NM=false
DEST MANAGED SERVERS.O.START MS USE NM=true
DEST MANAGED SERVERS.O.NM LISTEN ADDRESS=example.com
DEST MANAGED SERVERS.O.CLUSTER NAME DEST MANAGED SERVERS=Cluster 1
DEST MANAGED SERVERS.O.DOMAIN HOME DEST MANAGED SERVERS=/scratch/oracle wls/mwBase/
domains/base domain
DEST_MANAGED_SERVERS.O.NM_LISTEN_PORT=5558
DEST_MANAGED_SERVERS.O.FMW_HOME_DEST_MANAGED_SERVERS=/scratch/oracle_wls/mwBase/
middleware
DEST MANAGED SERVERS.O.WORK DIR LOC DEST MANAGED SERVERS=/tmp/scaleUpDest
DEST MANAGED SERVERS.O.MS PORT DETAILS=19030:Listen Port
DEST MANAGED SERVERS.O.DEST MANAGED SERVER LISTEN ADDRESS=example.com
```

Table A-2 describes the parameters used in the aforementioned example.

Table A-2 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 exaple, 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.

Table A-2 (Cont.) Description of the Parameters Used in a Properties File for Scaling Up or Scaling Out a WebLogic Server

Parameter	Description
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.
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.

Table A-2 (Cont.) Description of the Parameters Used in a Properties File for Scaling Up or Scaling Out a WebLogic Server

Parameter	Description
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.
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_DE ST_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_N M	If destination managed server is associated with a machine, then set to True. Else, set to False.
DEST_MANAGED_SERVERS.0.NM_LISTEN_ADDRE SS	Listen address of node manager of the machine that is associated with the destination managed server.
DEST_MANAGED_SERVERS.0.CLUSTER_NAME_D EST_MANAGED_SERVERS	If destination managed server is a part of the cluster, you specify the cluster name.
DEST_MANAGED_SERVERS.0.DOMAIN_HOME_DE ST_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_D EST_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_S ERVER_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 Enterprise Manager. This use case essentially covers, creating the UDDP using the Enterprise Manager 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 <code>abc.example.com</code> is created using the Enterprise Manager 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 **UDDPTest** 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.
- Once you have provided all the details, click Submit. Enter the a unique Submission name FirstUDDP.
- 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:

 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=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
```

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:defaultService: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:

- 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.
- 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:

38A69D2

```
Java EE Provisioning, Java EE Application Provisioning
Entities, oracle:defaultService:em:provisioning:1:cat:C771B5AAF4A4EED9E040E50AD38A
6E98
MultiOMS, List of Oracle shipped
Directives, oracle:defaultService:em:provisioning:1:cat:C771B5AAF1ACEED9E040E50AD3
8A6E98
myFolder, myFolder
description, oracle:defaultService:em:provisioning:1:cat:C771B5A38A484CE3E040E50AD
```



```
OSBProvisioning, OSBProvisioning
Entities, oracle:defaultService:em:provisioning:1:cat:C771B5AAF3F1EED9E040E50AD38A
6E98
```

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:C771B5A38A494CE3E040E50AD
38A69D2

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_ConfigTmpl
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_WLSTemlpate
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:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_Generic:C77200CA9DC1E7AAE
040E50AD38A1599:0.1'

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:C771B5A38A494CE3E040E50A
D38A69D2'
```

Output:

 $\label{eq:myEntity} \mbox{myEntity description, Ready, Component, Generic Component, Untested, SYSMAN}$



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_Generic:C77
200CA9DC1E7AAE040E50AD38A1599:0.1'
-file="/scratch/user1/patch13653908.zip;newfile1.zip"
-host="host.us.example.com"
-credential_name=mycred11
-credential_owner=sysman
```



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
'oracle:defaultService:em:provisioning:1:cmp:COMP_Component:SUB_Generic:C77200CA9DC1E7AAE
040E50AD38A1599:0.2'.
```

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_Generic:C77
200CA9DC1E7AAE040E50AD38A1599:0.1'
-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_Generic:C77200CA9DC1E7AAE 040E50AD38A1599:0.2'.

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 migratedfrom 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=OmsShared
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=OmsAgent
```

migrate to loc=secondLoc-migrate to type=OmsShared

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=OmsAgent
```

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.



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:

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>
```



```
<mqmt bcn txn with props>
            <mgmt bcn txn is representative="true" name="ats91 with databanks"</pre>
monitoring="true" txn type="OATS"/>
            cproperties>
               cproperty name="Collection Interval" num value="5.0" prop type="2"
encrypt="false"/>
               cproperty name="scriptDescription" string value="[1] Oracle Secure
Enterprise Search
[2] Oracle Secure Enterprise Search
[3] No Title
[4]
No Title" prop type="1" encrypt="false"/>
               cproperty name="fileUploadTime" string value="2010-09-10 13:49:46.0"
prop_type="1" encrypt="false"/>
               cproperty name="OpenScriptJwgName"
string_value="ats91_with_databanks.zip" prop_type="1" encrypt="false"/>
               cproperty name="thinkTimeMaxSec" num value="5.0" prop type="2"
encrypt="false"/>
               property name="usageOptions" string value="userDefined"
prop type="1" encrypt="false"/>
               cproperty name="thinkTimeOption" num value="1.0" prop type="2"
encrypt="false"/>
               cproperty name="fileSize" string value="93699" prop type="1"
encrypt="false"/>
               property name="beaconDistributionOverride"
string value="people=1; middle=1" prop type="1" encrypt="false"/>
               cproperty name="FilePropertyValue" prop type="7" encrypt="false"/>
               property name="databankFilesJar" prop type="7" encrypt="false"/>
               cproperty 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"</pre>
encrypt="false"/>
               cproperty name="thinkTimeMinSec" num value="0.0" prop type="2"
encrypt="false"/>
               cproperty name="oatsVersion" string value="9.1.0" prop type="1"
encrypt="false"/>
               cproperty name="databankValues"
string value="people.firstname=yang.,people.lastName=wang,middle.middlename col=x."
prop type="1" encrypt="false"/>
               property name="modules"
string value="oracle.oats.scripting.modules.utilities; version=2.4.0% #xA; oracle.oats.s
cripting.modules.http;version=2.4.0
oracle.oats.scripting.modules.basic;version=2
.4.0" prop type="1" encrypt="false"/>
               cproperty name="databankAliasMapping"
string value="middle=middle.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='databank: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="ATSTest1" -
testtype="OATS"
```

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



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 Lifecycle 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 before Oracle 12c can not be exported using the Partool export utility.



