Oracle® Enterprise Manager

Cloud Control Administrator's Guide

12*c* Release 1 (12.1.0.1)

E24473-09

February 2012

ORACLE®

Oracle Enterprise Manager Cloud Control Administrator's Guide, 12*c* Release 1 (12.1.0.1)

E24473-09

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

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

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

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

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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 is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

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

Contents

Preface xvii

Audience xvii
Documentation Accessibility xvii
Related Documents xvii
Conventions xviii

Part I Monitoring and Managing Targets

1 Discovering Targets

1.1 Configuring Automatic Discovery 1-1
1.1.1 Configuring Automatic Discovery of Un-Managed Host Machines Using IP Scan 1-2
1.1.2 Configuring Automatic Discovery of Targets On Managed Hosts 1-3
1.1.3 Checking For and Promoting Discovered Targets 1-5
1.2 Manually Adding Targets 1-5
1.2.1 Manually Adding Host Targets 1-6
1.2.2 Manually Adding Non-Host Targets 1-6

2 Overview of Systems Monitoring

2.1 Monitoring Overview 2-1
2.2 Monitoring Basics 2-1
2.2.1 Out-of-Box Monitoring 2-2
2.3 Monitoring Templates 2-3

3 Using Incident Management

3.1 Monitoring and Managing Via Incidents 3-1
3.2 Events 3-2
3.3 Incidents 3-4
3.3.1 Incident Composed of a Single Event 3-5
3.3.2 Incident Composed of Multiple Events 3-6
3.3.3 Incident Attributes 3-7
3.3.4 Event Prioritization 3-8
3.4 Incident Manager 3-8
3.5 Before Working with Incidents 3-9
3.6 Working with Incidents 3-10

3.6.1 Setting Up Views in Incident Manager ... 3-11
3.6.2 Responding and Working on a Simple Incident ... 3-12
3.6.3 Searching My Oracle Support Knowledge ... 3-13
3.6.4 Open Service Request ... 3-13
3.6.5 Suppressing Incidents and Problems ... 3-13
3.7 Incidents - Advanced Tasks ... 3-14
3.7.1 Creating an Incident Manually ... 3-14
3.7.2 Managing Workload Distribution of Incidents ... 3-14
3.8 Rule Sets ... 3-15
3.8.1 Out-of-Box Rule Sets ... 3-16
3.8.2 Rule Set Types ... 3-17
3.8.3 Rules ... 3-17
3.8.3.1 Rule Criteria ... 3-18
3.8.3.2 Rule Actions ... 3-24
3.8.4 Rule Set Guidelines ... 3-25
3.9 Before Using Rules ... 3-27
3.10 Working with Rules ... 3-29
3.10.1 Creating an Rule ... 3-29
3.10.2 Creating a Rule to Create an Incident ... 3-30
3.10.3 Creating a Rule to Manage Escalation of Incidents ... 3-30
3.10.4 Creating a Rule to Escalate a Problem ... 3-32
3.10.5 Setting Up Automated Notification for Private Rule ... 3-32
3.10.6 Creating a Rule to Receive Notification Regarding Incidents ... 3-33
3.11 Rules - Advanced Tasks ... 3-34
3.11.1 Setting Up a Rule to Send Different Notifications for Different Severity States of an Event 3-34
3.11.2 Creating a Rule to Create a Ticket for Incidents ... 3-35
3.11.3 Creating a Rule to Notify Different Administrators Based on the Event Type ... 3-36
3.11.4 Creating Notification Subscription to Existing Enterprise Rules ... 3-36
3.11.5 Manually Ensuring That There Are No Events That Should Be Incidents ... 3-37
3.12 Problems ... 3-37
3.13 Moving from Enterprise Manager 10/11g to 12c ... 3-38

4 Notifications

4.1 Setting Up Notifications ... 4-1
4.1.1 Setting Up a Mail Server for Notifications ... 4-2
4.1.1.1 Setting Up Repeat Notifications ... 4-3
4.1.2 Setting Up E-mail for Yourself ... 4-5
4.1.2.1 Defining E-mail Addresses ... 4-5
4.1.2.2 Setting Up a Notification Schedule ... 4-9
4.1.2.3 Subscribe to Receive E-mail for Incident Rules ... 4-9
4.1.3 Setting Up E-mail for Other Administrators ... 4-10
4.1.4 E-mail Customization ... 4-11
4.1.4.1 E-mail Customization Reference ... 4-12
4.2 Extending Notification Beyond E-mail ... 4-16
4.2.1 Custom Notification Methods Using Scripts and SNMP Traps ... 4-17
4.2.1.1 Adding a Notification Method based on an OS Command or Script ... 4-17

4.2.1.2 Adding a Notification Method Based on a PL/SQL Procedure 4-33
4.2.1.3 Adding a Notification Method Based on an SNMP Trap 4-48
4.3 Passing Corrective Action Status Change Information 4-49
4.3.1 Passing Corrective Action Execution Status to an OS Command or Script 4-50
4.3.2 Passing Corrective Action Execution Status to a PLSQL Procedure 4-50
4.4 Passing Job Execution Status Information 4-51
4.4.1 Passing Job Execution Status to a PL/SQL Procedure 4-51
4.4.2 Passing Job Execution Status to an OS Command or Script 4-54
4.5 Passing User-Defined Target Properties to Notification Methods 4-54
4.6 Management Information Base (MIB) 4-55
4.6.1 About MIBs 4-55
4.6.2 Reading the MIB Variable Descriptions 4-55
4.6.2.1 Variable Name 4-56
4.6.2.2 MIB Definition 4-56
4.7 Troubleshooting Notifications 4-57
4.7.1 General Setup 4-57
4.7.2 Notification System Errors 4-57
4.7.3 Notification System Trace Messages 4-57
4.7.4 E-mail Errors 4-59
4.7.5 OS Command Errors 4-60
4.7.6 SNMP Trap Errors 4-60
4.7.7 PL/SQL Errors 4-60

5 Managing with Groups

5.1 Introduction to Groups 5-1
5.2 Managing Groups 5-2
5.2.1 Using the Groups Page 5-2
5.2.2 About the Group Home Page 5-2
5.2.3 About the Group Charts Page 5-4
5.2.4 About the Group Members Page 5-4
5.2.5 Viewing Group Status History 5-4
5.2.6 About the System Dashboard 5-4
5.3 About Out-of-Box Reports 5-5
5.4 About Redundancy Systems 5-6
5.5 About Privilege Propagating Groups 5-6
5.5.1 Creating Privilege Propagating Groups 5-7
5.5.2 Using the Group Administration Privilege 5-7
5.5.3 Adding Members to Privilege Propagating Groups 5-7
5.5.4 Converting Conventional Groups to Privilege Propagating Groups 5-8

6 Administration Groups

6.1 What is an Administration Group? 6-1
6.1.1 Developing an Administration Group 6-3
6.2 Planning 6-3
6.3 Implementing Administration Groups and Template Collections 6-8
6.3.1 Creating an Administration Group 6-9

6.3.1.1 Accessing the Administration Group Home Page 6-9
6.3.1.2 Defining a Hierarchy 6-10
6.3.1.3 Defining Template Collections 6-13
6.3.1.4 Associating Template Collections with Administration Groups 6-16
6.4 Removing Administration Groups 6-22

7 Metric Extensions

7.1 What are Metric Extensions? 7-1
7.2 Metric Extension Lifecycle 7-3
7.3 Working with Metric Extensions 7-6
7.3.1 Administrator Privilege Requirements 7-6
7.3.2 Granting Create Metric Extension Privilege 7-7
7.3.3 Creating a New Metric Extension 7-7
7.3.4 Creating a New Metric Extension (Create Like) 7-9
7.3.5 Editing a Metric Extension 7-10
7.3.6 Creating the Next Version of an Existing Metric Extension 7-10
7.3.7 Importing a Metric Extension 7-11
7.3.8 Exporting a Metric Extension 7-12
7.3.9 Deleting a Metric Extension 7-12
7.3.10 Granting Edit/Full Access to Metric Extensions 7-12
7.3.11 Deploying Metric Extensions to a Group of Targets 7-13
7.3.12 Updating Older Versions of Metric Extensions Already deployed to a Group of Targets 7-13
7.4 Adapters 7-14
7.4.1 OS Command Adapter - Single Column 7-14
7.4.2 OS Command Adapter- Multiple Values 7-17
7.4.3 OS Command Adapter - Multiple Columns 7-18
7.4.4 SQL Adapter 7-19
7.4.5 SNMP (Simple Network Management Protocol) Adapter 7-20
7.4.6 JMX Adapter 7-21
7.5 Converting User-defined Metrics to Metric Extensions 7-22
7.5.1 Overview 7-22
7.5.2 Commands 7-23
7.6 Metric Extension Command Line Verbs 7-26

8 Utilizing the Job System and Corrective Actions

8.1 Job System Purpose and Overview 8-1
8.1.1 What Are Job Executions and Job Runs? 8-2
8.1.2 Operations on Job Executions and Job Runs 8-2
8.2 Preliminary Considerations 8-3
8.2.1 Creating Scripts 8-3
8.2.2 Sharing Job Responsibilities 8-4
8.2.3 Submitting Jobs for Groups 8-4
8.3 Creating Jobs 8-4
8.3.1 Selecting a Job Type 8-4
8.3.2 Creating an OS Command Job 8-5
8.3.2.1 Specifying a Single Operation 8-10

8.3.2.2 Specifying a Script ... 8-10
8.3.2.3 Access Level Rules ... 8-12
8.3.3 Creating a SQL Script Job ... 8-12
8.3.3.1 Specifying Targets ... 8-12
8.3.3.2 Specifying Options for the Parameters Page ... 8-12
8.3.3.3 Specifying Host and Database Credentials ... 8-13
8.3.3.4 Returning Error Codes from SQL Script Jobs ... 8-14
8.3.4 Creating a Multi-task Job ... 8-14
8.3.4.1 Job Capabilities ... 8-15
8.3.4.2 Specifying Targets for a Multi-task Job ... 8-15
8.3.4.3 Adding Tasks to the Job ... 8-15
8.4 Analyzing Job Activity ... 8-16
8.5 Generating Job Event Criteria ... 8-16
8.5.1 Enabling Events For Job Status and Targetless Jobs ... 8-17
8.5.2 Adding Targets To Generate Events For Job Status ... 8-18
8.6 Creating Event Rules For Job Status Change ... 8-18
8.6.1 Creating Job Status Change Event Rules For Jobs ... 8-18
8.6.2 Creating Job Status Change Event Rules For Targets ... 8-21
8.7 Creating Corrective Actions ... 8-24
8.7.1 Providing Credentials ... 8-25
8.7.2 Creating Corrective Actions for Metrics ... 8-25
8.7.3 Creating a Library Corrective Action ... 8-26
8.7.4 Specifying Access to Corrective Actions ... 8-27
8.7.4.1 Defining or Modifying Access ... 8-27
8.7.4.2 Access Level Rules ... 8-27
8.7.5 Setting Up Notifications for Corrective Actions ... 8-28
8.7.6 Providing Agent-side Response Actions ... 8-29
8.7.6.1 Specifying Commands and Scripts ... 8-29
8.7.6.2 Using Target Properties in Commands ... 8-29
8.7.6.3 Using Advanced Capabilities ... 8-30

9 Configuring Software Library

9.1 Overview of Software Library ... 9-1
9.2 Users, Roles, and Privileges ... 9-3
9.3 Software Library Storage ... 9-4
9.3.1 Upload File Locations ... 9-5
9.3.2 Referenced File Location ... 9-7
9.4 Prerequisites for Configuring Software Library ... 9-7
9.5 Configuring Software Library Storage Location ... 9-8
9.5.1 Configuring an OMS Shared Filesystem Location ... 9-8
9.5.2 Configuring an OMS Agent Filesystem Location ... 9-9
9.5.3 Configuring a Referenced File Location ... 9-9
9.6 Using Software Library Entities ... 9-11
9.7 Tasks Performed Using the Software Library Home Page ... 9-12
9.7.1 Organizing Entities ... 9-12
9.7.2 Creating Entities ... 9-13
9.7.2.1 Creating Generic Components ... 9-13

9.7.2.2 Creating Directives ... 9-15
9.7.3 Customizing Entities ... 9-17
9.7.4 Managing Entities ... 9-17
9.7.4.1 Granting or Revoking Privileges ... 9-18
9.7.4.2 Moving Entities ... 9-18
9.7.4.3 Changing Entity Maturity ... 9-19
9.7.4.4 Adding Notes to Entities ... 9-19
9.7.4.5 Adding Attachments to Entities ... 9-19
9.7.4.6 Viewing, Editing, and Deleting Entities ... 9-20
9.7.4.7 Searching Entities ... 9-20
9.7.4.8 Exporting Entities ... 9-21
9.7.4.9 Importing Entities ... 9-22
9.8 Maintaining Software Library ... 9-22
9.8.1 Periodic Maintenance Tasks ... 9-23
9.8.2 Re-Importing Oracle Owned Entity Files ... 9-23
9.8.3 Deleting Software Library Storage Location ... 9-23
9.8.4 Purging Deleted Entity Files ... 9-24
9.8.5 Backing Up Software Library ... 9-24

Part II Security

10 Configuring Security

10.1 About Oracle Enterprise Manager Security ... 10-1
10.2 Enterprise Manager Authentication ... 10-2
10.2.1 Repository-Based Authentication ... 10-3
10.2.2 Oracle Access Manager Single Sign-On ... 10-4
10.2.3 Single Sign-On Based Authentication ... 10-4
10.2.3.1 Registering Enterprise Manager as a Partner Application ... 10-5
10.2.3.2 Removing Single Sign-On Configuration ... 10-6
10.2.3.3 Registering Single Sign-On Users as Enterprise Manager Administrators ... 10-7
10.2.3.4 Bypassing the Single Sign-On Logon Page ... 10-9
10.2.3.5 Restoring the Default Authentication Method ... 10-9
10.2.4 Enterprise User Security Based Authentication ... 10-10
10.2.4.1 Registering Enterprise Users as Enterprise Manager Users ... 10-10
10.2.5 Microsoft Active Directory Based Authentication ... 10-11
10.2.5.1 Configuring WebLogic Server Authentication ... 10-13
10.2.5.2 Manage Active Directory Users with External Roles ... 10-15
10.2.5.3 Password Management for Active Directory Users ... 10-15
10.2.5.4 Remove Active Directory Users ... 10-15
10.2.5.5 Remove Active Directory Authentication ... 10-16
10.3 Enterprise Manager Authorization ... 10-16
10.3.1 Authentication Scheme ... 10-16
10.3.2 Classes of Users ... 10-16
10.3.3 Privileges and Roles ... 10-17
10.3.3.1 Granting Privileges ... 10-19
10.4 Configuring Secure Communication (SSL) for Cloud Control ... 10-30
10.4.1 About Enterprise Manager Framework Security ... 10-30

10.4.2 Enabling Security for the Oracle Management Service ... 10-30
10.4.2.1 Creating a New Certificate Authority ... 10-33
10.4.2.2 Viewing the Security Status and OMS Port Information ... 10-34
10.4.2.3 Configuring Transparent Layer Security ... 10-34
10.4.3 Securing the Oracle Management Agent ... 10-35
10.4.4 Enabling Security with Multiple Management Service Installations ... 10-37
10.4.5 Restricting HTTP Access to the Management Service ... 10-37
10.4.6 Managing Agent Registration Passwords ... 10-39
10.4.6.1 Using the Cloud Control Console to Manage Agent Registration Passwords 10-39
10.4.6.2 Using emctl to Add a New Agent Registration Password ... 10-40
10.4.7 Configuring the OMS with Server Load Balance ... 10-40
10.4.8 Enabling Security for the Management Repository Database ... 10-41
10.4.8.1 About Oracle Advanced Security and the sqlnet.ora Configuration File ... 10-41
10.4.8.2 Configuring the Management Service to Connect to a Secure Management Repository Database 10-42
10.4.8.3 Enabling Oracle Advanced Security for the Management Repository ... 10-43
10.4.8.4 Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database 10-44
10.4.9 Configuring Third Party Certificates ... 10-44
10.4.9.1 Configuring Third Party Certificate for HTTPS Upload Virtual Host ... 10-44
10.4.9.2 Configuring Third Party Certificate for HTTPS Console Virtual Host ... 10-45
10.5 Accessing Managed Targets ... 10-45
10.5.1 Credential Subsystem ... 10-46
10.5.1.1 **Named Credential** ... 10-46
10.5.1.2 Job Credentials ... 10-48
10.5.1.3 Monitoring Credentials ... 10-48
10.5.1.4 Collection Credentials ... 10-49
10.5.1.5 Preferred Credentials ... 10-49
10.5.1.6 Managing Credentials Using EM CLI ... 10-50
10.5.2 Setting Up SSH Key-based Host Authentication ... 10-51
10.5.3 Pluggable Authentication Modules (PAM) Support for Hosts ... 10-53
10.5.3.1 Configuring PAM for RHEL4 Users ... 10-53
10.5.3.2 Configuring PAM for AIX Users ... 10-54
10.5.4 Sudo and PowerBroker Support ... 10-54
10.5.4.1 Creating a Privilege Delegation Setting ... 10-55
10.6 Cryptographic Support ... 10-56
10.6.1 Configuring the emkey ... 10-56
10.6.2 emctl Commands ... 10-57
10.6.2.1 emctl status emkey ... 10-57
10.6.2.2 emctl config emkey -copy_to_credstore ... 10-58
10.6.2.3 emctl config emkey -copy_to_repos ... 10-58
10.6.2.4 emctl config emkey -copy_to_file_from_credstore ... 10-58
10.6.2.5 emctl config emkey -copy_to_file_from_repos ... 10-58
10.6.2.6 emctl config emkey -copy_to_credstore_from_file ... 10-59
10.6.2.7 emctl config emkey -copy_to_repos_from_file ... 10-59
10.6.2.8 emctl config emkey -remove_from_repos ... 10-59
10.6.3 Install and Upgrade Scenarios ... 10-59

10.6.3.1 Installing the Management Repository 10-60
10.6.3.2 Installing the First Oracle Management Service 10-60
10.6.3.3 Upgrading from 10.2 or 11.1 to 12.1 10-60
10.6.3.4 Recreating the Management Repository 10-60
10.7 Setting Up the Auditing System for Enterprise Manager 10-60
10.7.1 Configuring the Enterprise Manager Audit System 10-61
10.7.2 Configuring the Audit Data Export Service 10-61
10.7.3 Updating the Audit Settings 10-61
10.7.4 Searching the Audit Data 10-62
10.7.5 List of Operations Audited 10-64
10.8 Additional Security Considerations 10-69
10.8.1 Changing the SYSMAN and MGMT_VIEW Passwords 10-69
10.8.1.1 Changing the SYSMAN User Password 10-69
10.8.1.2 Changing the MGMT_VIEW User Password 10-70
10.8.2 Responding to Browser-Specific Security Certificate Alerts 10-71
10.8.2.1 Responding to the Internet Explorer Security Alert Dialog Box 10-71
10.8.2.2 Responding to Mozilla Firefox New Site Certificate Dialog Box 10-74
10.8.3 Configuring Beacons to Monitor Web Applications Over HTTPS 10-75
10.8.4 Patching Oracle Homes When the User is Locked 10-77
10.8.5 Cloning Oracle Homes 10-77

Part III Generating Reports

11 Using Information Publisher

11.1 About Information Publisher 11-1
11.2 Out-of-Box Report Definitions 11-2
11.3 Custom Reports 11-2
11.3.1 Creating Custom Reports 11-2
11.3.2 Report Parameters 11-3
11.3.3 Report Elements 11-3
11.4 Scheduling Reports 11-3
11.4.1 Flexible Schedules 11-3
11.4.2 Storing and Purging Report Copies 11-4
11.4.3 E-mailing Reports 11-4
11.5 Sharing Reports 11-4

Part IV Accessing Enterpise Manager via Mobile Devices

12 Cloud Control Mobile

12.1 Reviewing System Requirements 12-1
12.2 Performing Initial Setup 12-1
12.3 Connecting the First Time 12-2
12.4 Encountering the Login Screen 12-2
12.5 Managing Settings 12-3
12.6 Using Cloud Control Mobile in Incident Manager 12-4
12.7 Working in Cloud Control Mobile 12-6

12.7.1 Viewing Incidents and Problems 12-6
12.7.2 Changing Views 12-7
12.7.3 Performing Actions 12-8
12.8 Learning Tips and Tricks 12-8

Part V Administering Cloud Control

13 Personalizing Cloud Control

13.1 Personalizing a Cloud Control Page 13-1
13.2 Customizing a Region 13-2
13.3 Setting Your Homepage 13-3

14 Maintaining Enterprise Manager

14.1 Overview: Managing the Manager 14-1
14.2 Management Services and Repository 14-2
14.3 Viewing Enterprise Manager Topology and Charts 14-4
14.4 Viewing Enterprise Manager Services 14-6
14.5 Controlling and Configuring Management Agents 14-7
14.5.1 Management Agent Home Page 14-7
14.5.2 Controlling a Single Agent 14-8
14.5.3 Configuring Single Management Agents 14-9
14.5.4 Controlling Multiple Management Agents 14-9
14.5.5 Configuring Multiple Agents 14-11

15 Updating Cloud Control

15.1 Using Self Update 15-1
15.1.1 What Can Be Updated? 15-1
15.2 Setting Up Self Update 15-2
15.2.1 Setting Up Enterprise Manager Self Update Mode 15-2
15.2.2 Assigning Self Update Privileges to Users 15-3
15.2.3 Setting Up the Software Library 15-3
15.2.4 Setting Up the EM CLI Utility (Optional) 15-3
15.3 Applying an Update 15-4
15.3.1 Applying an Update in Online Mode 15-4
15.3.2 Applying an Update in Offline Mode 15-5
15.4 Acquiring or Updating Management Agent Software 15-5
15.4.1 Acquiring Management Agent Software in Online Mode 15-6
15.4.2 Acquiring Management Agent Software in Offline Mode 15-6
15.5 Deploying and Updating Plug-ins 15-7
15.5.1 Deploying a Plug-in 15-8
15.5.1.1 Downloading a Plug-in from the Enterprise Manager Store 15-8
15.5.1.2 Importing an External Archive into Enterprise Manager 15-8
15.5.1.3 Deploying a Plug-in to Oracle Management Service (OMS) 15-9
15.5.1.4 Adding Targets for the Plug-in to Monitor 15-10
15.5.1.5 Important Details Regarding Plug-in Deployment 15-11
15.5.2 Updating a Plug-in 15-11