For importing PAR files that contain Software Library entities, ensure that your Software Library is configured. For information on configuring software libraries, see Configuring a Software Library in *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 circumstances, 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 <code>\$ORACLE_HOME/bin/emctl</code> 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 parcedure guid> -file <file> -displayName <name> -
description <desc> -metadataOnly(optional)
emctl partool check
emctl partool help
```

Table A-3 describes the additional options that can be used with the emctl partool utility.

Table A-3 emctl partool Options

Option	Description
-repPasswd <reppassword></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></file>	Represents the PAR file.
-action <deploy view export></deploy view export>	Deploys, views, or exports PAR files.
-verbose	Indicates verbose mode.
help	Displays Help information.
-displayName <displayname></displayname>	Indicates PAR file name.
-parDir <dir></dir>	Directory where the PAR files are located.
-metadataOnly	Flag for metadata-only exports.
-guid <guid></guid>	Procedure GUID to export.
	To export multiple procedures, provide the GUIDs separated by ","
-parFile <file></file>	Path to the PAR file.
-description <description></description>	PAR file description.



Table A-3 (Cont.) emctl partool Options

Option	Description
-ssPasswd <secretstorepassword></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.</deploy view>

Checking Oracle Software Library

Before running the emctl partool utility to export or import PAR files, ensure that the <code>\$ORACLE_HOME</code> 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:

- In Cloud Control, from the Enterprise menu select Provisioning and Patching, and then click Procedure Library.
- On the Provisioning page, right click the deployment procedure name and from the menu select Copy Link Location.
- 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
```



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,09AEFF331025AA
D0EE40578C85FB5772" -file $ENV{T_WORK}/tvmgf_partool_multi_dp.par -displayName
"partool multi dp test" -description "partool multi dp test description" -
repPasswd sysman
```



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 Enterprise Manager, that is, the console. This section explains the following:

Importing Using Command Line Interface



Importing Using the Enterprise Manager Console



Importing an existing PAR file from previous releases (IE from Enterprise Manager 13.5 to Enterprise Manager 24.1) is not supported.

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 setup, 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 the Enterprise Manager Console

To import PAR files or deploy them to an OMS, you can use the emctl partool utility. Alternatively, you can import then from Enteprise Manager.

For importing the PAR files, follow these steps:

- 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.**
- 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.



When importing or exporting components and/or directives that contain properties with secret values, you must use the <code>-ssPasswd</code> 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 A-3.

Troubleshooting Issues

This appendix provides solutions to common issues you might encounter:

- Troubleshooting Database Provisioning Issues
- · Troubleshooting Patching Issues
- Refreshing Configurations
- Reviewing Log Files

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
- User Account and Role Issues
- Workaround for Errors



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

Issue

While performing a patch it fails with an unexpected error, although the credentials are correct and you have write permission on the EM stage location.

"Unexpected error occurred while checking the Normal Oracle Home Credentials"

Cause

The Software Library was setup using the OMS Agent file system. You do not have access to the named credentials that were used for setting up the Software Library.

Solution

Grant yourself access to the named credentials that were used to set up the Software Library.

Error Occurs While Uploading a Patch Set

Issue

When uploading a patch set using the **Upload Patches** to the Software Library page you see an error stating that the operation failed to read the patch attributes.

ERROR: Failed to read patch attributes from the uploaded patch file <filename>.zip

Cause

Occasionaly 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

Manually enter all the patch attributes under 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

Issue

OPatch Update job fails with the following error:

2021-11-28 10:31:19,127 RemoteJobWorker 20236 ERROR em.jobs startDownload.772-OpatchUpdateLatest: java.lang.NullPointerException: Category, 'Oracle Software Updates', has no child named, 'OPatch' at oracle.sysman.emInternalSDK.core.patch.util. ComponentUtil.getComponentCategory (ComponentUtil.java:854)

After applying the January 2012 Enterprise Manager patch you 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. When you run two patch upload or download jobs (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 Software Library does not display an error while creating these duplicate directories. But when you run the OPatch Update job it fails with a NullPointerException.

Solution

If you see two *Patch Components* directories in the Software Library, delete the one with fewer entries and retry the failed patch 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 see the error Oracle Software Updates already exists. 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

Issue

On the **Proxy Settings** page under **My Oracle Support and Proxy Connection tab**, when you provide the manual proxy configuration details, you get an exception error stating:

```
Error: -Proxy Test resulted in Exception. Error while reading data from URL https://updates.oracle.com.
Cannot establish proxy connection: 407 Proxy Authentication Required
```

Cause

You provided the configuration details of a proxy server that supports only a *Digest* authentication schema. By default, the proxy server is mapped only to the *Basic* authentication schema. Currently there is no support for *Digest* authentication schema.

Solution

Reconfigure your proxy server to the Basic authentication schema.

For a 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/nodemanager/
startup.properties
```



Append the following string to the value of the property Arguments:

```
-DHTTPClient.log.level\=ALL -DHTTPClient.log.verbose\=true
```

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
```

You can also use **grep**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

While entering credentials in Fleet Maintenance you set privileged credentials as normal Oracle home credentials you get an error stating that log files cannot be created in the EMStagedPatches directory.