15.5.2.1 Downloading the Latest Plug-in Archive from the Oracle Enterprise Manager Store 15-12
15.5.2.2 Updating a Plug-in on Oracle Management Service ... 15-12
15.5.2.3 Updating a Plug-in on a Management Agent ... 15-12
15.5.2.4 Un-deploying a Plug-in ... 15-13

16 Patching Enterprise Manager

16.1 Overview ... 16-1
16.2 Patching OMS and Management Repository ... 16-2
16.2.1 OMS Patches ... 16-2
16.2.2 Repository Patches ... 16-2
16.2.3 Applying OMS and Repository Patches ... 16-2
16.3 Patching Enterprise Manager Agents ... 16-3
16.3.1 Management Agent Patches ... 16-3
16.3.1.1 Patches and Updates Versus My Oracle Support ... 16-4
16.3.2 Automated Agent Patching ... 16-4
16.3.2.1 Accessing Patches and Updates ... 16-5
16.3.2.2 Viewing Patch Recommendations ... 16-5
16.3.2.3 Searching Patches ... 16-6
16.3.2.4 Applying Management Agent Patches ... 16-7
16.3.2.5 Verifying the Applied Agent Patches ... 16-12
16.3.2.6 Validating Agent Patch Errors ... 16-12
16.3.2.7 Deinstalling the Applied Agent Patches ... 16-13
16.3.3 Manual Agent Patching ... 16-13

17 Starting and Stopping Enterprise Manager Compone nts

17.1 Controlling the Oracle Management Agent ... 17-1
17.1.1 Starting, Stopping, and Checking the Status of the Management Agent on UNIX 17-1
17.1.2 Starting and Stopping the Management Agent on Windows ... 17-3
17.1.3 Checking the Status of the Management Agent on Windows ... 17-3
17.1.4 Troubleshooting Management Agent Startup Errors ... 17-4
17.1.4.1 Management Agent starts up but is not ready ... 17-4
17.1.4.2 Management Agent fails to start because of time zone mismatch between agent and OMS 17-4
17.1.4.3 Agent fails to start due to possible port conflict ... 17-4
17.1.4.4 Agent secure/unsecure fails ... 17-5
17.2 Controlling the Oracle Management Service ... 17-5
17.2.1 Controlling the Management Service on UNIX ... 17-5
17.2.2 Controlling the Management Service on Windows ... 17-6
17.2.3 Troubleshooting Oracle Management Service Startup Errors ... 17-6
17.3 Guidelines for Starting Multiple Enterprise Manager Components on a Single Host .. 17-7
17.4 Starting and Stopping Oracle Enterprise Manager 12c Cloud Control ... 17-8
17.4.1 Starting Cloud Control and All Its Components ... 17-8
17.4.2 Stopping Cloud Control and All Its Components ... 17-9
17.5 Additional Management Agent Commands ... 17-10
17.5.1 Uploading and Reloading Data to the Management Repository ... 17-10
17.5.2 Specifying New Target Monitoring Credentials ... 17-10

17.5.3 Listing the Targets on a Managed Host ... 17-11
17.5.4 Controlling Blackouts ... 17-11
17.5.5 Changing the Management Agent Time Zone ... 17-14
17.5.6 Reevaluating Metric Collections ... 17-14
17.6 emctl Commands ... 17-16
17.7 Using emctl.log File ... 17-22

18 Locating and Configuring Enterprise Manager Log Files

18.1 Managing Log Files ... 18-1
18.1.1 Viewing Log Files and Their Messages ... 18-3
18.1.2 Searching Log Files ... 18-3
18.1.2.1 Searching Log Files: Basic Searches ... 18-3
18.1.2.2 Searching Log Files: Advanced Searches ... 18-4
18.1.3 Downloading Log Files ... 18-5
18.2 Locating Management Agent Log and Trace Files ... 18-6
18.2.1 About the Management Agent Log and Trace Files ... 18-6
18.2.1.1 Structure of Agent Log Files ... 18-7
18.2.2 Locating the Management Agent Log and Trace Files ... 18-7
18.2.3 Setting Oracle Management Agent Log Levels ... 18-7
18.2.3.1 Setting gcagent.log ... 18-8
18.2.3.2 Setting gcagent_error.log ... 18-8
18.2.3.3 Setting the Log Level for Individual Classes and Packages ... 18-8
18.2.3.4 Setting gcagent_mdu.log ... 18-9
18.2.3.5 Setting the TRACE Level ... 18-11
18.3 Locating and Configuring Oracle Management Service Log and Trace Files ... 18-11
18.3.1 About the Oracle Management Service Log and Trace Files ... 18-11
18.3.2 Locating Oracle Management Service Log and Trace Files ... 18-12
18.3.3 Controlling the Size and Number of Oracle Management Service Log and Trace Files ... 18-12
18.3.4 Controlling the Contents of the Oracle Management Service Trace File ... 18-13
18.3.5 Controlling the Oracle WebLogic Server and Oracle HTTP Server Log Files ... 18-14

19 Maintaining and Troubleshooting the Management Repository

19.1 Management Repository Deployment Guidelines ... 19-1
19.2 Management Repository Data Retention Policies ... 19-2
19.2.1 Management Repository Default Aggregation and Purging Policies ... 19-2
19.2.2 Management Repository Default Aggregation and Purging Policies for Other Management Data 19-3
19.2.3 Modifying the Default Aggregation and Purging Policies ... 19-3
19.2.4 Modifying Data Retention Policies When Targets Are Deleted ... 19-4
19.2.5 How to Modify the Retention Period of Job History ... 19-5
19.2.6 DBMS_SCHEDULER Troubleshooting ... 19-6
19.3 Changing the SYSMAN Password ... 19-7
19.4 Dropping and Recreating the Management Repository ... 19-9
19.4.1 Dropping the Management Repository ... 19-9
19.4.2 Recreating the Management Repository ... 19-10

19.4.2.1 Using the RepManager Script to Create the Management Repository 19-10
19.4.2.2 Using a Connect Descriptor to Identify the Management Repository Database 19-11
19.5 Troubleshooting Management Repository Creation Errors 19-11
19.5.1 Package Body Does Not Exist Error While Creating the Management Repository 19-11
19.5.2 Server Connection Hung Error While Creating the Management Repository 19-12
19.5.3 General Troubleshooting Techniques for Creating the Management Repository 19-12
19.6 Cross Platform Enterprise Manager Repository Migration 19-12
19.6.1 Common Prerequisites 19-13
19.6.2 Methodologies 19-14
19.6.2.1 Cross Platform Transportable Tablespaces 19-14
19.6.2.2 Data Pump 19-16
19.6.2.3 Export/Import 19-19
19.6.3 Post Migration Verification 19-21

Part VI Configuring Enterprise Manager for High Availability

20 High Availability Solutions

20.1 Latest High Availability Information 20-1
20.2 Defining High Availability 20-2
20.2.1 Levels of High Availability 20-2
20.3 Comparing Availability Levels 20-3
20.4 Implementing High Availability Levels 20-4

21 Setting Up High Availability

21.1 Installation Best Practices for Enterprise Manager High Availability 21-1
21.1.1 Configuring the Management Agent to Automatically Start on Boot and Restart on Failure 21-1
21.1.2 Configuring Restart for the Management Agent 21-1
21.1.3 Installing the Management Agent Software on Redundant Storage 21-2
21.2 Installation Best Practices for Enterprise Manager Management Repository High Availability 21-2
21.3 Configuring a Standby Database for the Management Repository 21-3
21.4 Configuring Management Service to RAC Management Repository Communication 21-4
21.5 Configuring the First Management Service for High Availability 21-5
21.5.1 Management Service Install Location 21-5
21.6 Configuring the Cloud Control OMS in an Active/Passive Environment for High Availability Failover Using Virtual Host Names 21-6
21.6.1 Overview and Requirements 21-6
21.6.2 Installation and Configuration 21-6
21.6.3 Setting Up the Virtual Host Name/Virtual IP Address 21-7
21.6.4 Setting Up Shared Storage 21-7
21.6.5 Setting Up the Environment 21-7
21.6.6 Synchronizing Operating System IDs 21-8
21.6.7 Setting Up Shared Inventory 21-8
21.6.8 Installing the Software 21-8

21.6.9 Starting Up Services 21-9
21.7 Configuring the Software Library 21-9
21.8 Configuring a Load Balancer 21-9
21.8.1 SLB Setup 21-10
21.8.2 Enterprise Manager Side Setup 21-12
21.9 Configuring Standby Management Services on a Standby Site 21-13
21.9.1 Installing the First Standby Management Service 21-13
21.9.2 Installing Additional Standby Management Services 21-15
21.9.3 Validating Your Installation and Complete the Setup 21-16
21.9.3.1 Keeping the Standby Site in Sync 21-16

22 Backing Up Enterprise Manager

22.1 Backing Up Your Deployment 22-1
22.2 Management Repository Backup 22-1
22.3 Oracle Management Service Backup 22-2
22.4 Management Agent Backup 22-3

23 Enterprise Manager Outages

23.1 Recovery of Failed Enterprise Manager Components 23-1
23.1.1 Repository Recovery 23-1
23.1.2 Recovery Scenarios 23-3
23.1.2.1 Full Recovery on the Same Host 23-3
23.1.2.2 Incomplete Recovery on the Same Host 23-3
23.1.2.3 Full Recovery on a Different Host 23-4
23.1.2.4 Incomplete Recovery on a Different Host 23-4
23.1.3 Recovering the OMS 23-5
23.1.4 OMS Recovery Scenarios 23-6
23.1.4.1 Single OMS, No Server Load Balancer (SLB), OMS Restored on the same Host 23-6
23.1.4.2 Single OMS, No SLB, OMS Restored on a Different Host 23-8
23.1.4.3 Single OMS, No SLB, OMS Restored on a Different Host using the Original Hostname 23-9
23.1.4.4 Multiple OMS, Server Load Balancer, Primary OMS Recovered on the Same Host .. 23-11
23.1.4.5 Multiple OMS, Server Load Balancer configured, Primary OMS Recovered on a Different Host 23-13
23.1.4.6 Multiple OMS, SLB configured, additional OMS recovered on same or different host 23-15
23.1.5 Recovering Management Agents 23-15
23.1.6 Management Agent Recovery Scenarios 23-16
23.1.6.1 Management Agent Reinstall Using the Same Port 23-16
23.1.6.2 Management Agent Restore from Filesystem Backup 23-16
23.2 Recovering from a Simultaneous OMS-Management Repository Failure 23-17
23.2.1 Collapsed Configuration: Incomplete Management Repository Recovery, Primary OMS on the Same Host 23-17
23.2.2 Distributed Configuration: Incomplete Management Repository Recovery, Primary OMS and additional OMS on Different Hosts, SLB Configured 23-17

23.3 Switching Over or Failing Over to Standby Enterprise Manager Configurations 23-18
23.3.1 Switchover 23-18
23.3.2 Failover 23-20
23.3.3 Automatic Failover to the Standby Site in a Level 4 MAA Configuration 23-22

Part VII Engineered Systems Management

24 Discovering and Managing Exadata Targets and Systems

24.1 Automatically Discovering an Oracle Database Machine 24-2
24.2 Viewing the Topology of an Existing DB Machine Target 24-4
24.3 Drilling Down to Individual Targets 24-4
24.4 Viewing Critical Hardware Information for the DB Machine 24-5
24.5 Viewing DB Machine Alerts 24-5
24.6 Adding Exadata Components Manually 24-5
24.7 About Oracle Exadata Storage Server 24-6
24.7.1 Using Exadata As a Cloud Control Target 24-6
24.7.2 Performing Administration Tasks on Exadata Cells 24-7
24.7.3 Performing Administration Tasks on Infiniband Networks 24-8
24.7.4 Launching the IORM Performance Page 24-8
24.7.5 Viewing an Exadata Cell Configuration 24-9
24.7.6 Managing a Single I/O Resource Management Allocation 24-9
24.7.7 Accessing Oracle Support Workbench for Exadata Cell 24-10
24.7.8 Changing the IORM Mode and Updating the IORM Objective 24-10

25 Using Oracle Exalogic Elastic Cloud

25.1 Using the Exalogic Elastic Cloud Discovery Wizard 25-1
25.2 Displaying and Using the Exalogic Elastic Cloud Home Page and Dashboard 25-2
25.3 Viewing Application Deployments in Exalogic Elastic Cloud Targets 25-3
25.4 Viewing WebLogic Domains in Exalogic Elastic Cloud Targets 25-3
25.5 Viewing Coherence Clusters in Exalogic Elastic Cloud Targets 25-4
25.6 Viewing Hosts in Exalogic Elastic Cloud Targets 25-5

Index

Preface

This guide describes how to set up a Private Cloud, manage and deploy virtualization targets with Oracle Enterprise Manager 12c Release 1.

The preface covers the following:

- Audience
- Documentation Accessibility
- Related Documents
- Conventions

Audience

This document is intended for cloud administrators who want to setup and manage the cloud infrastructure. It is also intended for administrators and users of the Self Service Portal.

Documentation Accessibility

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

Access to Oracle Support

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

Related Documents

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

http://www.oracle.com/technetwork/documentation/index.html#em

Oracle Enterprise Manager also provides extensive Online Help. Click **Help** at the top of any Enterprise Manager page to display the online help window.

Conventions

The following text 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.
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

Monitoring and Managing Targets

This section contains the following chapters:

- Discovering Targets
- Overview of Systems Monitoring
- Using Incident Management
- Notifications
- Metric Extensions
- Managing with Groups
- Administration Groups
- Utilizing the Job System and Corrective Actions
- Configuring Software Library

1 Discovering Targets

Oracle components that are managed and monitored by Cloud Control, such as an Oracle Database or an Oracle WebLogic Server domain, are known as "managed targets".

Before a target can be managed, a Management Agent must first be installed on the host machine the target is running on. The target itself must then be assigned to a Management Agent, thereby promoting it to managed target status.

Cloud Control offers two modes for installing Management Agents to monitor and manage potential targets: Manually or using automatic target discovery.

The manual process is exactly that: You explicitly add a specific host or Oracle component as a target to bring under management.

In automatic discovery, you create an Enterprise Manager job that searches host machines for possible targets on a regularly-scheduled basis. A key benefit is that as new Oracle components are added to your infrastructure, they can be found and brought under management.

See the following sections for more information about discovering and adding new targets:

- Configuring Automatic Discovery
- Manually Adding Targets

1.1 Configuring Automatic Discovery

Automatic discovery refers to the process of scanning host machines for Oracle components that can be managed and monitored by Enterprise Manager Cloud Control. You can decide upon a schedule for discovery, the target types to be discovered and the hosts to scan for targets. Discovered targets can then be promoted to managed target status, enabling Enterprise Manager to manage the targets.

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

See the following sections for instructions on using the various self-discovery options:

- Configuring Automatic Discovery of Un-Managed Host Machines Using IP Scan
- Configuring Automatic Discovery of Targets On Managed Hosts
- Checking For and Promoting Discovered Targets

1.1.1 Configuring Automatic Discovery of Un-Managed Host Machines Using IP Scan

In automatic host discovery, a single Management Agent is tasked to scan the entire network based on IP address ranges that you specify. It then returns a list of "un-managed" host machines - that is, host machines that do not yet have a Management Agent installed - with a list of ports in use that match the ranges you specified. The name of the service using each port is also returned.

By looking at the list of services and ports, you should be able to determine what types of Oracle components have been discovered. For example, if a host is returned with port 7001 in use, you can assume that this port is associated with an Oracle WebLogic Server domain that can be promoted to managed target status.

The next step is to deploy Management Agents to the hosts you want to promote to managed status. Once a Management Agent is deployed to the host, any Oracle components running on the host will be discovered and reported as potential targets. These components can then be promoted to managed target status, enabling them to be managed and monitored by Cloud Control.

To be most effective, automatic discovery should ideally be run by a network administrator with an overview of what Oracle components are running on what ports.

Note that because the network will be scanned, the Sudo Privilege Delegation must be set on the Management Agent host that will perform the scan. Typically you will use the Management Agent installed by default on the Oracle Management Service host that Cloud Control is running, which means you can set this privilege on the Cloud Control machine. Privileges are managed via the Manage Privilege Delegation Settings page, which is accessed by selecting **Security**, then **Privilege Delegation**from the **Setup** menu. For more information on privileges, see Chapter 10, "Configuring Security".

To discover and configure hosts using IP scan, follow these steps:

1. Log into Enterprise Manager Cloud Control.
2. From the **Setup** menu, select **Add Target**, then select **Configure Auto Discovery**.
3. Click the **Configure** icon in the Hosts and Virtual Server Discovery Using IP Scan in the Configure Auto Discovery table.
4. Click **Create**. You will not create the discover job. By default, the Name field will be populated with a title including that date and time the job was created. Note that you can edit the discovery jobs and schedule discovery to run immediately or later.
5. Click **Add**. You will now select the Management Agent that will perform the network scan. You can select the Management Agent that is installed by default on the Oracle Management Service host, or can select another Agent if desired.

   Note that because the entire network will be scanned, the Sudo Privilege Delegation must be set on the Management Agent host that will perform the scan.
6. Select the agent in the IP Ranges for scan table, and enter the IP ranges to scan. You can specify any or even all of the following:
   - One or more absolute hostnames, each separated by a space; for example: `host1.example.com host3.example.com`
   - One or more IP addresses, each separated by a space

- A range of addresses; for example: `10.0.0-255.1-250`. Note that IP addresses and IP ranges must be separated by a comma; for example: `10.0.0-255.1-250`
- Classless Inter-Domain Routing (CIDR) notations; for example: `128.16.10.0/24`

Separate each value with a space; for example:

```
host1.example.com 192.168.0.1 128.16.10.0/24
10.0.0-255.1-250,254
```

7. A default list of ports to scan within the IP ranges you specified is listed in the Configure Ports table. These are default ports typically used by the listed Oracle components.

   To modify the port values for a component, select the component in the table and change the values accordingly. Up to 10 ports and/or port ranges can be specified.

8. If you want to add more component ports to the list, click **Add**. Enter the name of the service to include, and specify the port(s) or port range to scan.

9. Click **Save and Submit IP Scan** when finished. The Job Details panel now opens. This is where you can specify:

   - The schedule at which the discovery job will run. Note that you can start the job immediately.
   - The credentials set on the Management Agent that will perform the scan.

     As noted, the Sudo Privilege Delegation must be set on the Management Agent host that will perform the scan. The named credential that will be used must be configured to run as root.

10. After the discovery job executes, you can check for discovered hosts that may contain potential targets. You can do this two ways:

    - Select the job in the Host Discovery page, then click **View Discovered Targets**;

      or:

    - From the **Setup** menu, select **Add Target**, then select **Auto Discovery Results**.

11. Click the **Host Targets** tab. All discovered hosts are listed, with the open ports and identifiable service names shown. Based on your understanding of the Oracle components deployed on your network, you should be able to determine the types of potential targets that have been discovered.

12. Select a host in the table, then click **Promote** to promote the host to managed target status. The Add Host Targets wizard opens. You will use this wizard to install a Management Agent on the host.

    For instructions on installing a Management Agent, see "Installing Oracle Management Agent" in the *Enterprise Manager Cloud Control Basic Installation Guide*.

13. After the Management Agent has been successfully installed on the host, targets running on the host will be discovered as potential targets. See Section 1.1.3, "Checking For and Promoting Discovered Targets" for details on promoting targets.

1.1.2 Configuring Automatic Discovery of Targets On Managed Hosts

By default, automatic discovery of new Oracle targets on hosts that are already managed targets - that is, host machines that have a Management Agent installed - is

pre-configured to run on each host. Automatic target discovery is the most efficient way to discover potential targets on managed hosts, as Cloud Control can search one or more hosts for multiple types of targets at the same time.

Automatic discovery is enabled for all supported target types by default except for Oracle Fusion Middleware, which requires that a search paramter be provided.

The automatic discovery configuration is defined within a"discovery module", which you can modify to suit your requirements. You can schedule discovery to run on all hosts in the discovery module at the same interval, or can configure separate schedules for each host.

To configure automatic discovery on one or more managed hosts, follow these steps:

1. Login into Enterprise Manager Cloud Control.
2. From the **Setup** menu, select **Add Target**, then select **Configure Auto Discovery**.
3. Click the **Configure** icon in the **Multiple Target-type Discovery on Single Host** row in the Discovery table.
4. Expand **Search**, then enter the hostname for the host you want to check for targets in the Agent Host Name field. The host must have a Management Agent installed on it.
5. Click **Search**. The host will be added to the table below.
6. Select the host in the table and click **Configure**.
7. Set the schedule at which the discovery job will be run, in days. This schedule will be applied to all selected hosts. By default the job will run every 24 hours.
8. Select the Oracle component types you want to search the host for. Note that you must supply search parameters for some target types. To specify a parameter, select the target type in the Discovery Module column and click **Edit Parameters**.
   - Oracle Cluster and High Availability Service: No parameters required
   - Oracle Database, Listener and Automatic Storage Management: Specify the path to the Clusterware Home.
   - Oracle Home Discovery: No parameters required.
   - Oracle Secure Backup: No parameters required.
   - Oracle Fusion Middleware: Specify * (the "star" character) to search all Middleware Homes, or specify the path to one or more Middleware Homes on the host, each separated by a comma.
9. Click **OK** when finished. The host has been added to the discovery module.
10. Repeat these steps for each additional host you want to add to the discovery module.
11. Click **Run Discovery Now** to discover targets immediately.

    Note that the discovery job will also run at the scheduled daily interval.
12. Once targets have been discovered, you can promote them to managed status. See Section 1.1.3, "Checking For and Promoting Discovered Targets" for details on promoting targets.

1.1.3 Checking For and Promoting Discovered Targets

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

To promote discovered targets to managed status, follow these steps:

1. Log in into Enterprise Manager.
2. After the discovery job executes, you can check for discovered hosts that may contain potential targets. You can do this two ways:
   - Select the job in the Host Discovery page, then click **View Discovered Targets**; or
   - From the **Setup** menu, select **Add Target**, then select **Auto Discovery Results**.
3. Select a target to promote, then click **Promote**. A wizard specific to the target type you are promoting opens. Supply the required values.
4. Click the **Non-Host Targets**.You can choose one or several targets to promote.
5. Note that you can optionally click **Ignore** for a discovered target, essentially marking it to be processed at a later time.

   Ignored targets will be displayed in the Ignored Targets tab, and will remain in Cloud Control as un-managed targets until you decide to either promote or remove them.
6. Check the target type home page to verify that the target is promoted as an Enterprise Manager target. Once a target is successfully promoted, the Management Agent installed on the target host will begin collecting metric data on the target.

> **Note:** Enterprise Manager does not support simultaneous promotion of multiple targets. Additionally, multiple selection of database targets has been disabled to avoid a user selecting RAC databases across clusters. This is similar to the user-guided discovery feature where a user cannot discover targets across a cluster in the same session.

1.2 Manually Adding Targets

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

You must be able to specify the properties of a target to be managed and create an Enterprise Manager managed target.

Not all target types can be manually added. During registration with the discovery framework, the target type owner indicates whether a target type can be manually added or not.

See the following sections for instructions:

- Manually Adding Host Targets
- Manually Adding Non-Host Targets

1.2.1 Manually Adding Host Targets

A wizard guides you through the process of manually deploying a Management Agent to a new host target.

For instructions on installing a Management Agent, see "Installing Oracle Management Agent" in the *Enterprise Manager Cloud Control Basic Installation Guide*.

1.2.2 Manually Adding Non-Host Targets

A configuration page or wizard based on target type metadata listing all the instance properties required to manage target is displayed.

You can specify a name for the target and provide the required configuration information.

To add targets manually to Enterprise Manager, follow these steps:

1. Log in into Enterprise Manager.
2. From the **Setup,** select **Add Target**, then select **Add Targets Manually** from the drop-down menus. Enterprise Manager displays the Add Targets Manually page.
3. Under the Add Targets Manually page, go to the Add Targets Manually sub-section and choose an option:
   - Add Non-Host Targets Using Guided Process

     Choose one of the target types to add, such as **Oracle Cluster and High Availability Service**, **Oracle Database Machine**, or **WebLogic Domain Discovery**. This process will also add related targets.
   - Add Non-Host Targets by Specifying Target Monitoring Properties

     Choose one of the target types to add, such as **Fusion J2EE Application**, **Applications Utilities**, or **Supplier Portal**.
4. After you select the target type, you will follow a wizard specific to the target type to add the target.

   Upon confirmation, the target becomes a managed target in Enterprise Manager. Enterprise Manager simply accepts the information, performs validation of the supplied data where possible and starts monitoring the target.

2

Overview of Systems Monitoring

2.1 Monitoring Overview

Enterprise Manager monitoring features provide increased out-of-box value, automation, and monitoring support to enable IT organizations to maximize operational efficiencies and provide high quality services. For applications that are built on Oracle, Enterprise Manager offers the most comprehensive monitoring of the Oracle environment. To support the myriad and variety of applications built on Oracle, Enterprise Manager expands its monitoring scope to non-Oracle components, such as third-party application servers, hosts, firewalls, server load balancers, and storage.

Enterprise Manager provides the most comprehensive management features for all Oracle products. For example, Enterprise Manager's monitoring functionality is tightly integrated with Oracle Database manageability features such as server-generated alerts. These alerts are generated by the database itself about problems it has self-detected. Server-generated alerts can be managed from the Enterprise Manager console and include recommendations on how problems can be resolved. Performance problems such as poorly performing SQL and corresponding recommendations that are generated by the database's self-diagnostic engine, called Automatic Database Diagnostic Monitor (ADDM), are also captured and exposed through the Enterprise Manager console. This allows Enterprise Manager administrators to implement ADDM recommendations with ease and convenience.

Adding targets to monitor is simple. Enterprise Mangager provides you with the option of either adding targets manually or automatically discovering all monitorable targets on a host. Once your targets have been added to your monitored environment, Enterprise Manager also makes it easy to expand the scope of enterprise monitoring beyond individual components. Using Enterprise Manager's group management functionality, you can easily organize monitorable targets into groups, allowing you to monitor and manage many components as one. Enterprise Manager can also automatically and intelligently apply monitoring settings for every added target.

2.2 Monitoring Basics

Enterprise monitoring functionality permits unattended monitoring of your IT environment. Enterprise Manager comes with a comprehensive set of performance and health metrics that allows monitoring of key components in your environment, such as applications, application servers, databases, as well as the back-end components on which they rely (such as hosts, operating systems, storage).

The Management Agent on each monitored host monitors the status, health, and performance of all managed components (also referred to as targets) on that host. If a target goes down, or if a performance metric crosses a warning or critical threshold, an "event" is created and sent to Enterprise Manager and to Enterprise Manager administrators who have registered interest in receiving such notifications. Enterprise monitoring functionality and the mechanisms that support this functionality are discussed in the following sections.

When it is not practical to have a Management Agent present to monitor specific components of your IT infrastructure, as might be the case with an IP traffic controller or remote Web application, Enterprise Manager provides Extended Network and Critical URL Monitoring functionality. This feature allows the Beacon functionality of the Agent to monitor remote network devices and URLs for availability and responsiveness without requiring an Agent to be physically present on that device. You simply select a specific Beacon, and add key network components and URLs to the Network and URL Watch Lists. More information about using this feature is available in the Enterprise Manager online help and from the Oracle Technology Network Web site.

2.2.1 Out-of-Box Monitoring

Enterprise Manager's Management Agents automatically start monitoring their host's systems (including hardware and software configuration data on these hosts) as soon as they are deployed and started. Enterprise Manager provides auto-discovery scripts that enable these Agents to automatically discover all Oracle components and start monitoring them using a comprehensive set of metrics at Oracle-recommended thresholds. This monitoring functionality includes other components of the Oracle ecosystem such as NetApp Filer, BIG-IP load balancers, Checkpoint Firewall, and IBM WebSphere and Oracle WebLogic application servers. Metrics from all monitored components are stored and aggregated in the Management Repository, providing administrators with a rich source of diagnostic information and trend analysis data. When critical alerts are detected, notifications are sent to administrators for rapid resolution.

Out-of-box, Enterprise Manager monitoring functionality provides:

- In-depth monitoring with Oracle-recommended metrics and thresholds.
- Montioring of all components of your IT infrastructure (Oracle and non-Oracle) as well as the applications and services that are running on them.
- Access to real-time performance charts.
- Collection, storage, and aggregation of metric data in the Management Repository. This allows you to perform strategic tasks such as trend analysis and reporting.
- E-mail and pager notifications for detected critical events.

Enterprise Manager can monitor a wide variety of components (such as databases, hosts, and routers) within your IT infrastructure.

Some examples of monitored metrics are:

- Archive Area Used (Database)
- Component Memory Usage (Application Server)
- Segments Approaching Maximum Extents Count (Database)
- Network Interface Total I/O Rate (Host)

Some metrics have associated predefined limiting parameters called thresholds that cause alerts to be triggered when collected metric values exceed these limits. Enterprise Manager allows you to set metric threshold values for two levels of alert severity:

- **Warning** - Attention is required in a particular area, but the area is still functional.
- **Critical** - Immediate action is required in a particular area. The area is either not functional or indicative of imminent problems.

Hence, thresholds are boundary values against which monitored metric values are compared. For example, for each disk device associated with the Disk Utilization (%) metric, you might define a warning threshold at 80% disk space used and critical threshold at 95%.

2.3 Monitoring Templates

Monitoring templates simplify the task of standardizing monitoring settings across your enterprise by allowing you to specify the monitoring settings once and apply them to your monitored targets. This makes it easy for you to apply specific monitoring settings to specific classes of targets throughout your enterprise. For example, you can define one monitoring template for test databases and another monitoring template for production databases.

A monitoring template defines all Enterprise Manager parameters you would normally set to monitor a target, such as:

- Target type to which the template applies.
- Metrics (including user-defined metrics), thresholds, metric collection schedules, and corrective actions.

When a change is made to a template, you can reapply the template across affected targets in order to propagate the new changes. You can reapply the monitoring templates as often as needed. For any target, you can preserve custom monitoring settings by specifying metric settings that can never be overwritten by a template.

Comparing Differences Between Targets and Monitoring Templates

Deciding how and when to apply a template is simplified by using the Compare Monitoring Template feature. This feature allows you to see at a glance how metric and policy settings defined in a template differ from those defined on the destination target. Compare Monitoring Template is especially useful when working with aggregate targets such as groups and systems. For example, after you apply a Monitoring Template to a group, you want to verify that the group members now have the same monitoring settings as the template. The Compare Monitoring Template feature makes checking simple. You can also schedule this as a report, allowing you to check periodically if the group members still follow the template settings.

3

Using Incident Management

Incident management allows you to monitor and resolve service disruptions quickly and efficiently.

This chapter covers the following topics:

- Monitoring and Managing Via Incidents
- Events
- Incidents
- Incident Manager
- Before Working with Incidents
- Working with Incidents
- Incidents - Advanced Tasks
- Rule Sets
- Before Using Rules
- Working with Rules
- Rules - Advanced Tasks
- Problems
- Moving from Enterprise Manager 10/11g to 12c

3.1 Monitoring and Managing Via Incidents

Enterprise Manager Cloud Control 12c greatly expands target monitoring and management capability beyond previous releases by letting you focus on what is important from a broader monitoring/management perspective rather than focusing on numerous discrete events that may be relevant to a particular situation.

> **Note:** Also available is the mobile application for managing the incidents and problems on the go. For more information, see Chapter 12, "Cloud Control Mobile"

What is an event?

An event is a a significant occurrence on a managed target that typically indicates something has occurred outside normal operating conditions. Examples of events include: database target down, performance threshold violation, change in application

configuration files, or job failure. An event can also be raised to signal successful operations or a job successfully completed.

Previous versions of Enterprise Manger generated alerts for exception conditions (metric alerts). For Enterprise Manager 12*c*, metric alerts are a type of event , one of many different event types. This event model significantly raises the number of conditions in an IT infrastructure for which Enterprise Manager can detect and raise events across the different functional areas of Enterprise Manager (such as monitoring, compliance, or the job system). Events now provide a uniform way to indicate that something of interest has occurred in a datacenter managed by Enterprise Manager.

What is an incident?

An incident is a situation or issue you need to act on. Of all events raised within your managed environment, there is likely only a subset that you need to act on because they impact your business applications (such as a target down event). An incident is, therefore, a significant event or set of related significant events that need to be managed because they can potentially impact your business applications. To manage these significant subset of events, Enterprise Manager provides incident management features.

Managing incidents is carried out through Incident Manger, which provides you with a central location from which to view, manage, diagnose and resolve incidents as well as identify, resolve and eliminate the root cause of disruptions. See Section 3.4, "Incident Manager" for more information.

When you create an incident, you identify the event(s) for which you want an incident to be created. An incident may consist of a single event, as might be the case when you are only interested in whether a single database down, or something more complex consisting of multiple events as might be the case when monitoring host resources where multiple events such as CPU utilization, memory utilization, swap space utilization events are raised to indicate that machine load is high.

Managing by incidents allows you to focus on the smaller set of important issues in your managed environment. Incident Manager provides a rich set of features to help manage incidents such as the ability to assign incident ownership, track incident resolution status, set incident priority, or set incident escalation level.

3.2 Events

By definition, an event is a significant occurrence within your IT infrastructure that Enterprise Manager can detect and subsequently notify interested parties or take action on. An event has very specific attributes that allow Enterprise Manager (and ultimately an Enterprise Manager administrator) to identify, categorize, and manage the event. All events have the following attributes

- Type
- Severity
- Entity on which the event is raised
- Message
- Timestamp
- Category

Event Types

Previous versions of Enterprise Manager let you monitor and manage by discrete signals and notified you by raising a metric alert as a result of threshold violations. For Enterprise Manager 12*c*, a metric alert is now one of several types of event conditions for which Enterprise Manager can monitor. These event conditions are called *event types*. As shown in the following list, the range of events types greatly expands Enterprise Manager's monitoring flexibility.

The following table lists all available event types.

Event Type	Description
Application Dependency and Performance Alert	Alerts are raised by the Application Dependency and Performance (ADP) monitoring when metrics related to a J2EE application or component have crossed some thresholds.
Compliance Standard Rule Violation	Events are generated for compliance standard rule violations. Each event corresponds to a violation of a compliance rule on a specific target.
Compliance Standard Score Violation	Events are generated for compliance standard score violations. An event is generated when the compliance score for a compliance standard on a specific target falls below predefined thresholds.
High Availability	High Availability events are generated for database availability operations (shutdown and startup), database backups and Data Guard operations (switchover, failover, and other state changes).
Job Status Change	All changes to the status of an Enterprise Manager job are treated as events, and these events are made available via the Job Status Change event class. A prerequisite to creating Incident Rules, is to enable the relevant job status and add required targets to job event generation criteria. To change this criteria, visit Setup->Incidents->Job Events.
JVM Diagnostics Threshold Violation	A JVMD Diagnostics (JVMD) event is raised when a JVMD metric exceeds its threshold value on a Java Virtual Machine target.
Metric Alert	A metric alert event is generated when an alert occurs for a metric on a specific target (Example: CPU utilization for a host target) or metric on a target and object combination (Example: space usage on a specific tablespace of a database target)
Metric Evaluation Error	A metric evaluation error is generated when the collection for a specific metric group fails for a target.
Service Level Agreement Alert	These events are generated when an alert occurs for a Service Level Agreement or a Service Level Objective.

Event Type	Description
Target Availability	The Target Availability Event represents a target's availability status (Example: Up, Down, Agent Unreachable, or Blackout).
User-reported	These events are created by end-users.

Incidents allow you to manage many discrete event types by providing an intuitive way to combine them into meaningful issues that you can act upon.

Event Severity

Another event attribute is severity. Just as previous versions of Enterprise Manager utilized metric alert severity levels, this concept has been extended to all event types. The following table shows the various event severity levels along with the associated icon.

Icon	Severity	Description
	Fatal	The monitored target is down (target down event). A Fatal severity is the highest level severity and only applies to the Target Availability event type.
	Critical	Immediate action is required in a particular area. The area is either not functional or indicative of imminent problems.
	Warning	Attention is required in a particular area, but the area is still functional.
	Advisory	While the particular area does not require immediate attention, caution is recommended regarding the area's current state. This severity is used primarily for compliance standards.
	Informational	A specific condition has just occurred.

3.3 Incidents

You monitor and manage your Enterprise Manager environment via incidents and not discrete events (even though an incident can conceivably consist of a single event). Managing by incident means rather than managing discrete events for your system. You now manage an incident that may consist of one significant event (for example, a target down event) or combination of related events (for example, host CPU utilization, host memory utilization, and host swap utilization events when monitoring host capacity). Incidents add an intuitive layer of abstraction that allows you to manage your monitored systems more efficiently because there is a smaller set of more meaningful incidents to manage.

When an incident is created, Enterprise Manager makes available rich set of incident management workflow features that let you to manage and track the incident through its complete lifecycle. Incident management functions allow you to:

- Assign incident ownership.
- Track the incident resolution status.
- Set incident priority.
- Set incident escalation level.

- Access (in context) My Oracle Support knowledge base articles and other Oracle documentation to help resolve the incident.
- Access direct in-context diagnostic/action links to relevant Enterprise Manager functionality allowing you to quickly diagnose or resolve the incident.

All incident management/tracking operations are carried out from incident Manager. Creation of incidents for events, assignment of incidents to administrators, setting priority, sending notifications and other actions can be automated using (incident) rules..

The following examples illustrate how incidents are constructed and how attributes map to various stages of the incident lifecycle.

3.3.1 Incident Composed of a Single Event

The simplest incident is composed of a single event. In the following example, you are concerned whenever any production target is down. You can create an incident for the target down event which is raised by Enterprise Manager if it detects the monitored target is down. Once the incident is created, you will have available all incident management functionality required to track and manage its resolution.

Figure 3–1 Incident with a Single Event

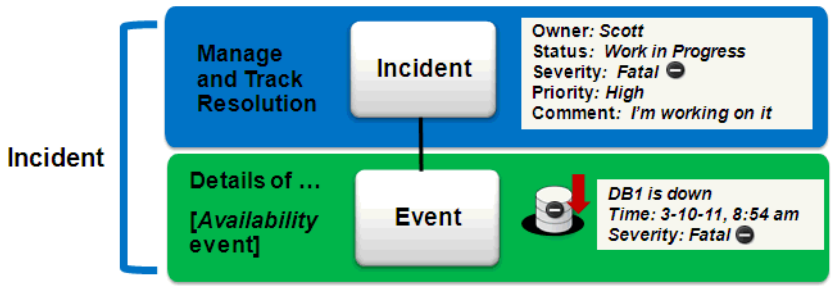

The figure shows how both the incident and event attributes are used to help you manage the incident. From the figure, we see that the database DB1 has gone down and an event of Fatal severity has been raised. An incident is opened and the owner/administrator Scott is currently working to resolve the issue. The incident severity is currently Fatal as the incident inherits the worst severity of all the events within incident. In this case there is only one event associated with the incident so the severity is Fatal.

As an open incident, you can use Incident Manager to track its ownership, its resolution status, set the priority and, if necessary, add annotations to the incident to share information with others when working in a collaborative environment. In addition, you have direct access to pertinent information from MOS and links to other areas of Enterprise Manager that will help you resolve the database problem. By drilling down on an open incident, you can access this information and modify it accordingly, as shown in the following graphic.

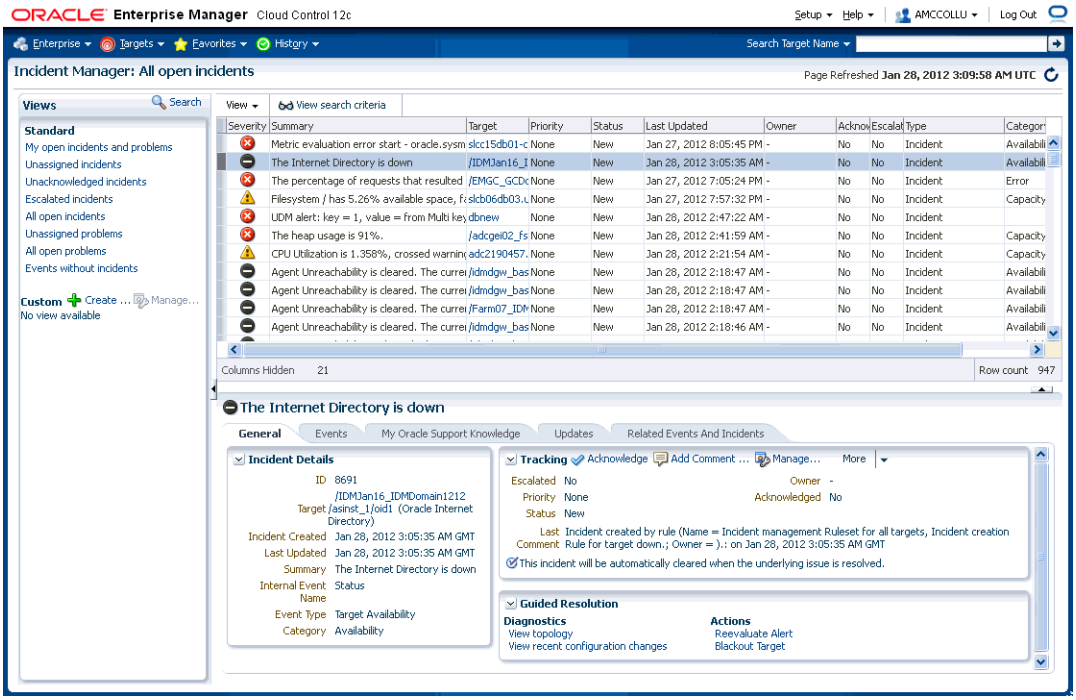

3.3.2 Incident Composed of Multiple Events

Situations of interest may involve more than a single event. It is an incident's ability to contain multiple events that allows you to monitor and manage complex and more meaningful issues. For example, if a monitored system is running out of space, separate multiple events such as *tablespace full* and *filesystem full* may be raised. Both, however, are related to running out of space. Another machine resource monitoring example might be the simultaneous raising of CPU utilization, memory utilization, and swap utilization events. Together, these events form an incident indicating extreme load is being placed on a monitored host. The following figure illustrates this example.

Figure 3–2 Incident with Multiple Events

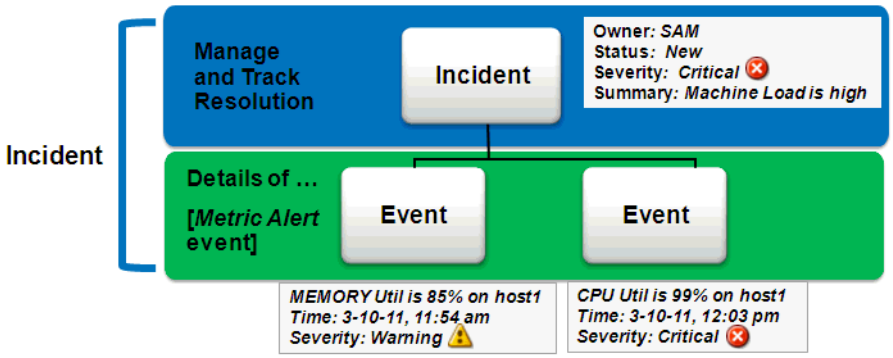

The incident severity is Critical even though one of the events (Memory Utilization) is only at a Warning severity level. Incidents inherit the worst severity of all the events within incident. The incident summary indicates why this incident should be of

interest, in this case, "Machine Load is high". This message is an intuitive indicator for all administrators looking at this incident. By default, the incident summary is pulled from the message of the event, however, this message can be changed by any administrator working on the incident.

Because administrators are interested in overall machine load, administrator Sam has created an incident for these two metric events because they are related—together these events represent a host overload situation. An administrator needs to take action because memory is filling up and consumed CPU resource is too high. In its current state, this condition will impact any applications running on the host.

Helpdesk Incident Resolution

If your IT process requires a helpdesk ticket be created to resolve incidents, then you can use the helpdesk connector to integrate the incident with a helpdesk ticket and have Enterprise Manager automatically open a ticket when the incident is created in addition to tracking the ticket ID, and status of the ticket. This provides administrators with a way to check the status of the ticket from within Incident Manager. Enterprise Manager also allows you to link out to a Web-based third-part console directly from the ticket so that you can launch the console in context directly from the ticket.