```
"Unable to create the file <RAC_HOME>/EMStagedPatches/PA_APPLY_PATCH_09_02_2011_14_27_13.log"
```

or:

```
ERROR: SharedDeviceException.
ACTION: Please check whether the configuration is supported or not.
```

Cause

When a Fleet Maintenance activity is executed it uses deployment procedures to orchestrate the patch deployment. Some steps in the deployment procedure run with normal Oracle home credentials, others with privileged Oracle home credentials. When normal Oracle home credentials are set as privileged Oracle home credentials, the deployment procedure runs those steps as a root user instead of the Oracle home owner, and encounters an error.

Solution

Return to the Fleet Maintenance page, in the Credentials tab set normal credentials as normal Oracle home credentials.

User Account and Role Issues

This section describes the following:

 Fleet Maintenance Patching Errors If Patch Designers and Patch Operators Do Not Have the Required Privileges

Fleet Maintenance Patching Errors If Patch Designers and Patch Operators Do Not Have the Required Privileges

Issue

When patching with Fleet Maintenance, the patch fails 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/

The error occurs because the user does not have the following roles:

- 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:

- 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

You can also grant the provisioning roles that include the previous roles within:

- EM_PROVISIONING_OPERATOR
- EM_PROVISIONING_DESIGNER

After granting the privileges, retry the out-of-place patching operation.

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

This section describes the possible missing property errors and the workarounds you can use to resolve them.



Table A-4 Missing Properties Error and Workaround

Problem Workaround

Empty target property version The target is not properly configured or is unavailable. Reconfigure the target or check for metric collection errors.

Inadequate or incomplete target information collected by Oracle Management Agent.

Recompute the dynamic properties and refresh the host configuration so that the Management Repository is updated with the latest configuration of the host.

- To recompute the dynamic properties:
 - Stop and restart the Management Agent (easiest method, incurrs in downtime).

emctl stop agent emctl start agent



For clusters restart the Management Agent on all nodes.

- Reload the dynamic properties of all the targets (preferred method avoids blackouts or downtime).
 - View the list of targets that are monitored by the Management Agent:

emctl config agent listtargets

ii. Reload the dynamic properties of a specific target:

emctl reload agent dynamic properties [<Target name>:<Target Type>]

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.

Rediscover the domain so all the targets are re-discovered effectively.

- Log in to the Domain Home page using appropriate credentials. For example, Farm01 base domain.
- On the domain home page go to the Farm menu, select Refresh WebLogic Domain, and click Ok on all following pages to complete the process. After successful completion of the process the domain home page is refreshed with the targets discovered.

Workaround for Unsupported Configuration Errors

Table A-5 Workaround for Unsupported Configuration Errors

Problem	Workarounds
The target is not mediated by the OMS	The target discovery is not appropriate. Remove the target from Enterprise Manger, and rediscover on all the Management Agents in the cluster.

Oracle Fleet Patching and Provisioning: Diagnostic and Troubleshooting

This section describes troubleshooting for FPP server and client issues within Enterprise Manager.

General diagnostics and log collection for FPP server and client

Check the status of the FPP server and client

- GI HOME/bin/srvctl status rhpserver
 - Try starting the FPP server if the status is stopped
- GI HOME/bin/srvctl start rhpserver
 - On failure:
 - Check MGMTDB is running
 - * GI HOME/bin/srvctl status mgmtdb
 - * Check MGMTDB Listener is running
 - * GI HOME/bin/srvctl status mgmtlsnr
 - Check RHP disk group is online
 - * srvctl status diskgroup -diskgroup <DGNAME>

Check status of FPP Client

- GI HOME/bin/srvctl status rhpclient
 - Try starting the FPP client if the status is stopped

FPP Server Side diagnostics

- tfactl diagcollect -rhp
- tfactl print repository For TFA log location

FPP Server Logs location

• GI BASE/crsdata/<host>/rhp/rhpserver.log

FPP Client side Diagnostics

- tfactl diagcollect
- tfactl print repository: For TFA log location

FPP Client Logs Location

• GI BASE/crsdata/<host>/rhp/rhpclient.log



Common Issues with FPP

Issue	Reason	Diagnostics	Workaround
Connection Errors in Deployment Procedure(DP) logs	Connection issues between EM and the FPP server or between FPP server and Client	Check if the server and client are up Check if the server REST API end point is configured correctly in EM by checking server properties Check for connectivity between server and client	If the server and client are up, run a on-demand sync in EM to make sure REST API end point is updated automatically in EM. <link api="" for="" job="" on-demand="" pointing="" rest="" sync="" to=""/> MISSING You can update the REST end point manually as well <link api="" for="" rest="" server<="" td="" to="" updating=""/>
DP step fails during sync job	Could be a temporary error due to concurrent calls to FPP server	N/A	properties> MISSING Retry the failed step from EM
DP steps fails during a FPP operation	This can happen due to various reasons	Check DP logs for the reason and also check server/client logs Follow the common diagnostics steps mentioned above.	If the job needs to be retried do the following 1. If the DP step that invokes the operation failed, then a simple retry of the step should re-trigger the call to the FPP server 2. If the DP step that invokes the operation went through and the next step that polls the FPP for status of the job failed, you can either retry the failed step to resume polling or if you want to retry the failed operation itself, then click on Update and Retry button, and for the parameter called RETRY_OPERATION, specify value as true. This will retry
Failures in the FPP calls from EM	The failures could be due to various reasons. If the above have been looked at, and nothing from logs seems to help, you might need more support from the Oracle support	Follow the common diagnostics steps mentioned above.	the operation itself. N/A



Troubleshooting Database Provisioning Issues

This section provides troubleshooting tips for common database provisioning issues.

Grid Infrastructure Root Script Failure

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
```

Retry the failed step from the Procedure Activity page.

SUDO Error During Deployment Procedure Execution

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 the Sudoers file contains the Default env_reset parameter, add the following entries after it:

Defaults env keep="JRE HOME PERL5LIB EMDROOT"

Prerequisites Checks Failure

Issue

Prerequisites checks fail when submitting a deployment procedure

Description



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

Issue

Oracle ASM disk creation fails.

Description

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

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**.



- In the Database Procedures page, select the Provision Oracle Database Deployment Procedure and click Create Like.
- 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

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 diagnosable 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

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 cluvy 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**.



Click Retry. The failed step will be retried.

Refreshing Configurations

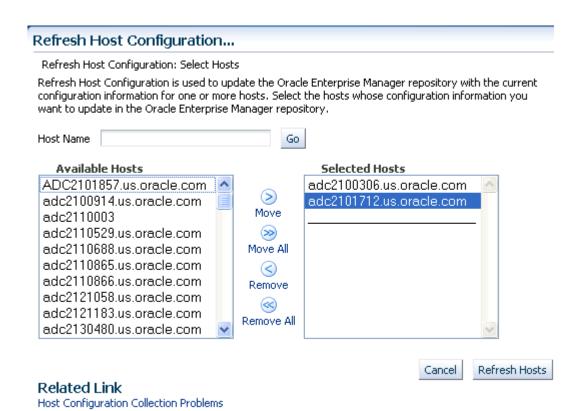
If you encounter issues and are expected to refresh the configurations in the host or the Oracle home, then follow the instructions outlined in the following sections:

- 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 Enterprise Manager, 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.



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.

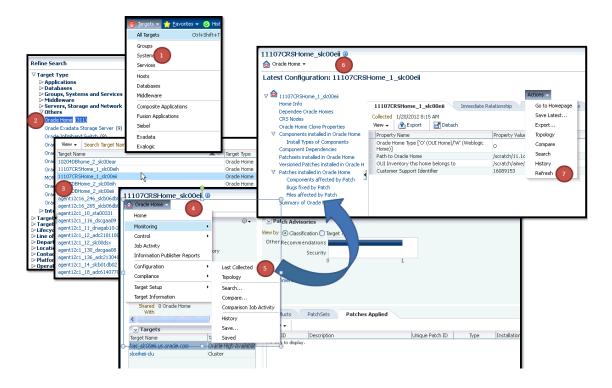
Note:

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.
- Click the Target name to select it.
- 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.



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	change plan, 25-15, 25-19
according	creating change plan, 25-15
accessing	submitting change plan, 25-19
compliance standards, 24-15	change plans
compliance violations, 24-47	overview, 25-14
accessing Software Library console, 2-2	clearing violations, 24-75
adding new target data collections, 23-57	clone
administering pluggable databases, 13-50	cloning a running Oracle database replay
agent-side rules	client, 11-2
compliance standard rules, 24-28	cloning a running Oracle RAC instance, 8-2
definition, 24-2	cloning database, 14-7
AIX Installed Packages parser, 23-77	RMAN backup, 14-7
analyzing configuration health, 23-61	cloning pluggable databases, 13-19
AND/OR logical operators	cluster verification utility, 3-38
in comparison rules, 23-22	collected configurations for targets, 23-1
Apache HTTPD parser, 23-77	columnar
Application Data Model, 30-3	parser parameters, 23-76
application signature, 30-2	parsers, 23-75
Oracle E-Business Suite, 30-2	comparing data, 25-22
Oracle Fusion Applications, 30-2	comparison
Oracle PeopleSoft Enterprise, 30-2	rules, <i>1-3</i>
associating compliance standard to a target,	comparison results, 23-23, 23-24
24-73	standard target, 23-24
associations. See relationships, 23-89	synchronizing configuration files, 23-24
auditing, how to set up compliance standards for,	comparison template
24-15	creating or editing, 23-9
automatic metadata discovery, 30-4	deleting, 23-12
Autosys parser, 23-77	exporting, 23-12
	importing, 23-13
В	managing, 23-12, 23-52
	Template Settings tab, 23-9
base parsers	viewing, 23-12
columnar, 23-75	comparison templates, 23-3, 23-9–23-11
format-specific, 23-71	creating templates, 23-9
properties, 23-77	editing templates, 23-9
BEA Tuxedo parser. See Ubb Config parser,	property settings, 23-10
23-71	rules for matching, 23-11
Blue Martini DNA parser, 23-71	template settings, 23-10
browsing	comparison wizard, 23-13
compliance frameworks, 24-13	comparisons, 23-3
compliance standard rules, 24-41	compliance
compliance standards, 24-23	accessing, 24-4
compliance standards, 27 20	compliance frameworks, 24-7
	compliance standard rule folders, 24-27
C	compliance standard rules, 24-28
change management 25.1	compliance standards, 24-15
change management, 25-1	33p



compliance <i>(continued)</i>	compliance standard rule folders (continued)
configuring, 24-6	creating, <i>24-27</i>
dashboard, 24-45	definition, 24-3
evaluating, 24-43	managing in compliance standard, 24-27
evaluation	compliance standard rules
results, 24-46, 24-47	about, 24-28
importance, 24-55	agent-side rules, 24-28
investigating evaluation errors, 24-54	definition, 24-2
managing, 24-6	operations on
overview, 24-1	browsing, 24-41
reports, 24-54	creating agent-side rules, 24-32, 24-66
score, 24-55	creating like, 24-39
statistics, accessing, 24-44	creating manual rules, 24-33, 24-66
summary information, 24-46	creating repository rules, 24-30
terms used in, 24-2	deleting, 24-40
violations	editing, <i>24-39</i>
accessing, 24-47	exporting, 24-40
examples of viewing, 24-50	importing, 24-41
managing, 24-47	searching, 24-41
manual rules, 24-47	repository rules, 24-29
of a target, 24-50	types, 24-28
suppressing, 24-47	compliance standards
unsuppressing, 24-47	about, <i>24-15</i>
compliance frameworks	accessing, <i>24-15</i>
about, <i>24-7</i>	adding to another compliance standard, 24-18
accessing, 24-7	adding to compliance framework, <i>24-9</i>
adding compliance standard to, 24-9	advantages of using, 24-15
advantages of using, 24-7	associating with targets, 24-24
compliance of database targets, 24-14	definition, 24-2
compliance score, 24-57	errors, 24-23, 24-24
definition, 24-2	evaluation results, 24-23
editing importance, 24-9	investigating violations, 24-49
errors, 24-14	operations on, 24-17
evaluation results, <i>24-13</i>	browsing, 24-23
lifecycle status, 24-9	creating, 24-18
operations on, 24-8	creating like, 24-21
browsing, 24-13	customizing, 24-21
creating, 24-19	deleting, 24-21
creating like, 24-10	editing, 24-21
deleting, <i>24-11</i>	exporting, 24-22
editing, 24-11	importing, 24-22
exporting, 24-12	·
•	searching, 24-23
importing, 24-12 searching, 24-13	security metrics, enabling, 24-26 setting up for auditing use, 24-15
provided by Oracle, 24-7	violations details, 24-50
· · · · · · · · · · · · · · · · · · ·	
reasons for using, 24-7	compliance violations, investigating, 24-49
user-defined, 24-7	configuration browser, 23-38, 23-39, 23-41, 23-42
violations details, 24-50	inventory and usage details, 23-42
compliance score	saved configurations, 23-41
compliance framework, 24-57	viewing configuration data, 23-39
compliance standard, 24-56	Configuration Browser, viewing configurations,
compliance standard rule, 24-56	23-38
definition, 24-3	configuration changes, 23-43
parent node, 24-57	configuration collections, 23-48
compliance standard rule folders	configuration comparison
about, <i>24-27</i>	ignore rule, 23-20