3.3.3 Incident Attributes

Every incident possesses attributes that provide information as identification, status for tracking, and ownership. The following table lists available incident attributes.

Incident Attribute	Definition
Escalated	Escalation Levels ■ None (Not escalated) ■ Level 1 ■ Level 2 ■ Level 3 ■ Level 4 ■ Level 5
Priority	Priority Options ■ None ■ Low ■ Medium ■ High ■ Very High ■ Urgent
Status	Incident Status ■ New ■ Work in Progress ■ Resolved
Comment	Annotations added by an administrator to communicate analysis information or actions taken to resolve the incident.
Owner	Administrator/user currently working on the incident..

Incident Attribute	Definition
Acknowledged	Yes or No. Acknowledging an incident stops any repeat notifications for that incident. When an incident is acknowledged, it will be implicitly assigned to the user who acknowledged it. When a user assigns an incident to himself, it is considered 'acknowledged'. Once acknowledged, an incident cannot be unacknowledged .

3.3.4 Event Prioritization

When working in a large enterprise it is conceivable that when systems are under heavy load, an extraordinarily large number of incidents and events will be generated. All of these need to be processed in a timely and efficient manner in accordance with your business priorities. To have them processed sequentially can result in long waits before incidents can be resolved: High priority events/incidents need to be addressed before those of low priority.

In order to determine which event/incidents are high priority, Enterprise Manager uses a prioritization protocol based on two incident/event attributes: Lifecycle Status of the target and the Incident/Event Type. Lifecycle Status is a target property that specifies a target's operational status. You can set/view a target's Lifecycle Status from the UI (from a target's **Target Setup** menu, select **Properties**). You can set target Lifecycle Status properties across multiple targets simultaneously by using the Enterprise Manager Command Line Interface (EM CLI) set_target_propert_value verb.

A target's Lifecycle Status is set when it is added to Enterprise Manager for monitoring. At that time, you determine where in the prioritization hierarchy that target belongs—the highest level being "mission critical" and the lowest being "development."

Target Lifecycle Status

- Mission Critical (highest priority)
- Production
- Stage
- Test
- Development (lowest priority)

Incident/Event Type

- Availability (highest priority)
- All events/incidents (Fatal severity)
- All events/incidents (Warning and Critical severities)
- All events/incidents (Informational) (lowest priority)

3.4 Incident Manager

Incident Manager provides, in one location, the ability to search, view, manage, and resolve incidents and problems impacting your environment. Use Incident Manager to perform the following tasks:

- Filter incidents, problems, and events by using views
- Respond and work on an incident

- Manage incident lifecycle including assigning, acknowledging, tracking its status, prioritization, and escalation

Figure 3–3 incident Manager

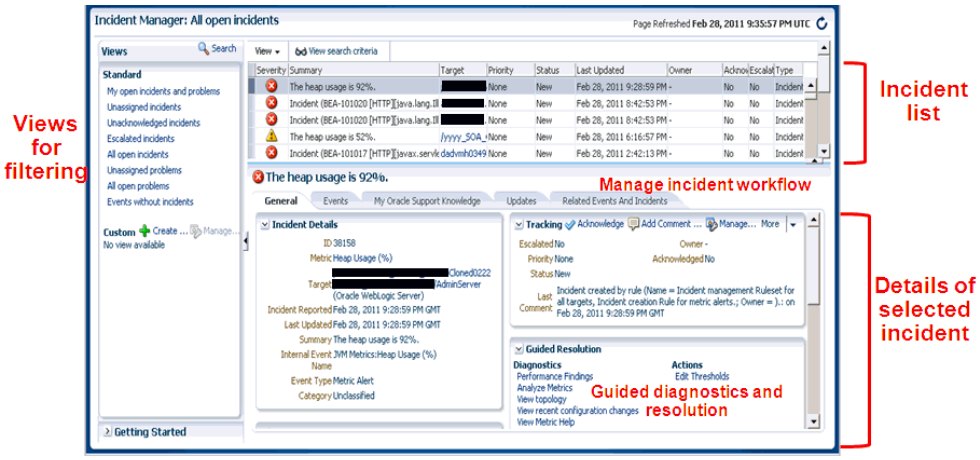

The advantages of using Incident Manager include:

- Ability to manage events, incidents, and problems from a central location.
- Ability to assign incidents to specific personnel (prioritize, escalate, and track incidents through various states of resolution).
- For the person working the incidents and problems:

  See what incidents and problems are assigned to him, acknowledge that he is working on the incidents and problems, and provide information to the user community regarding the progress of the resolution.
- Integration with My Oracle Support.
- In-context diagnostics and resolution links.

3.5 Before Working with Incidents

In order to work with incidents, ensure all relevant Enterprise Manager administrator accounts have been granted the appropriate privileges to manage incidents and ensure that the notification system is properly configured to send notification.

Granting User Privileges for Events, Incidents and Problems

Users are granted privileges for events, incidents, and problems in the following situations:

For events, two privileges are defined:

- The View Event privilege allows you to view an event and add comments to the event.
- The Manage Event privilege allows you to take update actions on an event such as closing a manually-cleared event, creating an incident for an event, and creating a ticket for an event. You can associate an event with an incident.

For incidents, two privileges are defined:

- The View Incident privilege allows you to view an incident, and add comments to the incident.
- The Manage Incident privilege allows you to take update actions on an incident. The update actions supported for an incident includes incident assignment and prioritization, resolution management, manually closing manually-clear events, manually create a problem for an incident, and create a ticket for an incident.

For problems, two privileges are defined:

- The View Problem privilege allows you to view a problem and add comments to the problem.
- The Manage Problem privilege allows you to take update actions on the problem. The update actions supported for a problem include problem assignment and prioritization, resolution management, manually closing the problem, and update customer-defined attribute values. It also includes ability to create a Service Request and gather diagnostics using Support Workbench.

To administer privileges from the Enterprise Manager UI, from the **Setup** menu, select **Security** and then **select Administrators**.

3.6 Working with Incidents

Data centers follow operational practices that enable them to manage events and incidents by business priority and in a collaborative manner. Enterprise Manager provides the following features to enable this management and automation:

- Sending notifications to the appropriate administrators.
- Assigning initial ownership of an incident and perhaps transferring ownership based on shift assignments or expertise.
- Tracking its resolution status.
- Assigning priorities based on the component affected and nature of the incident.
- Escalating incidents in order to meet service level agreements (SLA).
- Accessing My Oracle Support knowledge articles.
- Opening Oracle Service Requests to request assistance with problems with Oracle software.
- Generating management and operational reports to track the status of incidents

You can manage an incident by performing the following:

1. In the **All Open Incidents** view, select the incident.
2. In the resulting Details page, click the **General** tab, then click **Manage**. The **Manage** dialog displays.

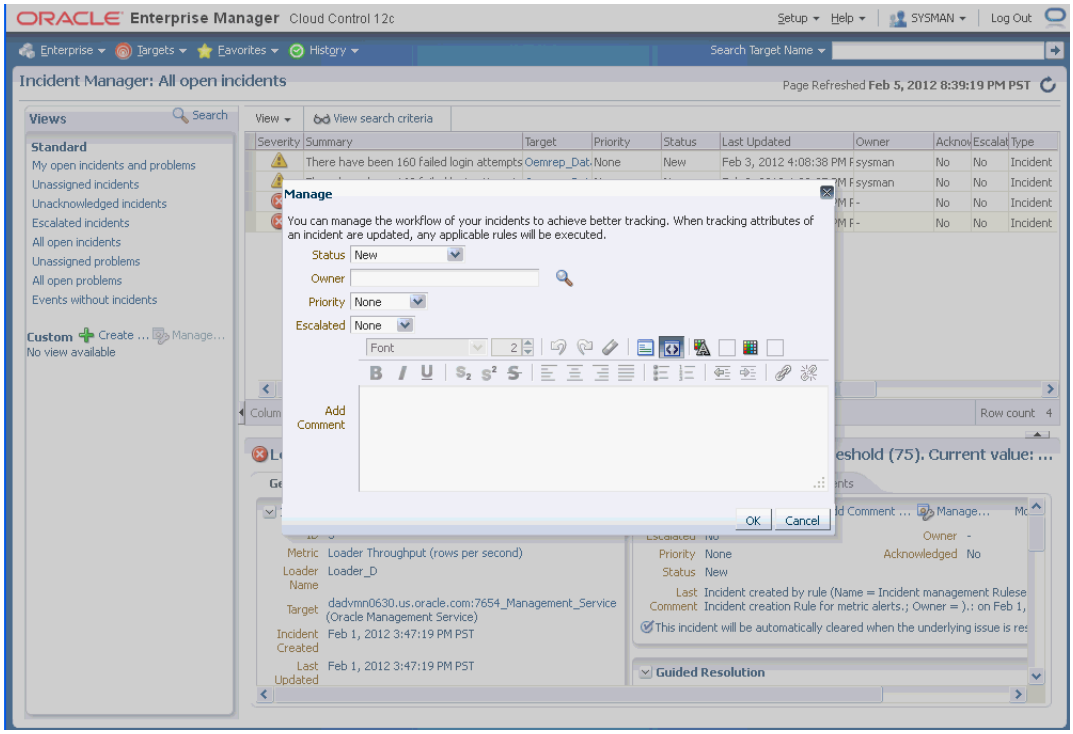

You can then adjust the priority, escalate the incident, and assign it to a specific engineer.

3.6.1 Setting Up Views in Incident Manager

A view is a set of search criteria for filtering incidents and problems in the system. You can define views to help you gain quick access to the incidents and problems on which you need to focus. For example, you may define a view to display all the incidents associated with the production databases that you own.

By specifying preferences to view the following for each of the incidents in the list: incident severity, incident message, acknowledgement flag, date the incident triggered, administrator assigned to it, resolution status, priority, escalated flag, ID, and category, you can filter extraneous incidents. Once the view preference is saved, Enterprise Manager will display only the list of matching incidents.

You can then search the incidents for only the ones with specific attributes, such as priority 1. The view allows easy access to pertinent incident for daily triaging activity. Accordingly, you can save the search criteria as a filter named "All priority 1 incidents for my targets". The filter becomes available in the UI for immediate use and will be available anytime you log in to access the specific incidents.

Note: The view you create is specific to your Enterprise Manager account and cannot currently be shared with other administrators.

Perform the following steps:

1. Navigate to the Incident Manager page.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.

2. In the **Views** region located on the left, click **Search**.

   a. In the **Search** region, search for Incidents using the **Type** list and select **Incidents**.

   b. In the Criteria region, choose all the criteria that are appropriate. To add fields to the criteria, click **Add Fields...** and select the appropriate fields.

   c. After you have provided the appropriate criteria, click **Get Results**.

   d. To view all the columns associated with this table, in the **View** menu, select **Columns**, then select **Show All**.

      Validate that the list of incidents match what you are looking for. If not, change the search criteria as needed.

   e. Click the **Create View...** button.

3.6.2 Responding and Working on a Simple Incident

Before you begin working on resolving an incident, ensure your Enterprise Manager account has been granted the appropriate privileges to manage incidents from your managed system.

- Privileges on events are calculated based on the privilege on the underlying source objects. For example, the user will have VIEW privilege on an event if he can view the target for the event.
- Privileges on an incident are calculated based on the privileges on participating events.
- Similarly, problem privileges are calculated based on privileges on underlying incidents.

Perform the following steps:

1. Navigate to Incident Manager.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.

2. Use a view to filter the list of incidents. For example, the administrator should use My Open Incidents and Problems view to see incidents and problems assigned to him. You can then sort the list by priority which you talk about next..

   To view incidents assigned to you, click on the predefined view **My Open Incidents and Problems**.

   Work on the incident with the highest priority. Be aware that as you are working on an individual incident, new incidents might be coming in. Update the list of incidents by clicking the **Refresh** icon.

3. To work on an incident, select the incident. In the General tab, click **Acknowledge** to acknowledge the incident and set yourself as owner.

4. If the solution for the incident is unknown, use one or all of the following methods made available in the Incident page:

   - Use the **Guided Resolution** region and access any recommendations, diagnostic and resolution links available.
   - Check My Oracle Support Knowledge base for known solutions for the incident.

- Study related incidents available through the Related Events and Incidents tab.

5. Once the solution is known and can be resolved right away, resolve the incident by using tools provided by the system, if possible.
6. In most cases, once the underlying cause has been fixed, the incident is cleared in the next evaluation cycle. However, in cases like log-based incidents, clear the incident.

3.6.3 Searching My Oracle Support Knowledge

To access My Oracle Support Knowledge base entries from within Incident Manager, perform the following steps:

1. Navigate to Incident Manager.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.
2. Select one of the standard views. Choose the appropriate incident or problem in the View table.
3. In the resulting details region, click My Oracle Support Knowledge. Sign in to My Oracle Support.
4. On the My Oracle Support page, click the **Knowledge** tab to browse the knowledge base.

   From this page, in addition to accessing formal Oracle documentation, you can also change the search string in to look for additional knowledge base entries.

3.6.4 Open Service Request

There are times when you may need assistance from Oracle Support to resolve a problem. To submit a service request (SR), perform the following steps:

1. Navigate to Incident Manager.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.
2. Use one of the views to find the problem or search for it or use one of your custom views. Select the appropriate problem from table.
3. Click on the **Support Workbench: Package Diagnostic** link.
4. Complete the workflow for opening an SR. Upon completing the workflow, a draft SR will have been created.
5. Sign in to My Oracle Support if you are not already signed in.
6. On the My Oracle Support page, click the **Service Requests** tab.
7. Click **Create SR** button. Click **Help** to learn how to create a new SR.

3.6.5 Suppressing Incidents and Problems

There are times when it is convenient to hide an incident or problem from the list in the All Open Incidents page or the All Open Problems page. For example, you may want to suppress an incident while the incident is being actively worked on and you do not need to be notified.

To suppress an incident or problem:

1. Navigate to Incident Manager.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.

2. Select either the All Open Incidents view or the All Open Problems view. Choose the appropriate incident or problem. Click the **General** tab.

3. In the resulting details region, click **More**, then select **Suppress**.

4. On the resulting Suppress pop-up, choose the appropriate suppression type. Add a comment if desired.

5. Click **OK**.

3.7 Incidents - Advanced Tasks

You can perform the following advanced tasks using Incident Manager:

- Creating an Incident Manually
- Managing Workload Distribution of Incidents

3.7.1 Creating an Incident Manually

If an event of interest occurs that is not covered by any rule and you want to convert that event to an incident, perform the following:

1. Using an available view, find the event of interest.

2. Select the event in the table.

3. From the **More...** drop-down menu, choose **Create Incident...**

4. Enter the incident details and click **OK**.

5. Should you decide to work on the incident, set yourself as owner of the incident and update status to *Work in Progress*.

Example Scenario

As per the operations policy, the DBA manager has setup rules to create incidents for all critical issues for his databases. The remainder of the issues are triaged at the event level by one of the DBAs.

One of the DBA receives e-mail for an "SQL Response" event (not associated with an incident) on the production database. He accesses the details of the event by clicking on the link in the e-mail. He reviews the details of the event. This is an issue that needs to be tracked and resolved, so he opens an incident to track the resolution of the issue. He marks the status of the incident as "Work in progress".

3.7.2 Managing Workload Distribution of Incidents

Incident Manager enables you to manage incidents and problems to be addressed by your team

Perform the following tasks:

1. Navigate to Incident Manager.

   From the **Enterprise** menu on the Enterprise Manager home page, select **Monitoring**, then select **Incident Manager**.

2. Use the standard or custom views to identify the incidents for which your team is responsible. You may want to focus on unassigned and unacknowledged incidents and problems.
3. Review the list of incidents. This includes: determining person assigned to the incident, checking its status, progress made, and actions taken by the incident owner.
4. Add comments, change priority, reassign the incident as needed by clicking on the Manage button in the Incident Details region.

Example Scenario

The DBA manager uses Incident Manager to view all the incidents owned by his team. He ensures all of them are correctly assigned; if not, he reassigns and prioritizes them appropriately. He monitors the escalated events for their status and progress, adds comments as needed for the owner of the incident. In the console, he can view how long each of the incidents has been open. He also reviews the list of unassigned incidents and assigns them appropriately.

3.8 Rule Sets

With previous versions of Enterprise Manager, you used notification rules to choose the individual targets and conditions for which you want to perform actions and/or receive notifications (send e-mail, page, open a trouble ticket) from Enterprise Manager. For Enterprise Manager 12c, the concept and function of notification rules has been replaced with rules and rule sets.

- **Rules**: A rule instructs Enterprise Manager to take specific actions when incidents, events, or problems occur, such as performing notifications. Beyond notifications, rules can also instruct Enterprise Manager to perform specific actions, such as raising additional incidents.
- **Rule Set**: Consists of one or more rules that operate on the same object such as a group of targets. A rule set is simply a group of rules which together automate a business process.

A rule set is a set of one or more rules that apply to a common set of objects such as targets (hosts, databases, groups), jobs, metric extensions, or self updates. A rule set allows you to logically combine different rules relating to the common set of objects (such as jobs, targets, applications) into a single manageable unit. Operationally, individual rules within a rule set are executed in a specified order as are the rule sets themselves. rule sets are executed in a specified order. By default, the execution order for both rules and rule sets is the order in which they are created, but can be reordered from the Incident Manager UI.

The following figure shows typical rule set structure and how the individual rules are applied to a heterogeneous group of targets.

Figure 3–4 Rule Set Application

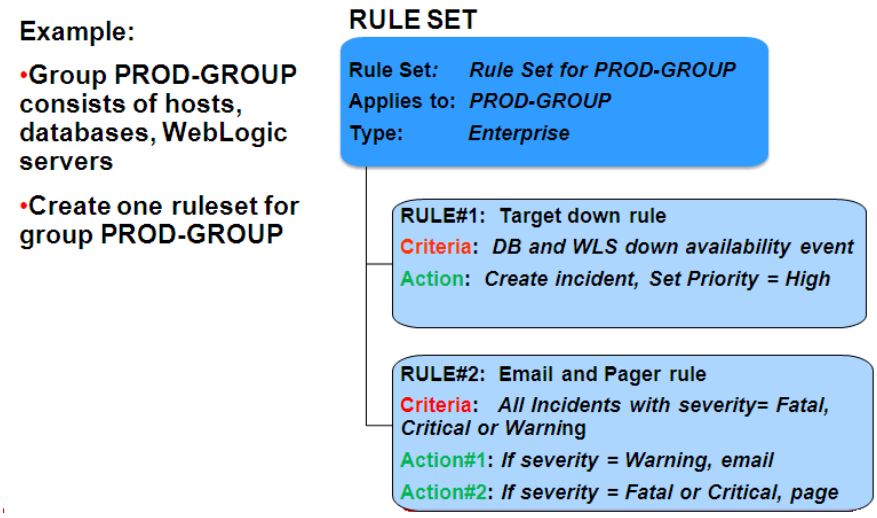

The graphic illustrates a situation where all rules pertaining to a group of targets can be put into a single rule set. In the above example, a group named *PROD-GROUP* consists of hosts, databases, and WebLogic servers exists as part of a company's managed environment. A single rule set is created to manage the group.

In addition to the actual rules contained within a rule set, a rule set possesses the following attributes:

- **Name**: A descriptive name for the rule set.
- **Description**: Brief description stating the purpose of the rule set.
- **Applies To**: Object to which all rules in the rule set apply: Valid rule set objects are targets, jobs, metric extensions, and self update.
- **Owner**: The Enterprise Manager user who created the rule set. Rule set owners have the ability to update or delete the rule set and the rules in the rule set.
- **Enabled**: Whether or not the rule set is actively being applied.
- **Type**: Enterprise or Private.

3.8.1 Out-of-Box Rule Sets

Enterprise Manager provides out-of-box rule sets for incident creation, event deletion based on typical scenarios. The following rule sets are immediately available upon installation.

Incident Management Rule Sets for All Targets

- Incident creation Rule for target down.
- Incident creation Rule for agent unreachable (for Agents and hosts).
- Incident creation Rule for metric alerts (for critical severity only).
- Out-of-box Incident creation rule for Service Level Agreement Alerts.
- Incident creation rule for compliance score violation
- Incident creation rule for high-availability events.
- Auto-clear Rule for metric alerts older than 7 days.

- Auto clear Rule for job status change terminal status events older than 7 days.
- Clear Application Dependency and Performance (ADP) alerts after without incidents after 7 days.

Event Management Rule Set for Self-Update

- Notification Rule for new updates

> **Note:** Out-of-Box rule sets cannot be deleted. They can only be enabled or disabled.

Some examples of the types of actions that a rule can perform are:

- Create an incident based on an event.
- Perform notification actions such as sending an e-mail or generating a helpdesk ticket.
- Perform actions to manage incident workflow notification via e-mail/PL/SQL methods/ SNMP traps. For example, if target down event occurs, create an incident and e-mail administrator Joe about the incident. If the incident is still open after two days, set the escalation level to one and e-mail Joe's manager.

3.8.2 Rule Set Types

There are two types of Rule Sets:

- **Enterprise**: Used to implement all operational practices within your IT organization. All supported actions are available for this type of rule set. However, because this type of rule set can perform all actions, there are restrictions as to who can create an enterprise rule set.

  In order to create or edit an enterprise rule set, an administrator must have been granted the "Create Enterprise Rule Set " privilege on the "Enterprise Rule Set" resource. An enterprise rule set can have multiple authors, however, if the originator of the rule set wants other administrators to edit the rule set, he will need to share access in order to work collaboratively. Rule sets are visible to all administrators.

  Important: When a rule set performs actions, the privileges of the rule set creator are used.

- **Private**: If an administrator does not have the Create Enterprise Rule Set resource privilege and consequently cannot create an enterprise rule set, but wants to be notified about something he is monitoring, he can create a private rule set. The only action a private rule set can perform is to send e-mail to the rule set owner. Any administrator can create a private rule set.

3.8.3 Rules

Rules are instructions within a rule set that automate actions on incoming events or incidents or problems (critical errors in Oracle software). Because rules operate on *incoming* incidents/events/problems, if you create a new rule, it will not act retroactively on incidents/events/problems that have already occurred.

Every rule is composed of two parts:

- **Criteria**: The events/incidents/problems on which rule applies.

- **Action**(s): The ordered set of one or more operations on the specified events, incidents, or problems. Each action can be executed based on additional conditions.

  Important: Rules are executed in a specified order. The rule execution order can be changed at any time. By default, rules are executed in the order they are created.