configuration comparison (continued)	configuration searches, 23-34, 23-36
ignore rule example, 23-22	managing configuration searches, 23-34
matching rule, 23-20	setting up, 23-36
matching rule example, 23-21	Configuration Topology Viewer, 23-58
rule examples, 23-21	configuration topology, viewing, 23-60
rule language, 23-18	configurations
rules, 23-17	hardware and software, 23-1
value constraint rule, 23-19	history, 1-3
wizard, 23-13	searching, 1-3, 23-34
configuration data collections, 23-56	viewing, 23-38
configuration data, extending collections, <i>23-56</i>	configure Grid Infrastructure, 6-7, 6-20
configuration extension	Connect((colon))Direct parser, 23-71
creating, 23-49	container databases (CDBs)
credentials, 23-51	prerequisites for creating, 4-4, 4-9, 4-13
database roles, 23-51	procedure, 4-5, 4-10, 4-14
deleting, 23-53	controlling appearance of information on a graph
deploying, 23-55	23-65
editing, 23-49	creating
editing deployment, 23-56	agent-side compliance standard rules, 24-32
encoding, 23-50	24-66
exporting, 23-53	comparison template, 23-9
Files ((amp)) Commands tab, 23-50	compliance frameworks, 24-9
importing, 23-53	compliance manual rules, 24-66
post-parsing rule, 23-85, 23-87, 23-89	compliance standard example, 24-72
privileges, 23-54	compliance standard rule folders, 24-27
roles, 23-54	compliance standards, 24-18
rules, 23-52	compliance standards, considerations, 24-26
sample non-XML parsed file, 23-86	configuration extension, 24-66
sample parsed SQL query, 23-88	custom target type, 23-48
sample XML parsed file, 23-85	manual compliance standard rules, 24-33
undeploying, 23-55	new relationships, 23-89
versioning, 23-53	repository compliance standard rules, 24-30
viewing collection data, 23-56	creating change plans using external clients,
viewing specification details, 23-53	25-18
XML parsed example (default), 23-69	creating database provisioning entities, 3-34
XML parsed example (generic), 23-70	creating database templates, 3-14
XML parsed example (modified), 23-70	creating databases, 4-4, 4-5, 6-22
XPath, 23-52	prerequisites, 4-4
configuration extension, creating, 24-66	procedure, 4-5
configuration extensions, 23-48–23-54	creating installation media, 3-13
commands tab, 23-50	creating like
creating, 23-49	compliance frameworks, 24-10
deployment, 23-54	compliance standard rules, 24-39
privileges, 23-54	compliance standards, 24-21
setting up credentials, 23-51	creating Oracle Database, 4-4
setting up rules, 23-52	prerequisites, 4-4
SQL tab, 23-51	creating Oracle RAC One Node Database, 4-12
versioning, 23-53	prerequisites, 4-13
	·
configuration history, 23-43, 23-44, 23-46	procedure, 4-14
accessing, 23-44	creating Oracle Real Application Clusters
annotating, 23-46	Database, 4-9
creating notification list, 23-46	prerequisites, 4-9
scheduling, 23-46	procedure, 4-10
searching history, 23-44	creating relationships to a target, 23-64
configuration history, job activity, 23-47	
configuration management, 23-1	

creating repository rule based on custom	database provisioning (continued)
configuration collections, 24-63	Oracle RAC database with file system on new
examples, 24-63	cluster, 6-18
creating users	Oracle Real Application Clusters One Node
designers, 2-7	databases, 7-1
operators, 2-7	prerequisites for designers, 3-7
credentials, for configuration extensions, 23-51	prerequisites for operators, 3-9
Cron Access parser, 23-75	provisioning and creating Oracle Databases,
Cron Directory parser, 23-75	4-17
CSV parser, 23-75	provisioning Oracle RAC, 8-1
Custom CFG parser, 23-77	setup, 3-5
custom configuration	supported targets, 3-2
custom target type, 23-48	usecases, 3-4
customizing topology views, 23-63	database provisioning overview, 3-1
CVU, 3-38	Database Query
	parser, 23-71
D	parser parameters, 23-72
D	Database Query Paired Colulmn
data comparison, 25-2, 25-20	parser, 23-71
overview, 25-19	parser parameters, 23-72
requirements, 25-20	database templates, 3-15
data discovery, 30-1	uploading to software library, 3-15
data discovery job, 30-5	Db2 parser, 23-71
data discovery results, 30-5	delete Oracle RAC, 10-2
Data Guard Rolling Upgrade, 19-4	prerequisites, 10-2
Data Redaction, 30-2	delete Oracle RAC nodes, 10-5
database credentials, 23-51	deleting
database host readiness, A-1	comparison template, 23-12
configuring SSH, A-2	compliance frameworks, 24-11
environment settings, <i>A-2</i>	compliance standard rules, 24-40
kernel requirements, A-3	compliance standards, 24-21
memory requirements, A-4	configuration extension, 23-53
network and IP requirements, A-5	custom topology views, 23-63
node time requirements, A-4	relationships from a target, 23-65
package requirements, A-4	deleting Oracle RAC, 10-2
installation directories and Oracle inventory,	deleting core components, 10-2
A-6	deleting entire Oracle RAC, 10-2
PDP setup, A-2	deleting pluggable databases, 13-37
setting user accounts, A-1	dependency analysis, 23-62
shell limits, A-2	deploying configuration extension, 23-55
storage requirements, A-6	deployment procedures, 7-2
database provisioning	editing the permissions, 15-17
administrator privileges, 3-6	phases and steps, 15-8
. •	target list, 15-6
deployment procedures, 3-2	tracking the status, 15-18
getting started, 4-16 Oracle Detabase software, 4-38	User, roles and privileges, 15-3
Oracle Database software, 4-28	variables, 15-7
Oracle Original Infrastructure and Oracle	viewing, editing, and deleting, 15-17
Oracle Grid Infrastructure and Oracle	deployments
Database software, 5-9	Management Repository, 23-1
Oracle Grid Infrastructure and Oracle	Designer and Operator Roles, 3-1
Databases with Oracle ASM, 5-2	determining configuration health compliance
Oracle Grid Infrastructure for Oracle Real	score, 23-61
Application Clusters Databases, 6-1	
Oracle RAC database with file system on	diagnosing compliance violations, 24-49
existing cluster, 6-12	compliance violations, <u>24-49</u>

Directory	emctl partool utility (continued)
parser, 23-71	exporting deployment procedure
parser parameters, 23-73	creating PAR file, A-41
	retrieving GUID, A-40
E	exporting deployment procedures, A-40
	importing PAR files, A-41
E-Business Suite	using cloud control, A-43
parser, 23-71	using command line, A-42
parser parameters, 23-73	overview, A-38
editing	software library, A-40
comparison template, 23-9	Enterprise Data Governance dashboard, 30-3
compliance frameworks, 24-10	enterprise manager users
compliance standard rules, 24-39	designers, 2-6
compliance standards, 24-21	operators, 2-6
configuration extension deployment, 23-56	super administrators, 2-6
EM ALL OPERATOR	errors
EM_ALL_VIEWER, 15-5	in compliance standards, 24-23, 24-24
EM_COMPLIANCE_DESIGNER role, 23-54	evaluating compliance, 24-43
EM_PLUGIN_AGENT_ADMIN role, 23-54	evaluation errors
EM PLUGIN OMS ADMIN role, 23-54	compliance, 24-54
EMCLI, 17-15, A-7	evaluation results
adding ATS service test, A-36	compliance, 24-46, 24-47
advantages, A-7	compliance standards, 24-23
creating a new generic component by	examples
associating a zip file, A-30	associating targets, 24-63
creating properties file, A-8	creating compliance standards, 24-63
launching a procedure with an existing saved	creating custom configuration, 24-63
procedure, A-14	creating custom-based repository rule based
migrate and remove a software library storage	on custom configuration collection,
location, A-35	24-63
overview, A-7	viewing compliance results, 24-63
provisioning Oacle WebLogic Server	viewing compliance violations, 24-50
using provisioning profile, A-22	excluding relationships from custom topology
provisioning Oracle Database software, <i>A-20</i>	views, 23-64
provisioning Oracle WebLogic Server, A-22	exporting
scaling up or scaling out, A-24	comparison template, 23-12
provisioning pluggable databases, <i>A-15</i>	compliance frameworks, 24-12
creating new pluggable databases, <i>A-15</i>	compliance standard rules, 24-40
migrating databases as pluggable	compliance standards, 24-22
databases, A-18	configuration extension, 23-53
using snapshot profiles, A-17	extend Oracle RAC, 9-1
provisioning user defined deployment	prerequisites, 9-1
procedure, A-29	extending Oracle RAC, 9-1, 9-2
adding steps/phases, A-29	prerequisites, 9-1
prerequisites, A-29	procedure, 9-2
running the procedure, A-30	'
provisioning verbs, A-7	Г
unplugging pluggable databases, <i>A-19</i>	F
using an existing properties file, A-11	file synchronization, 23-24
EMCLI verbs	finding See accessing, 24-47
create_pluggable_database, A-15, A-17	format-specific parsers, 23-71
migrate_noncdb_to_pdb, A-19	ioimat-specific parsers, 25-71
unplug_pluggable_database, <i>A-20</i>	
upload patches, 17-15	G
emctl partool utility, A-38	Colore CEC
emctl partool options, A-38	Galaxy CFG
emen partour options, A-30	parser, <i>23-71</i>

Galaxy CFG (continued) parser parameters, 23-73 generic system, and new relationships, 23-89 GNS settings, 6-21 gold image provisioning Oracle database replay client using gold image, 11-5 provisioning Oracle RAC using gold image, 8-9	lifecycle management (continued) solution areas change management, 1-3 compliance management, 1-4 configuration management, 1-3 discovery, 1-2 patching, 1-3 provisioning, 1-3 lifecycle status compliance frameworks, 24-9 Linux Directory List parser, 23-75 lock down, 15-22
hardware configuration, collecting information, 23-1 history job activity, 23-47	locking down feature, 3-1 logical operators AND/OR, 23-22
history, of configuration changes, 23-43 history, of configurations, 1-3 Hosts Access parser, 23-75	М
Tiosis Access parser, 23-73	Management Repository, 23-1 managing compliance standard rule folders, 24-27
ignore rule example, in comparisons, 23-22 ignore rule, in comparisons, 23-20 impact analysis, 23-62 importance definition, 24-3	mandatory infrastructure requirements creating user accounts, 2-5 manual rules definition, 24-2 violations, 24-47 matching rule, 23-21
editing in compliance standard, 24-9 importing comparison template, 23-13 compliance frameworks, 24-12 compliance standard rules, 24-41 compliance standards, 24-22 configuration extension, 23-53	examples, 23-21 matching rule example, for comparisons, 23-21 matching rule, in comparisons, 23-20 metadata discovery, 30-1 metadata discovery job, 30-3 metadata discovery results, 30-4
including relationships in custom topology views, 23-64	metadata XML files downloading files, 17-12
infrastructure requirements, 2-1 Introscope parser, 23-71 inventory and usage details, 23-42	migrating databases as pluggable databases, 13-31 Mime Types parser, 23-77 monitoring facets definition, 24-3
Java Baliau paraar 22 77	monitoring rules definition, 24-3
Java Policy parser, 23-77 Java Properties parser, 23-77 job activity history, 23-47	MQ-Series parser, 23-71 parser parameters, 23-74
K	<u>O</u>
Kernel Modules parser, 23-75	observation audit status, definition, 24-3 observation bundles, definition, 24-3 observations definition, 24-3
LDAP parser, 23-77	Odin parser, 23-71
lifecycle management overview, 1-2	optional infratructure requirements host configurations, A-54 self update for provisioning, 2-8