The following table shows how rule criteria and actions determine rule application. In this rule operation example there are three rules executed in order according to specified criteria.

Table 3–1 Rule Operation

Rule Name	Execution Order	Criteria	Action	
			Condition	**Actions**
Rule 1	First	CPU Util(%), Tablespace Used(%) metric alert events of warning or critical severity		Create incident.
Rule 2	Second	Incidents of warning or critical severity	If severity = critical	Notify by page
			If severity =warning	Notify by e-mail
Rule 3	Third	Incidents open for more than 7 days		Set escalation level to 1

In the rule operation example, *Rule 1* applies to two metric alert events: *CPU Utilization* and *Tablespace Used*. Whenever these events reach either Warning or Critical severity threshold levels, an incident is created.

When the incident severity level (the incident severity is inherited from the worst event severity) reaches Warning, *Rule 2* is applied according to its first condition and Enterprise Manager sends an e-mail to the administrator. If the incident severity level reaches Critical, *Rule 2*'s second condition is applied and Enterprise Manager sends a page to the administrator.

If the incident remains open for more than seven days, *Rule 3* applies and the incident escalation level is increased from None to Level 1.

3.8.3.1 Rule Criteria

Rules are applied to events, incidents, and problems according to criteria selected at the time of rule creation (or update).There are three rule applications:

- Incoming/updated events
- Newly created/updated incidents
- Newly created/updated problems

Available criteria varies depending on the rule application. The following tables list selectable criteria for each type.

Table 3–2 Rule Criteria: Events

Criteria	Description
Type	Rule applies to a specific event type. The following event types are available: ■ Application Dependency and Performance Alert ■ Compliance Standard Rule Violation ■ Compliance Standard Score Violation ■ High Availability ■ JVM Diagnostics Threshold Violation ■ Job Status Change ■ Metric Alert ■ Metric Evaluation Error ■ Service Level Agreement Alert ■ Target Availability ■ User-Reported
Severity	Rule applies to a specific event severity. The following event severities are available: ■ Fatal ■ Critical ■ Warning ■ Advisory ■ Informational ■ Clear
Category	Rule applies to a specific event category. The following event categories are available: ■ Availability ■ Business ■ Capacity ■ Configuration ■ Diagnostics ■ Error ■ Faults ■ Jobs ■ Load ■ Performance ■ Security

Table 3–2 (Cont.) Rule Criteria: Events

Criteria	Description
Target type	Rule applies to a specific target type. The following target types are available: ■ Agent ■ Application Deployment ■ Beacon ■ CSA Collector ■ Database Instance ■ Database System ■ EM Service ■ Host ■ Infrastructure Cloud ■ Metadata Repository ■ OMS Console ■ OMS Repository ■ OMS Platform ■ OMS and Repository ■ Oracle Authorization Policy Manager ■ Oracle Fusion Middleware Farm ■ Oracle HTTP Server ■ Oracle Home ■ Oracle Management Service ■ Oracle Web Logic Domain ■ Oracle Web Logic Server
Associated with incident	Typically, events are associated with incidents through rules. Specify Yes or No.
Event name	Rule applies to events with a specific name. The specified name can either be an exact match or a pattern match.
Root cause analysis result	Upon completion of Root Cause Analysis (RCA) event, the rule applies to the event that is marked either as root cause or symptom. Alternatively, the rule can act on an RCA event when it is no longer a symptom.
Associated incident acknowledged	Rule applies to an event that is associated with a specific incident when that incident is acknowledged by an administrator. Specify Yes or No.
Total occurrence count	For duplicated events, the rule is applies when the total number of event occurrences reaches a specified number.
Comment added	Rule applies to events where an administrator adds a comment.

For incidents, a rule can apply to all new and/or updated incidents, or newly created incidents that match specific criteria shown in the following table.

Table 3–3 Rule Criteria: Incidents

Criteria	Description
Rules that created the incidence	Rule applies to incidents raised by a specific rule.

Table 3–3 (Cont.) Rule Criteria: Incidents

Criteria	Description
Category	Rule applies to a specific incident category. The following incident categories are available: ■ Availability ■ Business ■ Capacity ■ Configuration ■ Diagnostics ■ Error ■ Faults ■ Jobs ■ Load ■ Performance ■ Security
Target Type	Rule applies to a specific target type. The following target types are available: ■ Agent ■ Application Deployment ■ Beacon ■ CSA Collector ■ Database Instance ■ Database System ■ EM Service ■ Host ■ Infrastructure Cloud ■ Metadata Repository ■ OMS Console ■ OMS Repository ■ OMS Platform ■ OMS and Repository ■ Oracle Authorization Policy Manager ■ Oracle Fusion Middleware Farm ■ Oracle HTTP Server ■ Oracle Home ■ Oracle Management Service ■ Oracle Web Logic Domain ■ Oracle Web Logic Server

Table 3–3 (Cont.) Rule Criteria: Incidents

Criteria	Description
Severity	Rule applies to a specific incident severity. The following incident severities are available: ■ Fatal ■ Critical ■ Warning ■ Advisory ■ Informational ■ Clear
Acknowledged	Rule applies if the incident has been acknowledged by an administrator. Specify Yes or No.
Owner	Rule applies for a specified incident owner.
Priority	Rule applies when incident priority matches a selected priority. Available priorities are: ■ Urgent ■ Very High ■ High ■ Medium ■ Low ■ None
Status	Rule applies when the incident status matches a selected incident status. Available statuses: ■ New ■ Work in Progress ■ Resolved ■ Closed
Escalation Level	Rule applies when the incident escalation level matches the selected level. Available escalation levels: None, Level 1, Level 2, Level 3, Level 4, Level 5
Associated with Ticket	Rule applies when the incident is associated with a helpdesk ticket. Specify Yes or No.
Associated with Service Request	Rule applies when the incident is associated with a service request. Specify Yes or No.
Diagnostic Incident	Rule applies when the incident is a diagnostic incident. Specify Yes or No.
Unassigned	Rule applies if the newly raised incident does not have an owner.
Comment Added	Rule applies if an administrator adds a comment to the incident.

For problems, a rule can apply to all new and/or updated problems, or newly created problems that match specific criteria shown in the following table.

Table 3–4 Rule Criteria: Problems

Criteria	Description
Problem key	Each problem has a problem key, which is a text string that describes the problem. It includes an error code (such as ORA 600) and in some cases, one or more error parameters. Rule can apply to a specific problem key or a key matching a specific pattern (using a wildcard character).
Category	Rule applies to a specific problem category. The following problem categories are available: ■ Availability ■ Business ■ Capacity ■ Configuration ■ Diagnostics ■ Error ■ Faults ■ Jobs ■ Load ■ Performance ■ Security
Target Type	Rule applies to a specific target type. The following target types are available: ■ Agent ■ Application Deployment ■ Beacon ■ CSA Collector ■ Database Instance ■ Database System ■ EM Service ■ Host ■ Infrastructure Cloud ■ Metadata Repository ■ OMS Console ■ OMS Platform ■ OMS and Repository ■ Oracle Authorization Policy Manager ■ Oracle Fusion Middleware Farm ■ Oracle HTTP Server ■ Oracle Home ■ Oracle Management Service ■ Oracle WebLogic Domain ■ Oracle WebLogic Server
Acknowledged	Rule applies when the problem is acknowledged.
Owner	Rule applies for a specified problem owner.

Table 3–4 (Cont.) Rule Criteria: Problems

Criteria	Description
Priority	Rule applies when problem priority matches a selected priority. Available priorities are: ■ Urgent ■ Very High ■ High ■ Medium ■ Low ■ None
Status	Rule applies when the problems matches a specific status. The following statuses are available: ■ New ■ Work in Progress ■ Resolved ■ Closed
Escalation Level	Rule applies when the problem escalation level matches the selected level. Available escalation levels: None, Level 1, Level 2, Level 3, Level 4, Level 5
Incident Count	Rule applies when the number of incidents related to the problem reaches the specified count limit. The problem owner and the Operations manager are notified via e-mail.
Associated with Service Request	Rule applies if the incoming problem is has an associated Service Request. Specify Yes or No.
Associated with Bug	Rule applies if the incoming problem is has an associated bug. Specify Yes or No
Unassigned	Rule applies if the newly raised incident does not have an owner.
Comment Added	Rule applies if an administrator adds a comment to the problem.

3.8.3.2 Rule Actions

For each rule condition, Enterprise Manager allows you to define specific actions. The following table summarizes available actions for each to rule application.

Table 3–5 Available Rule Actions

Action	Event	Incident	Problem
E-mail	Yes	Yes	Yes
Page	Yes	Yes	Yes
Advanced Notification Method	Yes	Yes	Yes
Create an Incident	Yes	No	No
Update Incident/Problem Attributes	No	Yes	Yes
Create a Helpdesk Ticket	Yes	Yes	Yes

3.8.4 Rule Set Guidelines

When creating rule sets, adhering to the following guideline will result in efficient use of system resource as well as operational efficiency.

- For rule sets that operate on targets (for example, hosts and databases), use groups to consolidate all targets into a single target for the rule set.
- Consolidate all rules that apply to the group members within the same rule set and make the group the target of the rule set.
- Leverage the execution order of rules within the rule set.

When creating a new rule, you are given a choice as to what object the rule will apply— events, incidents or problems. Use the following rule usage guidelines to help guide your selection.

Table 3–6 Rule Usage Guidelines

Rule Usage	Application
Rules on Events…	To create incidents for the alerts/events managed in Enterprise Manager
	To create tickets for incidents managed by helpdesk analysts , you want to create an incident for an event, then create a ticket for the incident.
	Send events to third party management systems
	To send notifications on events (no incident created)
Rules on Incidents	Automate management of incident workflow operations (assign owner, set priority, escalation levels..) and send notifications
	Create tickets based on incident conditions. For example, create a ticket if the incident is escalated to level 2.
Rules on Problems	Automate management of problem workflow operations (assign owner, set priority, escalation levels..) and send notifications

Rule Set Example

The following example illustrates many of the implementation guidelines just discussed. All targets have been consolidated into a single group, all rules that apply to group members are part of the same rule set, and the execution order of the rules has been set. In this example, the rule set applies to a group (Production Group G) that consists of the following targets:

- DB1 (database)
- Host1 (host)
- WLS1 (WebLogic Server)

All rules in the rule set perform three types of actions: incident creation, notification, and escalation.

- Example *Rule Set*:
  - Rule set applies to target: Group target G
  - Rules in the *rule set*:
    1. Rule(s) to create incidents for specified events
    2. Rule(s) that send notifications on the incidents
    3. Rule(s) that escalate incidents based on some condition (e.g. length of time incident is open)

In a more detailed view of the rule set, we can see how the guidelines have been followed.

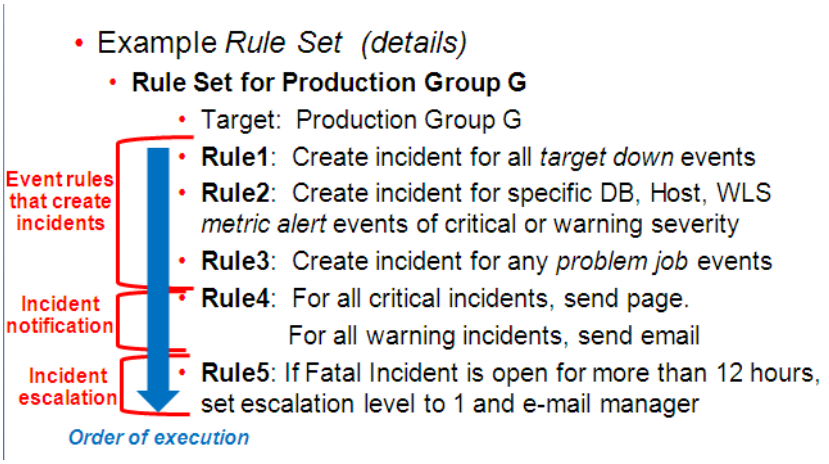

In this detailed view, there are five rules that apply to all group members. The execution sequence of the rules (rule 1 - rule 5) has been leveraged to correspond to the three types of rule actions in the rule set: Rules 1-3

- Rules 1-3: Incident Creation
- Rule 4: Notification
- Rule 5: Escalation

By synchronizing rule execution order with the progression of rule action categories, maximum efficiency is achieved. As shown in this example, by consolidating all notifications in one rule, it is easier to make notification changes in the future when the notifications operations are defined in one place than in multiple places. Note: This assumes that the notification requirements for all the incidents (from rules 1 - 3) are the same.

The following table illustrates explicit rule set operation for this example.

Table 3–7 Example Rule Set for Production Group G

Rule Name	Execution Order	Criteria	Action	
			Condition	Actions
Rule Set: Targets within Production Group G				

Table 3–7 (Cont.) Example Rule Set for Production Group G

Rule Name	Execution Order	Criteria	Action	
			Condition	Actions
Rule 1	First	DB1 goes down . Host1 goes down. WLS1 goes down.		Create incident.
Rule 2	Second	**DB1** Tablespace Full (%) Warning=85, Critical=97 **Host1** CPU Utilization (%): Warning=65, Critical=85 **WLS1** Heap Usage (%) Warning=80, Critical=90	If severity=Warning If severity=Critical	Create incident.
Rule 3	Third	Event generated for problem job status changes for DB1, Host1, and WLS1.		Create incident.
Rule 4	Fourth	All incidents for Production Group G	Severity=Warning Severity=Critical	Send e-mail Send page
Rule 5	Fifth	Incident remains open for more than 12 days.	Status=Fatal	Increase escalation level to 1.

3.9 Before Using Rules

Before you use rules, ensure the following prerequisites have been set up:

- User's Enterprise Manager account has notification preferences (e-mail and schedule).
- If you decide to use connectors, tickets, or advanced notifications, you need to configure them before using them in the actions page.
- Ensure that the SMTP gateway has been properly configured to send e-mail notifications.
- User's Enterprise Manager account has been granted the appropriate privileges to manage incidents from his managed system.

  To perform user account administration, click **Setup** on the Enterprise Manager home page, select **Security**, then select **Administrators** to access the Administrators page.

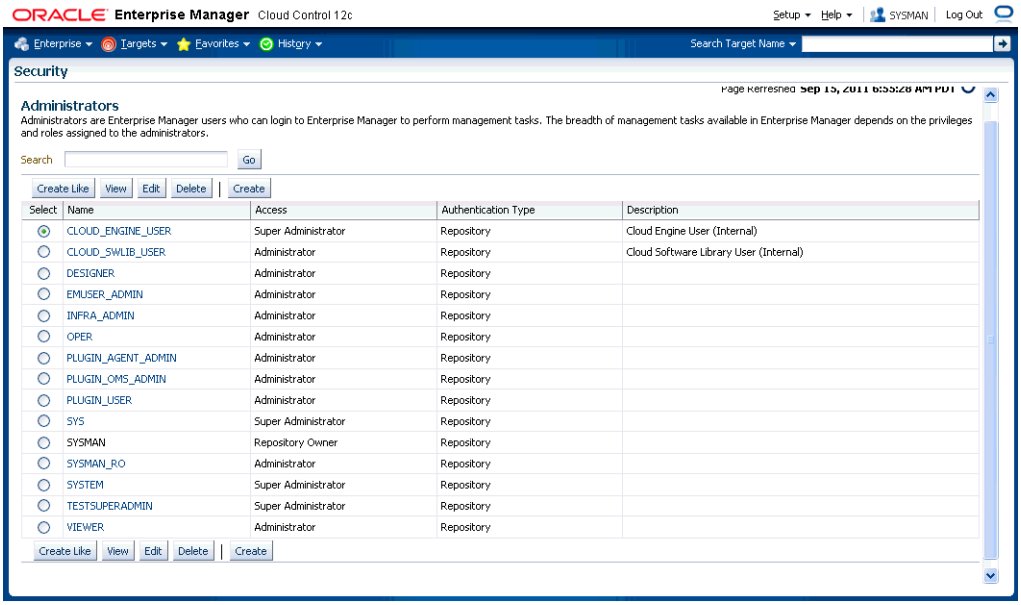

Privileges Required for Enterprise Rule Sets

As the owner of the rule set, an administrator can perform the following:

- Update or delete the rule set, and add, modify, or delete the rules in the rule set.
- Assign co-authors of the rule set. Co-authors can edit the rule set the same as the author. However, they cannot delete rule sets nor can they add additional co-authors.
- When a rule action is to update an event, incident, or problem (for example, change priority or clear an event), the action succeeds only if the owner has the privilege to take that action on the respective event, incident, or problem.
- Additionally, user must be granted privilege to create an enterprise rule set.

If an incident or problem rule has an update action (for example, change priority), it will take the action only if the owner of the respective rule set has manage privilege on the matching incident or problem.

To acquire privileges, click **Setup** on the Enterprise Manager home page, select **Security**, then select **Administrators** to access the Administrators page. Select an administrator from the list, then click **Edit** to access the Administrator properties wizard as shown in the following graphic.

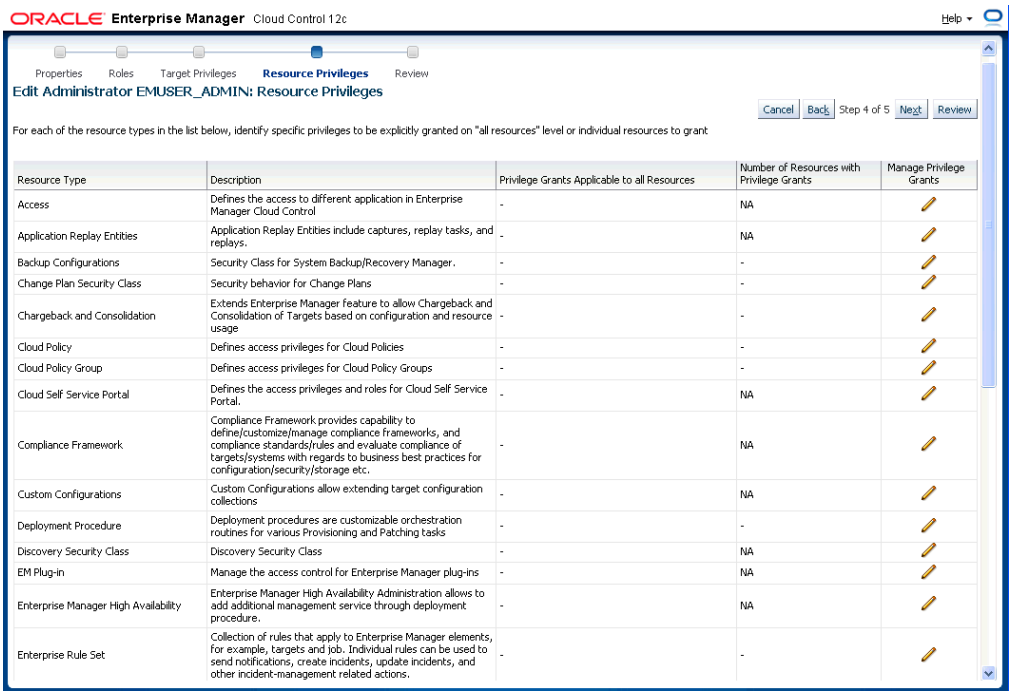

3.10 Working with Rules

You can perform the following tasks using Rules:

- Creating an Rule
- Creating a Rule to Create an Incident
- Creating a Rule to Manage Escalation of Incidents
- Creating a Rule to Escalate a Problem
- Setting Up Automated Notification for Private Rule
- Creating a Rule to Receive Notification Regarding Incidents

3.10.1 Creating an Rule

To create an rule, perform the following steps:

1. From the **Setup** menu, select **Incidents** then select **Rules**.
2. On the Incident Rules - All Enterprise Rules page, edit the existing rule set (highlight the rule set and click **Edit...**) or create a new rule set. Rules are created in the context of a rule set.
3. In the Rules tab of the Edit Rule Set page, click **Create...** and select the type of rule to create (Event, Incident, Problem) on the Select Type of Rule to Create page. Click **Continue**.
4. In the Create New Rule wizard, provide the required information.
5. Once you have finished defining the rule, click **Continue** to add the rule to the rule set. Click **Save** to save the changes made to the rule set.

3.10.2 Creating a Rule to Create an Incident

To create a rule that creates an incident, perform the following steps:

1. From the **Setup** menu, select **Incidents**, then select **Rules**.
2. Determine whether there is an existing rule set that contains a rule that manages the event. In the Rules page, use the Search option to find the events for the target and the associated rule set.

   Note: In the case where there is no existing rule set, create a rule set by clicking **Create Rule Set...** You then create the rule as part of creating the rule set.
3. Select the rule set that will contain the new rule. Click **Edit...** In the Rules tab of the Edit Rule Set page,

   1. Click **Create ...**
   2. Select "Incoming events and updates to events"
   3. Click **Continue**.

   Provide the rule details using the Create New Rule wizard.

   a. Select the Event Type the rule will apply to, for example, Metric Alert. (Metric Alert is available for rule sets of the type Targets.) You can then specify metric alerts by selecting **Specific Metrics**. The table for selecting metric alerts displays. Click the **+Add** button to launch the metric selector. On the Select Specific Metric Alert page, select the target type, for example, Database Instance. A list of relevant metrics display. Select the ones in which you are interested. Click **OK**.

      You also have the option to select the severity and corrective action status.

   b. Once you have provided the initial information, click **Next**. Click **+Add** to add the actions to occur when the event is triggered. One of the actions is to **Create Incident**.

      As part of creating an incident, you can assign the incident to a particular user, set the priority, and create a ticket. Once you have added all the conditional actions, click **Continue**.

   c. After you have provided all the information on the Add Actions page, click **Next** to specify the name and description for the rule. Once on the Review page, verify that all the information is correct. Click **Back** to make corrections; click **Continue** to return to the Edit (Create) Rule Set page.

   d. Click **Save** to ensure that the changes to the rule set and rules are saved to the database.
4. Test the rule by generating a metric alert event on the metrics chosen in the previous steps.

3.10.3 Creating a Rule to Manage Escalation of Incidents

To create a rule to manage incident escalation, perform the following steps:

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. Determine whether there is an existing rule set that contains a rule that manages the incident. You can add it to any of your existing rule sets on incidents.

   Note: In the case where there is no existing rule set, create a rule set by clicking **Create Rule Set...** You then create the rule as part of creating the rule set.

3. Select the rule set that will contain the new rule. Click **Edit...** in the Rules tab of the Edit Rule Set page, and then:
   1. Click **Create ...**
   2. Select "Newly created incidents or updates to incidents"
   3. Click **Continue**.
4. For demonstration purposes, the escalation is in regards to a production database.

   As per the organization's policy, the DBA manager is notified for escalation level 1 incidents where a fatal incident is open for 48 hours. Similarly, the DBA director is paged if the incident has been escalated to level 2, the severity is fatal and it has been open for 72. The operations VP is paged for fatal incidents open for more than 72 hours when the incident has been escalated to level 3.

   Provide the rule details using the Create New Rule wizard.

   a. Select **Specific Incidents** where the rule applies to all newly created incidents or incidents where severity=fatal.

   b. In the Conditions for Actions region located on the Add Actions page, select **Execute the actions on the conditions specified**.

      Select **How long the incident is open and in a particular state (select time and optional expressions)**

      Select the Time to be 48 hours and the Attribute Name to be **Escalation Level** with a value of 1. Click **Continue**.

   c. In the Basic Notification region, type the name of the administrator to be notified by e-mail or page.

   d. Repeat steps b and c to page the DBA director (escalation level=2, severity=fatal, and open for 72 hours). Page Operations VP (escalation level=3, severity=fatal, and open for 72 hours).

      you have to specify different duration condition. We are also missing the action to set the escalation level.

   e. Review the summary and save the rule.

   f. Click **Next** until you get to the Summary screen. Verify that the information is correct and click **Save**.
5. Review the sequence of existing enterprise rules and position the newly created rule in the sequence.

   On the Edit Rule Set page, select **Actions**, then select **Reorder Rules**. Click **Save** to save the change to the sequence.

Example Scenario

In many companies, the operations team handles incidents at different escalation levels. An incident is escalated to a higher level based on how long the incident remains unresolved.

To facilitate this process, the administration manager creates a rule to escalate unresolved incidents based on their age:

- To level 1 if the incident is open for 30 minutes
- To level 2 if the incident is open for 1 hour
- To level 3 if the incident is open for 90 minutes

As per the organization's policy, the DBA manager is notified for escalation level 1. Similarly, the DBA director and operations VP are paged for incidents escalated to levels "2" and "3" respectively.

Accordingly, the administration manager inputs the above logic and the respective Enterprise Manager administrator IDs in a separate rule to achieve the above notification requirement. Enterprise Manager administrator IDs represents the respective users with required target privileges and notification preferences (that is, e-mail addresses and schedule).

3.10.4 Creating a Rule to Escalate a Problem

To create a rule to escalate a problem, perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu, select **Incidents**, then select **Incident Rules**.

2. On the Incident Rules - All Enterprise Rules page, either create a new rule set (click **Create Rule Set...**) or edit an existing rule set (highlight the rule set and click **Edit...**). (Rules are created in the context of a rule set.)

3. In the Rules section of the Edit Rule Set page, select **Create...** to create an enterprise rule to automate actions on the problem. Select Problem Rule on the **Select Type of Rule to Create** page. Click **Continue**.

4. On the Create New Rule page, select **Specific problems** and add the following criteria:

   The Attribute Name is **Incident Count**, the Operator is **Greater than or equals** and the Values is **20**. Click **Next**.

5. In the Conditions for Actions region on the Add Actions page select **Always execute the action**. As the actions to take when the rule matches the condition:

   - In the Notifications region, send e-mail to the owner of the problem and to the Operations Manager.
   - In the Update Problem region, enter the e-mail address of the appropriate administrator in the **Assign to** field.

   Click **Continue**.

6. Review the rules summary. Make corrections as needed. Click **Save**.

Example Scenario

In an organization, whenever an unresolved problem has more than 20 occurrences of associated incidents, the problem should be auto-assigned to the appropriate administrator based on target type of the target on which the problem has been raised.

Accordingly, a problem rule is created to observe the count of incidents attached to the problem and notify the appropriate administrator handling that specific target type.

The problem owner and the Operations manager are notified by way of e-mail.

3.10.5 Setting Up Automated Notification for Private Rule

A DBA has setup a backup job on the database that he is administering. As part of the job, the DBA has subscribed to e-mail notification for "completed" job status. Before you create the rule, ensure that the DBA has the requisite privileges to create jobs.

Perform the following steps:

1. Navigate to the Rules page.

   From the **Setup** menu, select **Incidents**, then select **Rules**.

2. On the Incident Rules - All Enterprise Rules page, either edit an existing rule set (highlight the rule set and click **Edit...**) or create a new rule set.

   Note: The rule set must be defined as a Private rule set.

3. In the Rules tab of the Edit Rule Set page, select **Create...** and select **Event Rule**. Click **Continue**.

4. On the Select Events page, select **Job Status Change** as the Event Type. Select the job in which you are interested either by selecting a specific job or selecting a job by providing a pattern, for example, Backup Management.

   Add additional criteria by adding an attribute: Target Type as Database Instance.

5. Add conditional actions: Event matches the following criteria (Severity is Informational) and E-mail Me for notifications.

6. Review the rules summary. Make corrections as needed. Click **Save**.

7. Create a database backup job and subscribe for e-mail notification when the job completes.

When the job completes, Enterprise Manager publishes the informational event for "Job Complete" state of the job. The newly created rule matches the rule and e-mail is sent out to the DBA.

The DBA receives the e-mail and clicks the link to access the details section in Enterprise Manager console for the event.

3.10.6 Creating a Rule to Receive Notification Regarding Incidents

To create a rule to receive notification on incidents, perform the following steps:

1. Navigate to Incident Rules page.

   From the **Setup menu**, select **Incidents**, then select **Incident Rules**.

2. Edit an existing enterprise rule set.

   Highlight the rule set and click **Edit...**

3. In the Rules section of the Edit Rules Set page:

   1. Click **Create ...**

   2. Select "Newly created incidents or updates to incidents"

   3. Click **Continue**.

   Select the type of incidents to which the rule should apply (**All new incidents and updated incidents**, **All new incidents**, or **Specific incidents**) and click **Next**. .

4. To be notified of the incident, define additional actions using the **Add Actions** page. For the conditions under which the incident occurs, select Always Execute the Actions. For the notifications, provide information in the Basic Notifications region.

5. When you receive the e-mail regarding the incident, click on the link to access the details section in Enterprise Manager console for the incident.

3.11 Rules - Advanced Tasks

You can perform the following advanced tasks using Rules:

- Setting Up a Rule to Send Different Notifications for Different Severity States of an Event
- Creating a Rule to Create a Ticket for Incidents
- Creating a Rule to Notify Different Administrators Based on the Event Type
- Creating Notification Subscription to Existing Enterprise Rules
- Manually Ensuring That There Are No Events That Should Be Incidents

3.11.1 Setting Up a Rule to Send Different Notifications for Different Severity States of an Event

Before you perform this task, ensure the DBA has set appropriate thresholds for the metric so that a critical metric alert is generated as expected.

Consider the following example:

The Administration Manager sets up a rule to page the specific DBA when a critical metric alert event occurs for a database in a production database group and to e-mail the DBA when a warning metric alert event occurs for the same targets. This task occurs when a new group of databases is deployed and DBAs request to create appropriate rules to manage such databases.

Perform the following tasks to set appropriate thresholds:

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, highlight a rule set and click **Edit...**. (Rules are created in the context of a rule set. If there is no existing rule set to manage the newly added target, create a rule set.)
3. In the Edit Rule Set page, locate the Rules section. From the **Actions** menu, select **Add Event Rule**.
4. Provide the rule details as follows:
   a. For Type, select **Metric Alerts** as the Type.
   b. In the criteria section, select **Severity**. From the drop-down list, check and **Critical** and **Warning** as the selected values. Click **Next**.
   c. On the Add Actions page, click **+Add**.

      In the Create Incident section, check the **Create Incident** option. Click **Continue**. The Add Action page displays with the new rule. Click **Next**.
   d. Specify a name for the rule and a description. Click **Next**.
   e. On the Review page, ensure your settings are correct and click **Continue**. A message appears informing you that the rule has been successfully created. Click **OK** to dismiss the message.

      Next, you need to create a rule to perform the notification actions.
5. From the Rules section on the Edit Rules page, click **Create**.
6. Select **Newly created incidents or updates to incidents** as the rule type and click **Continue**.
7. Check **Specific Incidents.**

8. Check **Severity** and from the drop-down option selector, check **Critical** and **Warning**. Click **Next**.
9. On the Add Actions page, click **Add**. The Conditional Actions page displays.
10. In the **Conditions for actions** section, choose **Only execute the actions if specified conditions match**.
11. From the **Incident matches the following criteria** list, choose **Severity** and then **Critical** from the drop-down option selector.
12. In the **Notifications** section, enter the DBA in the **Page** field. Click **Continue**. The Add Actions page displays.
13. Click **Add** to create a new action for the Warning severity.
14. In the **Conditions for actions** section, choose **Only execute the actions if specified conditions match**.
15. From the **Incident matches the following criteria** list, choose **Severity** and then **Warning** from the drop-down option selector.
16. In the **Notifications** section, enter the DBA in the **E-mail to** field. Click **Continue**. The Add Actions page displays with the two conditional actions. Click **Next**.
17. Specify a rule name and description. Click **Next**.
18. On the Review page, ensure your rules have been defined correctly and click **Continue**. The Edit Rule Set page displays.
19. Click **Save** to save your newly defined rules.

3.11.2 Creating a Rule to Create a Ticket for Incidents

According to the operations policy of an organization, all critical incidents from a production database should be tracked by way of Remedy tickets. A rule is created to invoke the Remedy ticket connector to generate a ticket when a critical incident occurs for the database. When such an incident occurs, the ticket is generated by the rule, the incident is associated with the ticket, and the operation is logged for future reference to the updates of the incident. While viewing the details of the incident, the DBA can view the ticket ID and, using the attached URL link, access the Remedy to get the details about the ticket.

Before you perform this task, ensure the following prerequisites are met:

- Monitoring support has been set up.
- Remedy ticketing connector has been configured.

Perform the following steps:

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, select the appropriate rule set and click **Edit...**. (Rules are created in the context of a rule set. If there is no applicable rule set , create a new rule set.)
3. Select the appropriate rule that covers the incident conditions for which tickets should be generated and click **Edit...**.
    a. Specify that a ticket should be generated for incidents covered by the rule.
    b. Specify the ticket template to be used.
4. Repeat step 3 until all appropriate rules have been edited.

5. Click **Save**.

3.11.3 Creating a Rule to Notify Different Administrators Based on the Event Type

As per operations policy for production databases, the incidents that relate to application issues should go to the application DBAs and the incidents that relate to system parameters should go to the system DBAs. Accordingly, the respective incidents will be assigned to the appropriate DBAs and they should be notified by way of e-mail.

Before you set up rules, ensure the following prerequisites are met:

- DBA has setup appropriate thresholds for the metric so that critical metric alert is generated as expected.
- Rule has been setup to create incident for all such events.
- Respective notification setup is complete, for example, global SMTP gateway, e-mail address, and schedule for individual DBAs.

Perform the following steps:

1. Navigate to the Incident Rules page.

   From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, highlight a rule set and click **Edit...**. (Rules are created in the context of a rule set. If there is no existing rule set , create a rule set.)
3. Search the list of enterprise rules matching the events from the production database.
4. Select the rule which creates the incidents for the metric alert events for the database. Click **Edit**.
5. Enter the specific metrics that identify *application* issues, as condition to match the incidents.
6. Enter the specific metrics, which identifies issues with *system parameters*, as condition to match the incidents.
7. Type a summary message, for example: Assign the incident to Cindy (Enterprise Manager administrator handling the system parameter issues). For the action, select to e-mail her.
8. Review the rules summary and make corrections as needed. Click **Save**.

3.11.4 Creating Notification Subscription to Existing Enterprise Rules

A DBA is aware that incidents owned by him will be escalated when not resolved in 48 hours. The DBA wants to be notified when the rule escalates the Incident. The DBA can subscribe to the Rule, which escalates the Incident and will be notified whenever the rule escalates the Incident.

Before you set up a notification subscription, ensure the following prerequisites are met:

- There exists an open incident for a database.
- There exists a rule that escalates High Priority Incidents for databases that have not been resolved in 48 hours.

Perform the following steps:

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, click on the rule set containing incident escalation rule in question and click **Edit...** (Rules are created in the context of a rule set. If there is no existing rule set , create a rule set.)
3. In the Rules section of the Edit Rule Set page, highlight the escalation rule and click **Edit...**.
4. Navigate to the Add Actions page.
5. Select the action that escalates the incident and click **Edit...**
6. In the Notifications section, add the DBA to the **E-mail cc** list.
7. Click **Continue** and then navigate back to the **Edit Rule Set** page and click **Save**.

As a result of the edit to the enterprise rule, when an incident stays unresolved for 48 hours, the rule marks it to escalation level 1. An e-mail is sent out to the DBA notifying him about the escalation of the incident.

3.11.5 Manually Ensuring That There Are No Events That Should Be Incidents

Oracle recommends managing via incidents in order to focus on important events or groups of related events. Due to the variety and sheer number of events that can be generated, it is possible that not all important events will be covered by incidents. To help you find these important yet untreated events, Enterprise Manager provides the **Events without incidents** predefined view.

Perform the following steps:

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. In the Views region, click **Events without incidents**.
3. Select the desired event in the table. The event details display.
4. In the details area, choose **More** and then either **Create Incident** or **Add Event to Incident**.

Example Scenario

During the initial phase of Enterprise Manager uptake, every day the DBA manager reviews the un-acknowledged events on the databases his team is responsible for and filters them to view only the ones which are not tracked by ticket or incident. He browses such events to ensure that none of them requires incidents to track the issue. If he feels that one such event requires an incident to track the issue, he creates an incident directly for this event.

If there are certain events he triages and feels nobody else has to follow-up on the event, he marks it as acknowledged. Enterprise Manager filters out events from the Incident Manager that have been acknowledged.

3.12 Problems

For Enterprise Manager 12*c*, problems focus on the diagnostic incidents and problem diagnostic incidents/problems generated by Advanced Diagnostic Repository (ADR), which are automatically raised by Oracle software when it encounters critical errors in the software. A problem, therefore, represents the root cause of all the Oracle software incidents. For these diagnostic incidents, in order to address root cause, a problem object in Enterprise Manager is created that represents the root cause of these diagnostic incidents. A problem is identified by a *problem key* which uniquely identifies

the particular error in software. Each occurrence of this error results in a diagnostic incident which is then associated with the problem object.

When a problem is raised for Oracle software, Oracle has determined that the recommended recourse is to open a service request (SR), send support the diagnostic logs, and eventually provide a solution from Oracle. As an incident, Enterprise Manager makes available all tracking, diagnostic, and reporting functions for problem management. Whenever you view all open incidents and problems, whether you are using Incident Manager, or in context of a target/group home page, you can easily determine what issues are actually affecting your monitored target.

To manage problems, you should use Support Workbench to open the SR. Access to Support Workbench functionality is available through Incident Manager (**Guided Resolution** area) in context of the problem.

The following figure shows the tracking and diagnostic functionality available for problems from Incident Manger.

Figure 3–5 Viewing Problems from Incident Manager

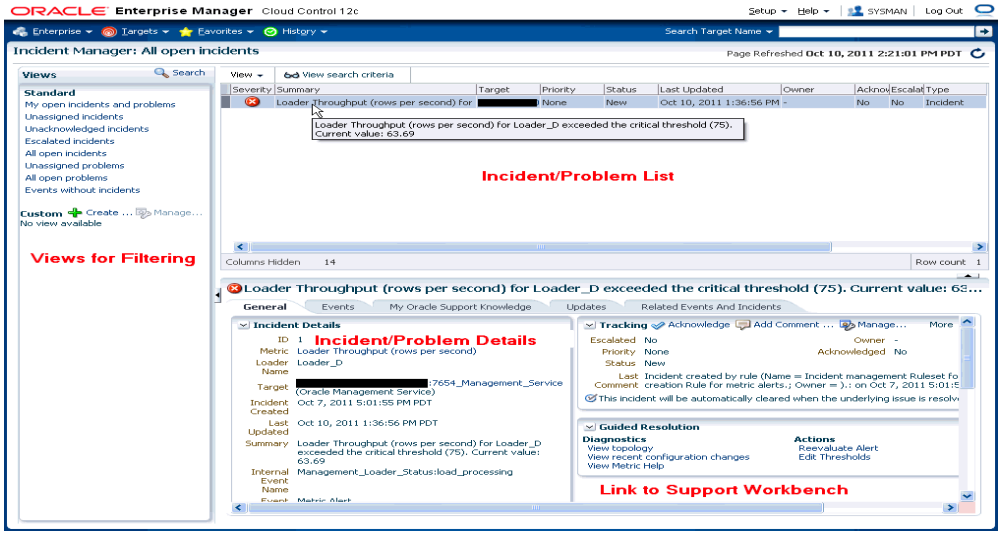

3.13 Moving from Enterprise Manager 10/11g to 12c

Enterprise Manager 12c incident management functionality leverages your existing pre-12c monitoring setup out-of-box. Migration is seamless and transparent. For example, if your Enterprise Manager 10/11g monitoring system sends you e-mails based on specific monitoring conditions, you will continue to receive those e-mails without interruption. To take advantage of 12c features, however, you may need to perform additional migration tasks.

Important: Alerts that were generated pre-12c will still be available.

Rules

When you migrate to Enterprise Manger 12c, all of your existing notification rules are automatically converted to rules. Technically, they are converted to event rules first with incidents automatically being created for each event rule.

In general, event rules allow you to define which events should become incidents. However, they also allow you to take advantage of the Enterprise Manager's increased monitoring flexibility.

For more information on rule migration, see the following documents:

- Appendix A, " Overview of Notification in Enterprise Manager Cloud Control" section "Migrating Notification Rules to Rule Sets" in the *Enterprise Manager Cloud Control Upgrade Guide*.
- Chapter 29 "Updating Rules" in the *Enterprise Manager Cloud Control Upgrade Guide*.

Privilege Requirements

The *Create Enterprise Rule Set* resource privilege is now required in order to edit/create enterprise rule sets and rules contained within. The exception to this is migrated notification rules. When pre-12c notification rules are migrated to event rules, the original notification rule owners will still be able to edit their own rules without having been granted the Create Enterprise Rule Set resource privilege. However, they must be granted the *Create Enterprise Rule Set* resource privilege if they wish to create new rules. Enterprise Manager Super Administrators, by default, can edit and create rule sets.

- Privileges on events are calculated based on the privilege on the underlying source objects. For example, the user will have VIEW privilege on an event if he can view the target for the event.
- Privileges on an incident are calculated based on the privileges on participating events.
- Similarly, problem privileges are calculated based on privileges on underlying incidents.

4 Notifications

The notification system allows you to notify Enterprise Manager administrators when specific incidents, events, or problems arise.

> **Note:** This chapter assumes that you are familiar with incident management. For information about monitoring and managing your IT infrastructure via incident management, see Chapter 3, "Using Incident Management".

As an integral part of the management framework, notifications can also perform actions such as executing operating system commands (including scripts) and PL/SQL procedures when specific incidents, events, or problems occur. This capability allows you to automate IT practices. For example, if an incident (such as monitoring of the operational (up/down) status of a database) arises, you may want the notification system to automatically open an in-house trouble-ticket using an OS script so that the appropriate IT staff can respond in a timely manner.

By using Simple Network Management Protocol (SNMP) traps, the Enterprise Manager notification system also allows you to send traps to SNMP-enabled third-party applications such as HP OpenView for events published in Enterprise Manager. Some administrators may want to send third-party applications a notification when a certain metric has exceeded a threshold.

This chapter covers the following:

- Setting Up Notifications
- Extending Notification Beyond E-mail
- Passing Corrective Action Status Change Information
- Passing Job Execution Status Information
- Passing User-Defined Target Properties to Notification Methods
- Management Information Base (MIB)
- Troubleshooting Notifications

4.1 Setting Up Notifications

All Enterprise Manager administrators can set up e-mail notifications for themselves. Super Administrators also have the ability to set up notifications for other Enterprise Manager administrators.

4.1.1 Setting Up a Mail Server for Notifications

Before Enterprise Manager can send e-mail notifications, you must first specify the Outgoing Mail (SMTP) servers to be used by the notification system. Once set, you can then define e-mail notifications for yourself or, if you have Super Administrator privileges, other Enterprise Manager administrators.

You specify the Outgoing Mail (SMTP) server on the Notification Methods page (Figure 4–1). To display the Notification Methods page, from the **Setup** menu, select **Notifications**, then select **Notification Methods**.

Note: You must have Super Administrator privileges in order to set up SMTP servers.

Specify one or more outgoing mail server names, the mail server authentication credentials (User Name, Password, and Confirm Password), if required, the name you want to appear as the sender of the notification messages, and the e-mail address you want to use to send your e-mail notifications. This address, called the Sender's Mail Address, must be a valid address on each mail server that you specify. A message will be sent to this e-mail address if any problem is encountered during the sending of an e-mail notification. Example 4–1 shows sample notification method entries.

Example 4–1 Mail Server Settings

- **Outgoing Mail (SMTP) Server -** smtp01.example.com:587, smtp02.example.com
- **User Name -** myadmin
- **Password -** ******
- **Confirm Password -** ******
- **Identify Sender As -** Enterprise Manager
- **Sender's E-mail Address -** mgmt_rep@example.com
- **Use Secure Connection** - *No*: E-mail is not encrypted. *SSL*: E-mail is encrypted using the Secure Sockets Layer protocol. *TLS, if available*: E-mail is encrypted using the Transport Layer Security protocol if the mail server supports TLS. If the server does not support TLS, the e-mail is automatically sent as plain text.