Oracle clusterware clone, 3-36	parsers (continued)
Oracle Clusterware Clone, 3-37	Tuxedo, 23-77
Oracle Database Clone, 3-35	UbbConfig, 23-71
Oracle Database topology, 4-1	Unix Config, 23-77
Oracle Label Security, 30-2	Unix Crontab, 23-75
Oracle ORA parser, 23-71	Unix Directory List, 23-75
Oracle RAC database topology, 6-3	Unix Groups, 23-75
Oracle Real Application Clusters Database	Unix GShadow, 23-75
·	Unix Hosts, 23-75
topology, 6-3	, , , , , , , , , , , , , , , , , , ,
	Unix INETD, 23-75
P	Unix Installed Patches, 23-72
	Unix Login, 23-77
PAM Configuration parser, 23-75	Unix Passwd, 23-75
parsers	Unix PROFTPD, 23-77
AIX Installed Packages, 23-77	Unix Protocols, 23-76
Apache HTTPD, 23-77	Unix Recursive Directory List, 23-72
Autosys, 23-77	Unix Resolve, 23-78
Blue Martini DNA, 23-71	Unix Services, 23-76
columnar, 23-75	Unix Shadow, 23-76
Connect((colon))Direct, 23-71	Unix SSH Config, 23-78
Cron Access, 23-75	Unix System, 23-78
Cron Directory, 23-75	Unix System Crontab, 23-76
CSV, 23-75	Unix VSFTPD, 23-78
Custom CFG, 23-77	Unix XINETD, 23-78
Database Query, 23-71	WebAgent, 23-78
Database Query Paired Column, 23-71	WebLogic (attribute-keyed), 23-68
Db2, 23-71	WebSphere (attribute-keyed), 23-68
Directory, 23-71	WebSphere (generic), 23-69
E-Business Suite, 23-71	Windows Checksum, 23-78
format-specific, 23-71	XML (generic), 23-69
Galaxy CFG, 23-71	XML default (attribute-keyed), 23-68
Hosts Access, 23-75	patch management solution
Introscope, 23-71	downloading catalog file, 17-12
Java Policy, 23-77	overview, 16-1
Java Properties, 23-77	patching modes
Kernel Modules, 23-75	online mode, <i>17-11</i>
	patching workflow, 16-1
LDAP, 23-77	registering proxy details, 17-11
Linux Directory List, 23-75	searching for patches, 17-7
Mime Types, 23-77	setting up infrastructure, 17-5
MQ-Series, 23-71	uploading patches to Software Library, 17-13
Odin, 23-71	patching
Oracle ORA, 23-71	identifying applicable patches
PAM Configuration, 23-75	searching in software library, 17-8
Process Local, 23-75	searching on MOS, 17-7
properties, 23-77	
Radia, 23-77	phases
Sectioned Properties, 23-77	parallel, 15-8
Secure TTY, 23-75	rolling, 15-8
Siebel, 23-71	pluggable database administration
SiteMinder Agent, 23-77	altering pluggable database state, 13-51
SiteMinder Registry, 23-77	opening/closing pluggable databases, 13-51
SiteMinder Report, 23-77	switching between pluggable databases,
SmWalker, 23-77	13-51
Solaris Installed Packages, 23-75	pluggable database jobs
Sun ONE Magnus, 23-77	create, 13-10, 13-19, 13-47
Sun ONE Obj., 23-77	creates, <i>13-28</i>

pluggable database jobs (continued)	provision Oracle RAC databases, 6-18
delete, 13-46, 13-49	provisioning
unplug, 13-42, 13-48	deleting or scaling down Oracle RAC (Real
pluggable database management requirements,	Application Cluster), 10-1
13-3	extending Oracle RAC (Real Application
pluggable database provisioning	Cluster), 9-1
cloning pluggable databases	provisioning Oracle database replay client,
using Full Clone, 13-19	11-1
creating new pluggable databases, 13-4	provisioning database client
migrating databases as pluggable databases,	getting started, 11-1
13-31	provisioning Oracle Database client, 11-8
plugging in pluggable databases, 13-10	provisioning Oracle RAC, 8-9, 8-16, 8-23
pluggable database removal	archived software binaries, 8-16
deleting pluggable databases, 13-42	gold image, 8-9
unplugging pluggable databases, 13-37	no root credentials, 8-23
pluggable databases	provisioning Oracle Real Application Clusters One
administering, 13-50	database, 7-2
getting started, 13-1	provisioning Oracle standby database, 12-1, 12-5,
jobs, 13-47	12-8, 12-9
overview, 13-2	creating logical standby database, 12-5
provisioning, 13-3	creating physical standby database, 12-1
removing, 13-37	creating primary database backup, 12-9
post-parsing rules, 23-84	managing existing standby database, 12-8
privileges, for configuration extensions, 23-54	provisioning Oracle standby databases, 12-1
Process Local parser, 23-75	provisioning pluggable databases, 13-3
properties parser constructs	provisioning profiles, 3-1
delimited section, 23-83	
delimited structure, 23-82	D
element cell, 23-84	R
explicit property, 23-81	Radia parser, 23-77
implicit property, 23-81	relationships, 23-89
INI section, 23-83	Relocating Pluggable Databases, 13-35, 13-36
keyword name property, 23-81	removing pluggable databases, 13-37
keyword property, 23-81	reports, compliance, 24-54
reserved directive, 23-82	repository rules
reserved function, 23-82	definition, 24-3
simple property, 23-80	
structure, 23-83	in compliance standards, 24-29
XML structure, 23-82	results
properties parsers, 23-77	compliance evaluation, 24-46, 24-47
advanced constructs, 23-80	reverse transform, 23-26
advanced parameters, 23-79	roles, for configuration extensions, <i>23-54</i>
basic parameters, 23-78	Rolling Upgrade, 19-4
Protection Policy	Rolling Upgrade procedure, 19-5
Data Redaction, 30-2	rollup options, 23-43
Oracle Label Security, 30-2	root cause analysis. See dependency analysis,
	23-62
Transparent Data Encryption, 30-2	rule examples, for comparisons, 23-21
Virtual Private Database, 30-2	rule expressions, in comparisons, 23-18
provision database Grid Infrastructure and Oracle RAC Database,	rules, 23-17, 23-19, 23-20
	include or exclude rule, 23-20
6-4	matching rule, 23-20
Oracle Grid Infrastructure and Oracle Real	value constraint rule, 23-19
Application Clusters, 6-1	rules expression, 23-18
provision Oracle RAC database	rules syntax, 23-18
with file system on a new cluster, 6-18	rules, in comparisons, 1-3, 23-17
with file system on an existing cluster, 6-12	rules, in configuration extensions, 23-52