Figure 4–1 Defining a Mail Server

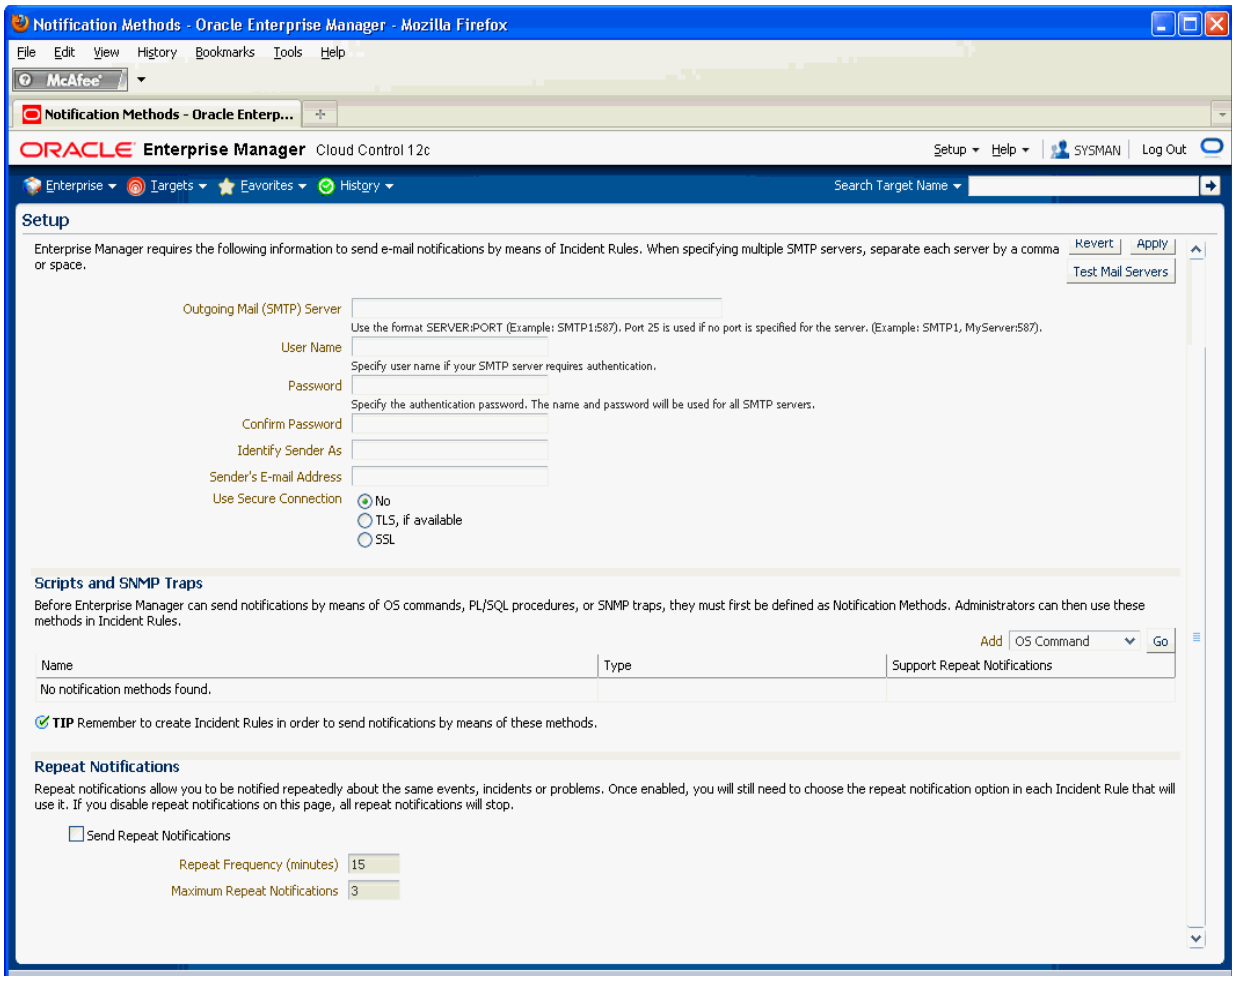

Note: The e-mail address you specify on this page is not the e-mail address to which the notification is sent. You will have to specify the e-mail address (where notifications will be sent) from the Password and E-mail page. From the **Setup** menu, choose **MyPreferences** and then **Enterprise Manager Password & E-mail**.

After configuring the e-mail server, click **Test Mail Servers** to verify your e-mail setup. You should verify that an e-mail message was received by the e-mail account specified in the **Sender's E-mail Address** field.

Defining multiple mail servers will improve the reliability of e-mail notification delivery. E-mail notifications will be delivered if at least one e-mail server is up. The notification load is balanced across multiple e-mail servers by the OMS, which switches through them (servers are allocated according to availability) after 20 e-mails have been sent. Switching is controlled by the *em.notification.emails_per_connection* emoms property.

4.1.1.1 Setting Up Repeat Notifications

Repeat notifications allow administrators to be notified repeatedly until an incident is either acknowledged or the number of **Maximum Repeat Notifications** has been reached. Enterprise Manager supports repeat notification for all notification methods

(e-mail, OS command, PL/SQL procedure, and SNMP trap). To enable this feature for a notification method, select the **Send Repeat Notifications** option. In addition to setting the maximum number of repeat notifications, you can also set the time interval at which the notifications are sent.

Important: For Oracle database versions 10 and higher, it is recommend that no modification be made to *aq_tm_processes* init.ora parameter. If, however, this parameter must be modified, its value should be at least one for repeat notification functionality. If the Enterprise Manager Repository database version is 9.2, the *aq_tm_processes* init.ora parameter must be set to at least one to enable repeat notification functionality.

Repeat Notifications for Incident Rules

Setting repeat notifications globally at the notification method level may not be provide sufficient flexibility. For example, you may want to have different repeat notification settings based on event type. Enterprise Manager accomplishes this by allowing you to set repeat notifications for individual incident rule sets or individual rules within a rule set. Repeat notifications set at the rule level take precedence over those defined at the notification method level.

Important: Repeat notifications for rules will only be sent if the **Send Repeat Notifications** option is enabled in the Notification Methods page.

For PL/SQL, OS command, and SNMP trap notification methods, you must enable each method to support repeat notifications. You can select **Supports Repeat Notifications** option when adding a new notification method or by editing an existing method.

Figure 4–2 Enabling Repeat Notification for an OS Command Notification Method

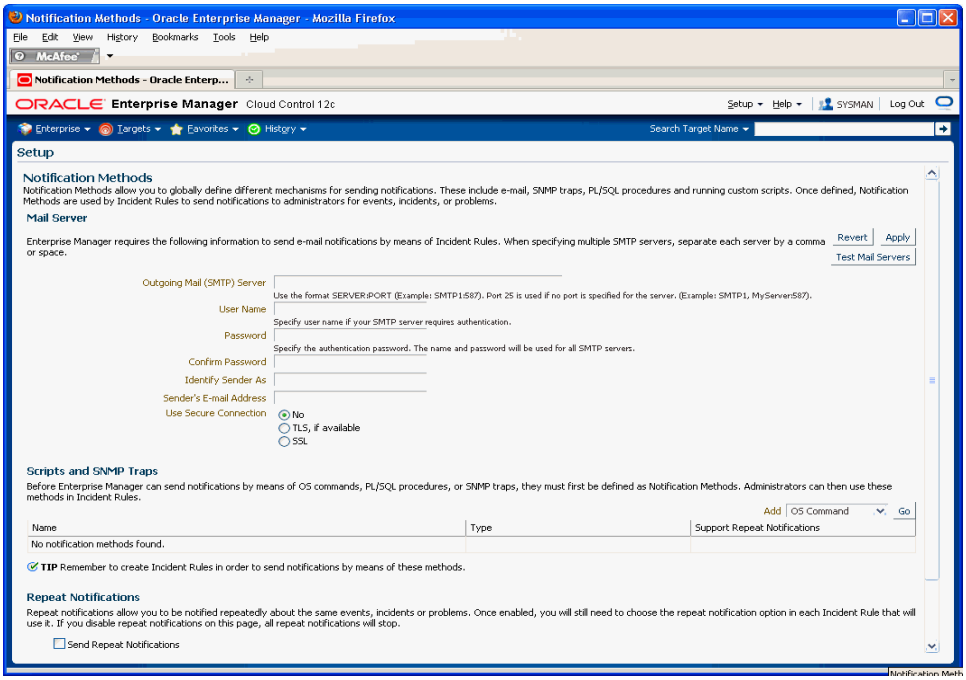

4.1.2 Setting Up E-mail for Yourself

If you want to receive notifications by e-mail, you will need to specify your e-mail address(s) in the Password & E-mail page (from the Setup menu, select **MyPreferences**, then select **Enterprise Manager Password & E-mail**). In addition to defining notification e-mail addresses, you associate the notification message format (long, short, pager) to be used for your e-mail address.

Setting up e-mail involves three steps:

Step 1: Define e-mail addresses.

Step 2: Set up a Notification Schedule.

Step 3: Subscribe to incident rules in order to receive e-mails.

4.1.2.1 Defining E-mail Addresses

An e-mail address can have up to 128 characters. There is no upper limit with the number of e-mail addresses.

To add an e-mail address:

1. From the **Setup** menu, select **MyPreferences**, then select **Enterprise Manager Password & E-mail**.
2. Click **Add Another Row** to create a new e-mail entry field in the **E-mail Addresses** table.
3. Specify the e-mail associated with your Enterprise Manager account. All e-mail notifications you receive from Enterprise Manager will be sent to the e-mail addresses you specify.

   For example, `user1@oracle.com`

Select the *E-mail Type* (message format) for your e-mail address. *E-mail (Long)* sends a HTML formatted e-mail that contains detailed information. Example 4–2 shows a typical notification that uses the long format.

E-mail (Short) and *Pager(Short)* (Example 4–3) send a concise, text e-mail that is limited to a configurable number of characters, thereby allowing the e-mail be received as an SMS message or page. The content of the message can be sent entirely in the subject, entirely in the body or split across the subject and body.

For example, in the last case, the subject could contain the severity type (for example, Critical) and the target name. The body could contain the time the severity occurred and the severity message. Since the message length is limited, some of this information may be truncated. If truncation has occurred there will be an ellipsis end of the message. *Pager(Short)* addresses are used for supporting the paging feature in incident rules. Note that the incident rules allow the rule author to designate some users to receive a page for critical issues.

4. Click **Apply** to save your e-mail address.

Example 4–2 Long E-mail Notification for Metric Alerts

```
Target type=Host
Target name=machine6140830.example.com
Message=Filesystem / has 54.39% available space, fallen below warning (60) or
critical (30) threshold.
Severity=Warning
Event reported time=Apr 28, 2011 2:33:55 PM PDT
Event Type=Metric Alert
Event name=Filesystems:Filesystem Space Available (%)
Metric Group=Filesystems
Metric=Filesystem Space Available (%)
Metric value=54.39
Key Value=/
Key Column 1=Mount Point
Rule Name=NotifRuleSet1,Event rule1
Rule Owner=SYSMAN
```

Example 4–3 Short E-mail Notification for Alerts

```
Subject is :
EM:Unreachable Start:myhost
Body is :
Nov 16, 2006 2:02:19 PM EST:Agent is Unreachable (REASON = Connection refused)
but the host is UP
```

More about E-mail(Short) and Pager(Short) Formats

Enterprise Manager does not directly support message services such as paging or SMS, but instead relies on external gateways to, for example, perform the conversion from e-mail to page. Beginning with Enterprise Manager 12c, the notification system allows you to tag e-mail addresses explicitly as 'page' or 'e-mail'. Explicit system differentiation between these two notification methods allows you to take advantage of the multiple action capability of incident rules. For example, the e-mail versus page distinction is required in order to send you an e-mail if an event severity is 'warning' or page you if the severity is 'critical'. To support this capability, a Pager format has been made available that sends an abbreviated version of the short format e-mail.

EMOMS properties can be used for controlling the size and format of the short e-mail. The following table lists emoms properties for Notification System.

Table 4–1 EMOMS Properties for Notifications

Property Name	Default Value	Description
em.notification.emails_per_minute	250	Email delivery limits per minute. The Notification system uses this value to throttle number of Email delivery per minutes. Customer should set the value lower if doesn't want to over flow the Email server, or set the value higher if the Email server can handle high volume of Emails.
em.notification.cmds_per_minute		OS Command delivery limits per minute. The Notification system uses this value to throttle number of OS Command delivery per minutes.
em.notification.os_cmd_timeout	30	OS Command delivery timeout in seconds. This value indicates how long to allow OS process to execute the OS Command delivery. Set this value higher if the OS command script requires longer time to complete execution.
em.notification.plsql_per_minute	250	PL/SQL delivery limits per minute. The Notification system uses this value to throttle number of PL/SQL delivery per minutes.
em.notification.java_per_minute	500	JAVA delivery limits per minute. The Notification system uses this value to throttle number of Java delivery per minutes.
em.notification.ticket_per_minute	250	Ticket delivery limits per minute. The Notification system uses this value to throttle number of Ticket delivery per minutes.
em.notification.traps_per_minute	250	SNMP delivery limits per minute. The Notification system uses this value to control the number of SNMP Trap per minutes.
em.notification.locale.plsql	OMS Locale	This property specifies the Locale delivered by advanced PL/SQL notification. The customer can define this property to overwrite the default Locale where the OMS is installed.
em.notification.locale.email	OMS Locale	This property specifies the Locale delivered by Email. Customer can define this property to overwrite the default Locale where the OMS is installed.
em.notification.locale.osmcd	OMS Locale	This property specifies the Locale delivered by OS Command. Customer can define this property to overwrite the default Locale where the OMS is installed.
em.notification.locale.snmp	OMS Locale	This property specifies the Locale delivered by SNMP trap. Customer can define this property to overwrite the default Locale where the OMS is installed.
em.notification.oscmd.max_env_var_length	512	The maximum length of OS Common environment variable value.
em.notification.snmp.max_oid_length	2560	The maximum length of SNMP OID value.

Table 4–1 (Cont.) EMOMS Properties for Notifications

Property Name	Default Value	Description
em.notification.min_delivery_threads	6	The minimum number of active threads in the thread pool initially and number of active threads are running when system is in low activities. Setting the value higher will use more system resources, but will deliver more notifications.
em.notification.max_delivery_threads	24	The maximum number of active threads in the thread pool when the system is in the high activities. This value should greater than em.notification.min_delivery_threads. Setting the value higher will use more system resources and deliver more notifications.
em.notification.short_format_length	>=1 (155)	The size limit of the total number of characters in short email format. The customer should modify this property value to fit their Email or Pager limit content size.
em.notification.snmp_packet_length	>=1 (5120)	The maximum size of SNMP Protocol Data unit.
em.notification.email_content_transfer_encoding	8-bit, 7-bit(QP), 7-bit(BASE64) (8-bit)	The character set that can encode the Email. Oracle supports three character sets : 8-bit, 7-bit(QP), and7-bit(BASE64).
em.notification.emails_per_connection	>=1 (20)	The maximum number of emails delivered to same email gateway before switching to the next available email gateway (assumes customers have configured multiple email gateways). This property is used for email gateway load balance.
em.notification.short_format	both, subject, body (both)	Use short format on both subject and body, subject only, or body only..

You must establish the maximum size your device can support and whether the message is sent in subject, body or both.

You can modify the EMOMS properties by using the Enterprise Manager command line control `emclt get/set/delete/list` property command.

Get Property Command

```
emctl get [-sysman_pwd "sysman password"]-name em.notification.short_format_length
```

Set Property Command

```
emctl set  property -name em.notification.short_format_length -value 155
```

Emoms Properties Entries for a Short E-mail Format

```
emctl set  property -name em.notification.short_format_length -value 155
emctl set property -name em.notification.short_format -value both
```

4.1.2.2 Setting Up a Notification Schedule

Once you have defined your e-mail notification addresses, you will need to define a notification schedule. For example, if your e-mail addresses are user1@oracle.com, user2@oracle.com, user3@oracle.com, you can choose to use one or more of these e-mail addresses for each time period in your notification schedule.

> **Note:** When you enter e-mail addresses for the first time, a 24x7 weekly notification schedule is set automatically. You can then review and modify the schedule to suit your monitoring needs.

A notification schedule is a repeating schedule used to specify your on-call schedule—the days and time periods and e-mail addresses that should be used by Enterprise Manager to send notifications to you. Each administrator has exactly one notification schedule. When a notification needs to be sent to an administrator, Enterprise Manager consults that administrator's notification schedule to determine the e-mail address to be used. Depending on whether you are Super Administrator or a regular Enterprise Manager administrator, the process of defining a notification schedule differs slightly.

If you are a regular Enterprise Manager administrator and are defining your own notification schedule:

1. From **Setup** menu, select **Notifications**, then select **My Notification Schedule**.
2. Follow the directions on the Notification Schedule page to specify when you want to receive e-mails.

4.1.2.3 Subscribe to Receive E-mail for Incident Rules

An incident rule is a user-defined rule that specifies the criteria by which notifications should be sent for specific events that make up the incident. An incident rule set, as the name implies, consists of one or more rules associated with the same incident.

When creating an incident rule, you specify criteria such as the targets you are interested in, the types of events to which you want the rule to apply. Specifically, for a given rule, you can specify the criteria you are interested in and the notification methods (such as e-mail) that should be used for sending these notifications. For example, you can set up a rule that when any database goes down or any database backup job fails, e-mail should be sent and the "log trouble ticket" notification method should be called. Or you can define another rule such that when the CPU or Memory Utilization of any host reach critical severities, SNMP traps should be sent to another management console.

Notification flexibility is further enhanced by the fact that with a single rule, you can perform multiple actions based on specific conditions. Example: When monitoring a condition such as machine memory utilization, for an incident severity of 'warning' (memory utilization at 80%), send the administrator an e-mail, if the severity is 'critical' (memory utilization at 99%), page the administrator immediately.

You can subscribe to a rule you have already created.

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. On the Incident Rules page, select the desired rule.
3. From the **Actions** menu, select **Notifications**, then select **Basic Notifications**.

Out-of-Box Incident Rules

Enterprise Manager comes with two incident rule sets that cover the most common monitoring conditions, they are:

- Incident Management Ruleset for All Targets
- Event Management Ruleset for Self Update

If the conditions defined in the out-of-box incident rules meet your requirements, you can simply subscribe to receive e-mail notifications for the conditions defined in the rule using the subscribe procedure shown in the previous section.

Creating Your Own Incident Rules

You can define your own custom rules. The following procedure documents the process of incident rule creation for non-Super Administrators.

To create your own incident rule:

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.

   The Incident Rules page displays. From this page you can create a new rule set, to which you can add new rules. Alternatively, if you have the requisite permissions, you can add new rules to existing

2. Click **Create Rule Set...**

   The create rule set page displays.

3. Specify the **Name**, **Description**, and the **Targets** to which the rules set should apply.
4. Click the **Rules** tab, then click **Create**.
5. Choose the incoming incident, event or problem to which you want the rule to apply. See "Working with Rules" for more information.
6. Click **Continue**.

   Enterprise Manager displays the Create Incident Rule pages. Enter the requisite information on each page to create your incident rule.

7. Follow the wizard instructions to create your rule.

   Once you have completed defining your rule, the wizard returns you to the create rule set page.

8. Click **Save** to save the incident rule set.

4.1.3 Setting Up E-mail for Other Administrators

If you have Super Administrator privileges, you can set up e-mail notifications for other Enterprise Manager administrators. To set up e-mail notifications for other Enterprise Manager administrators, you need to:

Step 1: Ensure Each Administrator Account has an Associated E-mail Address

Each administrator to which you want to send e-mail notifications must have a valid e-mail address.

1. From the **Setup** menu, select **Security** and then **Administrators**.
2. For each administrator, define an e-mail address. This sets up a 24x7 notification schedule for this user that uses all the e-mail addresses specified.

Enterprise Manager also allows you to specify an administrator address when editing an administrator's notification schedule.

Step 2: Define Administrators' Notification Schedules

Once you have defined e-mail notification addresses for each administrator, you will need to define their respective notification schedules. Although a default 24x7 notification schedule is created when you specify an e-mail address for the first time, you should review and edit the notification schedule as needed.

1. From the **Setup** menu, select **Notifications**, then select **Notification Schedule**.

   From the vertical navigation bar, click Schedules (under Notification). The Notification Schedule page appears.

2. Specify the administrator who's notification schedule you wish to edit and click **Change**.
3. Click **Edit Schedule Definition**. The Edit Schedule Definition: Time Period page appears. If necessary, modify the rotation schedule.
4. Click **Continue**. The Edit Schedule Definition: E-mail Addresses page appears.
5. Follow the directions on the Edit Schedule Definition: E-mail Addresses page to modify the notification schedule.
6. Click **Finish** when you are done.
7. Repeat steps three through seven for each administrator.

Step 3: Assign Incident Rules to Administrators

With the notification schedules set, you now need to assign the appropriate incident rules for each designated administrator.

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. Select the desired **Ruleset** and click **Edit**.
3. Click on the **Rules** tab, select the desired rule and click **Edit**.
4. Click **Add Actions**, select desire action and click **Edit**.
5. Enter the **Administrator** name on either **E-mail To** or **E-mail Cc** field in the **Basic Notification** region.
6. Click **Continue**, click **Next**, click **Next**, click **Continue**, and finally click **Save**.

4.1.4 E-mail Customization

Enterprise Manager allows Super Administrators to customize global e-mail notifications for the following types: All events, incidents, problems, and specific event types installed. You can alter the default behavior for all events by customizing *Default Event Email Template*. In addition, you can further customize the behavior for a specific event type by customizing the template for the event type. For instance, you can customize the *Metric Alert Events* template for the metric alert event type. Using predefined building blocks (called attributes and labels) contained within a simple script, Super Administrators can customize alert e-mails by selecting from a wide variety of information content.

To customize an e-mail:

1. From the **Setup** menu, select **Notifications,** then select **Customize Email Formats**.
2. Choose the **Type** and **Format**.

3. Click **Customize**. The Customize E-mail Template page displays.

From the Customize E-mail Template page, you can modify the content of the e-mail template Enterprise Manager uses to generate e-mail notifications. Extensive information on script formatting, syntax, and options is available from the Edit E-mail Template page via imbedded assistance and online help.