S consider the section and are in a constant	Solaris Installed Packages parser, 23-75
	specify OS users, 6-18
save as draft, configuration extension, 23-53	specifying rules, 23-17
saved configurations, 23-41	statistics, accessing compliance, 24-44
scale down Oracle RAC, 10-5	steps
schema baseline, 25-1 multiple versions, 25-4	action, 15-8
•	computational, 15-8
schema baseline version, 25-3 schema baselines	file transfer, 15-8
	host command, 15-10
export, <i>25-4</i> import, <i>25-4</i>	job, <i>15-8</i> library component, <i>15-8</i>
overview, 25-2	library directive, 15-9
schema change plans, 25-2	manual, 15-8
schema comparison, 25-1	Sun ONE Magnus parser, 23-77
schema comparison versions, 25-8	Sun ONE Obj parser, 23-77
schema comparisons, 25-6, 25-7	suppressing violations, 24-47, 24-74
comparison options, 25-7	synchronizing files, 23-24
overview, 25-5	system component structure, 23-61
schema map, 25-6	System component structure, 25-01
scope specification, 25-6	_
schema synchronization, 25-2	Т
overview, 25-8	torget type guetom 22 40
versions, 25-11	target type, custom, 23-48
schema synchronization version, <i>25-11</i>	targets associating compliance standard with, 24-24,
schema synchronization cycle, 25-11	24-73
schema synchronizations, 25-8–25-10	TDE
schema map, 25-9	prerequisites, 29-2
scope specification, 25-8, 25-9	topology
synchronization mode, 25-10	Oracle RAC database topology, 6-3
synchronization options, 25-9	topology viewer
scope specification, 25-3	controlling appearance of information on a
search configurations	graph, 23-65
predefined, 1-3, 23-34	creating
user-defined, 1-3, 23-34	relationships to a target, 23-64
searching	customizing views, 23-63
compliance frameworks, 24-13	deleting custom views, 23-63
compliance standard rules, 24-41	deleting relationships from a target, 23-65
compliance standards, 24-23	dependency analysis, 23-62
Sectioned Properties parser, 23-77	excluding relationships from custom views,
Secure TTY parser, 23-75	23-64
security metrics	impact analysis, 23-62
enabling, 24-26	including relationships in custom views, 23-64
in compliance standards, 24-26	tracking configuration changes, 23-43
sensitive column type, 30-4	Transparent Data Encryption, 30-2
sensitive data discovery, 30-3	meeting prerequisites, 29-2
Setting Up MOS, 2-8	troubleshooting missing property errors, A-47
Siebel	troubleshooting unsupported configuration errors,
parser, 23-71	A-48
parser parameters, 23-74	Tuxedo parser, 23-77
SiteMinder Agent parser, 23-77	
SiteMinder Registry parser, 23-77	U
SiteMinder Report parser, 23-77	<u></u>
SmWalker parser, 23-77	UbbConfig parser, 23-71
software library	undeploying configuration extension, 23-55
uploading patches, 17-13	Unix Config parser, 23-77
Software Library Administration, 2-3	Unix Crontab parser, 23-75

Unix Directory List parser, 23-75 Unix Groups parser, 23-75 Unix GShadow parser, 23-75 Unix Hosts parser, 23-75 Unix INETD parser, 23-75 Unix Installed Patches parser, 23-72 parser parameters, 23-74 Unix Login parser, 23-77 Unix PROFTPD parser, 23-77 Unix Protocols parser, 23-76 Unix Recursive Directory List parser, 23-72 parser parameters, 23-74 Unix Resolve parser, 23-78 Unix Services parser, 23-76 Unix Shadow parser, 23-76 Unix Shadow parser, 23-76 Unix System Crontab parser, 23-78 Unix System Crontab parser, 23-78 Unix System parser, 23-78 Unix VSFTPD parser, 23-78 Unix VSFTPD parser, 23-78 Unix XINETD parser, 23-78 Unix XINETD parser, 23-78 Unix Junix Allour parser, 23-78 Unix System parser, 23-78 Unix System parser, 23-78 Unix System parser, 23-78 Unix System parser, 23-78 Unix Description parser, 23-78 Unix FTPD parser, 23-79 Unix FTPD parser, 23-79 Unix FTPD parser, 23-79 Unix FTPD parser, 23-75 Unix FTPD parser, 23-76 Unix FTPD parser, 23-75 Unix FTPD parser, 23-75 Unix FTPD parser, 23-76 Unix FTPD parser, 23-76 Unix FTPD parser, 23-75 Unix FTPD parser, 23-76 Unix FTPD parser, 23-75 Unix FTPD parser, 23-75 Unix FTPD parser, 23-75 Unix FTPD	violation details compliance standard, 24-50 violations clearing, 24-75 compliance managing, 24-47 viewing examples, 24-50 details compliance framework, 24-50 compliance standard, 24-50 of a target, 24-50 suppressing, 24-47, 24-74 unsuppressing, 24-47 Virtual Private Database, 30-2
	WebAgent parser, 23-78 WebLogic parser (attribute-keyed), 23-68 WebSphere parser (attribute-keyed), 23-68 WebSphere parser (generic), 23-69 Windows Checksum parser, 23-78
	XML default parser (attribute-keyed), 23-68 XML parser (generic), 23-69 XPath conditions and expressions, 23-84 configuration extension, 23-52
value constraint rule, in comparisons, 23-19 Verifying Rolling Upgrade, 19-8 viewing comparison template, 23-12 configuration data, 23-39 configuration extension specification details, 23-53 configuration health problem details, 23-61	