Figure 4–3 E-mail Customization

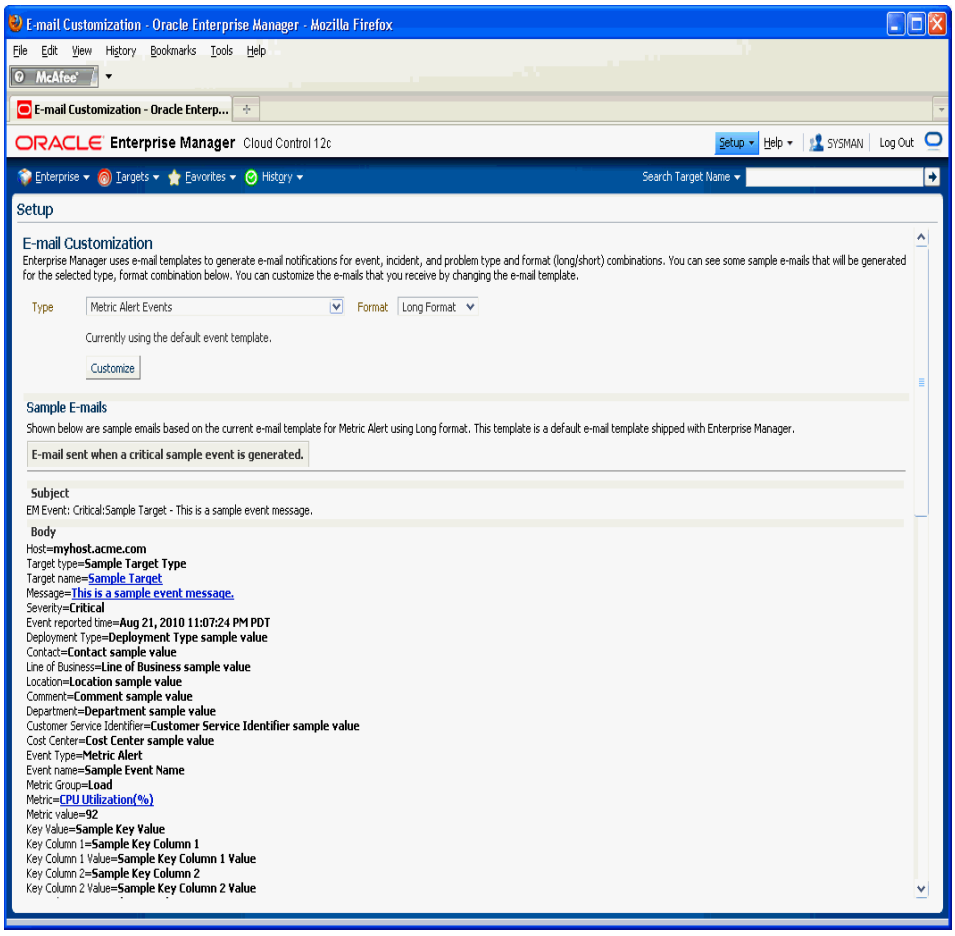

4.1.4.1 E-mail Customization Reference

The following reference summarizes the semantics and component syntax of the pseudo-language used to define e-mails. The pseudo-language provides you with a simple, yet flexible way to customize e-mail notifications. The following is a summary of pseudo-language conventions/limitations:

- You can add comments (or any free-form text) using separate lines beginning with "--" or at end of lines.
- You can use attributes.
- You can use IF & ELSE & ENDIF control structures. You can also use multiple conditions using "AND" or "OR". Nested IF statements are not supported.
- You can insert spaces for formatting purposes. Spaces at the beginning of a line will be ignored in the actual e-mail. To insert spaces at the beginning of a line, use the [SP] attribute.

- Use "/" to escape and "[" or "]" if you want to add attribute names, operators, or IF clauses to the actual e-mail.
- HTML is not supported.

Reserved Words and Operators

The following table lists all reserved words and operators used when modifying e-mail scripts.

Table 4–2 Reserved Words and Operators

Reserved Word/Operator	Description
IF, ELSIF, ENDIF, ELSE	Used in IF-ELSE constructs.
AND, OR	Boolean operators – used in IF-ELSE constructs only.
NULL	To check NULL value for attributes - used in IF-ELSE constructs only.
|	Pipe operator – used to show the first non-NULL value in a list of attributes. For example: `METRIC_NAME|SEVERITY`
EQ, NEQ	Equal and Not-Equal operators – applicable to NULL, STRING and NUMERIC values.
/	Escape character – used to escape reserved words and operators. Escape characters signify that what follows the escape character takes an alternative interpretation.
[,]	Delimiters used to demarcate attribute names and IF clauses.

Syntax Elements

Literal Text

You can specify any text as part of the e-mail content. The text will be displayed in the e-mail and will not be translated if the Oracle Management Services (OMS) language setting is changed. For example, 'my Oracle Home' appears as 'my Oracle Home' in the generated e-mail.

Predefined Attributes

Predefined attributes/labels will be substituted with actual values in a specific context. To specify a predefined attribute/label, use the following syntax:

```
[PREDEFINED_ATTR]
```

Attribute names can be in either UPPER or LOWER case. The parsing process is case-insensitive.

A pair of square brackets is used to distinguish predefined attributes from literal text. For example, for a job e-mail notification, the actual job name will be substituted for `[EXECUTION_STATUS]`. For a metric alert notification, the actual metric column name will be substituted for `[METIRC_COLUMN]`.

You can use the escape character "/" to specify words and not have them interpreted as predefined labels/attributes. For example, "`/[NEW/]`" will not be considered as the predefined attribute `[NEW]` when parsed.

Operators

`EQ, NEQ` – for text and numeric values

`NULL`- for text and numeric values

`GT, LT, GE, LE` – for numeric values

Control Structures

The following table lists acceptable script control structures.

Table 4–3 Control Structures

Control Structure	Description
Pipe "|"	Two or more attributes can be separated by '|' character. For example, `[METRIC_NAME|SEVERITY]` In this example, only the applicable attribute within the current alert context will be used (replaced by the actual value) in the e-mail. If more than one attributes are applicable, only the left-most attribute is used.

Table 4–3 (Cont.) Control Structures

Control Structure	Description

IF	Allows you to make a block of text conditional. Only one level of IF and ELSIF is supported. Nested IF constructs are not supported.
	
	All attributes can be used in IF or ELSIF evaluation using EQ/NEQ operators on NULL values. Other operators are allowed for "SEVERITY" and "REPEAT_COUNT" only.
	
	Inside the IF block, the values need to be contained within quotation marks " ". Enterprise Manager will extract the attribute name and its value based on the position of "EQ" and other key words such as "and", "or". For example,
	
	`[IF REPEAT_COUNT EQ "1" AND SEVERITY EQ "CRITICAL" THEN]`
	
	The statement above will be true when the attributes of the alert match the following condition:
	
	■ Attribute Name: REPEAT_COUNT
	■ Attribute Value: 1
	■ Attribute Name: SEVERITY
	■ Attribute Value: CRITICAL
	
	Example IF Block:
	
	`[IF EXECUTION_STATUS NEQ NULL]`
	`  [JOB_NAME_LABEL]=[EXECUTION_STATUS]`
	`  [JOB_OWNER_LABEL]=[JOB_OWNER]`
	`[ENDIF]`
	
	`[IF SEVERITY_CODE EQ CRITICAL ]`
	`  [MTRIC_NAME_LABEL]=[METRIC_GROUP]`
	`  [METRIC_VALUE_LABEL]=[METRIC_VALUE]`
	`  [TARGET_NAME_LABEL]=[TARGET_NAME]`
	`  [KEY_VALUES]`
	`[ENDIF]`
	
	Example IF and ELSEIF Block:
	
	`[IF SEVERITY_CODE EQ CRITICAL]`
	`  statement1`
	
	`[ELSIF SEVERITY_CODE EQ WARNING]`
	`  statement2`
	
	`[ELSIF SEVERITY_CODE EQ CLEAR]`
	`  statement3`
	
	`[ELSE]`
	`  statement4`
	
	`[ENDIF]`

Comments

You can add comments to your script by prefacing a single line of text with two hyphens "--". For example,

```
-- Code added on 8/3/2009
    [IF REPEAT_COUNT NEQ NULL]
    . . .
```

Comments may also be placed at the end of a line of text.

```
[IF SEVERITY_SHORT EQ W] -- for Warning alert
```

HTML Tags in Customization Content

Use of HTML tags is not supported.

When Enterprise Manager parses the e-mail script, it will convert the “<” and “>” characters of HTML tags into encoded format (< and >). This ensures that the HTML tag is not treated as HTML by the destination system.

Examples

E-mail customization template scripts support three main operators.

- Comparison operators: EQ/NEQ/GT/LT/GE/LE
- Logic operators: AND/OR
- Pipeline operator: |

4.2 Extending Notification Beyond E-mail

Notification Methods are the mechanisms by which notifications are sent. Enterprise Manager Super Administrators can set up e-mail notifications by configuring the 'e-mail' notification method. Most likely this would already have been set up as part of the Oracle Management Service installation.

Enterprise Manager Super Administrators can also define other custom notification methods. For example, event notifications may need to be forwarded to a 3rd party trouble-ticketing system. Assuming APIs to the third-party trouble-ticketing system are available, a custom notification method can be created to call a custom OS script that has the appropriate APIs. The custom notification method can be named in a user-friendly fashion, for example, "Log trouble ticket". Once the custom method is defined, whenever an administrator needs to send alerts to the trouble-ticketing system, he simply needs to invoke the now globally available notification method called "Log trouble ticket".

Custom notification methods can be defined based on any custom OS script, any custom PL/SQL procedure, or by sending SNMP traps. A fourth type of notification method (Java Callback) exists to support Oracle internal functionality and cannot be created or edited by Enterprise Manager administrators.

Only Super Administrators can define OS Command, PL/SQL, and SNMP Trap notification methods. However, any Enterprise Manager administrator can add these notification methods (once defined by the Super Administrator) as actions to their incident rules.

Through the Notification Methods page, you can:

- Set up the outgoing mail servers if you plan to send e-mail notifications through incident rules
- Create other custom notification methods using OS and PL/SQL scripts and SNMP traps.
- Set global repeat notifications.

4.2.1 Custom Notification Methods Using Scripts and SNMP Traps

You can create other custom notification methods based on OS scripts, PL/SQL procedures, or SNMP traps. Any administrator can then use these methods in incident rules.

The length of the SNMP OID value is limited to 2560 bytes by default. Configure emoms property *em.notification.snmp.max_oid_length* to change the default limit.

For Enterprise Manager 12c, SNMP traps are delivered for event notifications only. SNMP trap notifications are not supported for incidents or problems.

Note: SNMP advanced notification methods defined using previous versions of Enterprise Manager (pre-12c) will continue to function without modification. Traps will conform to the older Enterprise Manager MIB definition.

4.2.1.1 Adding a Notification Method based on an OS Command or Script

Notification system invokes the custom script when an incident rule matches the OS Command advanced notification action. Custom script receives notifications for matching events, incidents and problem through environment variables.

The length of any environment variable's value is limited to 512 characters by default. Configure emoms property named em.notification.oscmd.max_env_var_length for changing the default limit.

Note: Notification methods based on OS commands must be configured by an administrator with Super Administrator privileges.

Step 1: Define your OS command or script.

You can specify an OS command or script that will be called by the notification system when an incident rule matches the OS Command advanced notification action. You can use incident, event, or problem context information, corrective action execution status and job execution status within the body of the script. Passing this contextual information to OS commands/scripts allows you to customize automated responses specific event conditions. For example, if an OS script opens a trouble ticket for an in-house support trouble ticket system, you will want to pass severity levels (critical, warning, and so on) to the script to open a trouble ticket with the appropriate details and escalate the problem. For more information on passing specific types of information to OS Command or Scripts, see:

- "Passing Event, Incident, Problem Information to an OS Command or Script" on page 4-19
- "Passing Corrective Action Execution Status to an OS Command or Script" on page 4-50
- "Passing Job Execution Status to an OS Command or Script" on page 4-50

Step 2: Deploy the script on each Management Service host.

You must deploy the OS Command or Script on each Management Service host machine that connects to the Management Repository. The OS Command is run as the user who started the Management Service.

The OS Command or Script should be deployed on the same location on each Management Service host machine. The OS Command should be an absolute path, for example, /u1/bin/logSeverity.sh. The command is run by the user who started the Management Service. If an error is encountered during the running of the OS Command, the Notification System can be instructed to retry the sending of the notification to the OS Command by returning an exit code of 100. The procedure is initially retried after one minute, then two minutes, then three minutes and so on, until the notification is a day old, at which point it will be purged.

Example 4–4 shows the parameter in emoms.properties that controls how long the OS Command can execute without being killed by the Management Service. This is to prevent OS Commands from running for an inordinate length of time and blocking the delivery of other notifications. By default the command is allowed to run for 30 seconds before it is killed. The *em.notification.os_cmd_timeout* emoms property can be configured to change the default timeout value.

Example 4–4 Changing the em.notification.os_cmd_timeout EMOMS Property

```
emctl set property -name em.notification.os_cmd_timeout value 30
```

Step 3: Register your OS Command or Script as a new Notification Method.

Add this OS command as a notification method that can be called in incident rules. Log in as a Super Administrator. From the **Setup** menu, select **Notifications**, then select **Notification Methods**. From this page, you can define a new notification based on the 'OS Command' type. See "Adding a Notification Method based on an OS Command or Script".

The following information is required for each OS command notification method:

- Name
- Description

  Both Name and Description should be clear and intuitive so that the function of the method is clear to other administrators.

- OS Command

You must enter the full path of the OS command or script in the OS command field (for example, `/u1/bin/myscript.sh`). For environments with multiple Management Services, the path must be exactly the same on each machine that has a Management Service. Command line parameters can be included after the full path (for example, /u1/bin/myscript.sh arg1 arg2).

Example 4–5 shows information required for the notification method.

Example 4–5 OS Command Notification Method

```
Name Trouble Ticketing
Description Notification method to log trouble ticket for a severity occurrence
OS Command /private/mozart/bin/logTicket.sh
```

> **Note:** There can be more than one OS Command configured per system.

Step 4: Assign the notification method to an instance rule.

You can edit an existing rule (or create a new instance rule), then go to the Methods page. From the **Setup** menu, choose **Incidents** and then **Incident Rules**. The Incident Rules page provides access to all available rule sets.

Passing Event, Incident, Problem Information to an OS Command or Script

The notification system passes information to an OS script or executable using system environment variables.

Conventions used to access environmental variables vary depending on the operating system:

- UNIX: $ENV_VARIABLE
- Windows:%ENV_VARIABLE%

The notification system sets the following environment variables before calling the script. The script can then use any or all of these variables within the logic of the script.

Environment Variables Common to Event, Incident and Problem

Table 4–4 Generic Environment Variables

Environment Variable	Description
NOTIF_TYPE	Type of notification and possible values NOTIF_NORMAL, NOTIF_RETRY, NOTIF_DURATION, NOTIF_REPEAT, NOTIF_CA, NOTIF_RCA
REPEAT_COUNT	How many times the notification has been sent out before this notification.
RULESET_NAME	The name of the ruleset that triggered this notification.
RULE_NAME	The name of the rule that triggered this notification.
RULE_OWNER	The owner of the ruleset that triggered this notification.
MESSAGE	The message of the event, incident, or problem.
MESSAGE_URL	EM console URL for this message.

Table 4–5 Category-Related Environment Variables

Environment Variable	Description
CATEGORIES_COUNT	Number of categories in this notification. This value is equal to1 if one category is associated with event, incident or problem. It is equal to 0 if no category associated with event, incident or problem.
CATEGORY_CODES_ COUNT	Number of category codes in this notification.
CATEGORY_n	Category is translated based on locale defined in OMS server. Valid values for the suffix "_n" are between 1.. $CATEGORIES_ COUNT
CATEGORY_CODE_n	Codes for the categories. Valid values for the suffix "_n" are between 1..$CATEGORY_CODES_COUNT

The following lists the common environment variables for User Defined Target Properties. They will be populated under the following cases: (a) When an event has a related target, (b) When an incident or a problem have single event source and have a related target.

Table 4–6 User-Defined Target Property Environment Variables

Environment Variable	Description
ORCL_GTP_COMMENT	Comment
ORCL_GTP_CONTACT	Contact
ORCL_GTP_COST_ CENTER	Cost Center
ORCL_GTP_ DEPARTMENT	Department
ORCL_GTP_ DEPLOYMENT_TYPE	Deployment type
ORCL_GTP_LINE_OF_ BUS	Line of Business
ORCL_GTP_LOCATION	Location

Event Notification-Specific Environment Variables

Table 4–7 Event Notification-Specific Environment Variables

Environment Variable	Description
EVENT_NAME	Event Name.
EVENT_REPORTED_ TIME	Event reported date.
EVENT_SOURCE_ COUNT	Number of Sources associated with this event.
EVENT_TYPE	Event type.
EVENT_OCCURRENCE_ TIME	Event occurrence time.
EVENT_TYPE_ATTRS	The list of event type specific attributes.
EVENT_CONTEXT_ ATTRS	Event context data.
LAST_UPDATED_TIME	Last updated time
SEQUENCE_ID	Event sequence global unique identifier.
SEVERITY	Severity of event, it is translated.
SEVERITY_CODE	Code for event severity. Possible values are the following. FATAL, CRITICAL, WARNING, MINOR_WARNING, INFORMATIONAL, and CLEAR
ACTION_MSG	Message describing the action to take for resolving the event.
TOTAL_OCCURRENCE_ COUNT	Total number of duplicate occurrences
SEQUENCE_ID	Event sequence ID
RCA_DETAILS	If RCA is associated with this events.

The following tables lists the environment variables for the incident associated with an event. They are populated when the event is associated with an incident.

Table 4–8 Associated Incident Environment Variables

Environment Variable	Description
ASSOC_INCIDENT_ ACKNOWLEDGED_BY_ OWNER	Set to yes, if associated incident was acknowledged by owner
ASSOC_INCIDENT_ ACKNOWLEDGED_ DETAILS	The details of associated incident acknowledgement. For example: No - if not acknowledged Yes By userName - if acknowledged
ASSOC_INCIDENT_ STATUS	Associated Incident Status
ASSOC_INCIDENT_ID	Associated Incident ID
ASSOC_INCIDENT_ PRIORITY	Associated Incident priority. Supported value are Urgent, Very High, High, Medium,Low, None.
ASSOC_INCIDENT_ OWNER	Associated Incident Owner if it is existed.
ASSOC_INCIDENT_ ESCALATION_LEVEL	Escalation level of the associated incident has a value between 0 to 5.

Following lists the common environment variables related to the Source Object. They are populated when $SOURCE_OBJ_TYPE is not TARGET.

Table 4–9 Source Object-Related Environment Variables

Environment Variable	Description
SOURCE_OBJ_TYPE	Type of the Source object. For example, JOB, TEMPLATE.
SOURCE_OBJ_NAME	Source Object Name.
SOURCE_OBJ_NAME_ URL	Source's event console URL.
SOURCE_OBJ_SUB_TYPE	Sub-type of the Source object. For example, it provides the underlying job type for job status change events.
SOURCE_OBJ_OWNER	Owner of the Source object.

Following lists the common environment variables for the target, associated with the given issue. They are populated when the issue is related to a target.

Table 4–10 Target-Related Environment Variables

Environment Variable	Description
TARGET_NAME	Name of Target
TARGET_TYPE	Type of Target
TARGET_OWNER	Owner of Target
HOST_NAME	The name of the host on which the target is deployed upon.
TARGET_URL	Target's Enterprise Manager Console URL.

Table 4–10 (Cont.) Target-Related Environment Variables

Environment Variable	Description
TARGET_LIFECYCLE_STATUS	Life Cycle Status of the target. Possible values: Production, MissionCritical, Stage, Test, and Development. It is null if not defined.
TARGET_VERSION	Target Version of the target

Events are classified into multiple types. For example, the mertc_alert event type is used for modeling metric alerts. Following SQL query lists the environment variables corresponding to the event type specific attributes.

```
Select event_class as event_type, upper(name) as env_var_name
from em_event_class_attrs
where notif_order != 0
and event_class is not null
union
select event_class as event_type, upper(name) || '_NLS' as env_var_name
from em_event_class_attrs
where notif_order != 0
and event_class is not null
and is_translated = 1
order by event_type, env_var_name;
```

There is environment variable payload specific to each event type which can be accessed from the OS scripts.The following tables list notification attributes for the most critical event types.

Table 4–11 Environment variables specific to Metric Alert event type

Environment Variable	Description
COLL_NAME	The name of the collection collecting the metric.
COLL_NAME_NLS	The translated name of the collection collecting the metric
KEY_COLUMN_X	Internal name of Key Column X where X is a number between 1 and 7.
KEY_COLUMN_X_NLS	Translated name of Key Column X where X is a number between 1 and 7.
KEY_COLUMN_X_VALUE	Value of Key Column X where X is a number between 1 and 7.
KEY_VALUE	Monitored object for the metric corresponding to the Metric Alert event.
METRIC_COLUMN	The name of the metric column
METRIC_COLUMN_NLS	The translated name of the metric column.
METRIC_DESCRIPTION	Brief description of the metric.
METRIC_DESCRIPTION_NLS	Translated brief description of the metric.
METRIC_GROUP	The name of the metric.
METRIC_GROUP_NLS	The translated name of the metric
NUM_KEYS	The number of key metric columns in the metric.

Table 4–11 (Cont.) Environment variables specific to Metric Alert event type

Environment Variable	Description
SEVERITY_GUID	The guid of the severity record associated with this metric alert.
VALUE	Value of the metric when the event triggered.

Table 4–12 Environment variables specific to Target Availability event type

Environment Variable	Description
AVAIL_SEVERITY	The transition severity that resulted in that status of the target to change to the current availability status..
AVAIL_SUB_STATE	The sub-status of a target for the current status.
CYCLE_GUID	The guid of the first severity record in this availability cycle.
METRIC_GUID	Metric Guid of response metric.
SEVERITY_GUID	The guid of the severity record associated with this availability status.
TARGET_STATUS	The current availability status of the target.
TARGET_STATUS_NLS	The translated current availability status of the target.

Table 4–13 Environment variables specific to Job Status Change event type

Environment Variable	Description
EXECUTION_ID	Unique ID of the job execution..
EXECUTION_LOG	The job output of the last step executed.
EXECUTION_STATUS	The internal status of the job execution.
EXECUTION_STATUS_ NLS	The translated status of the job execution.
EXEC_STATUS_CODE	Execution status code of job execution.
STATE_CHANGE_GUID	Unique ID of last status change

You can use SQL queries to list the deployed event types in your deployment and the payload specific to each one of them. The following SQL can be used to list all internal event type names which are registered in the Enterprise Manager.

```
select class_name as event_type_name from em_event_class;
```

Following SQL lists environment variables specific to metric_alert event type.

```
select env_var_name
  from
    ( Select event_class as event_type, upper(name) as env_var_name
     from em_event_class_attrs
    where notif_order != 0
    and event_class is not null
    union
    select event_class as event_type, upper(name) || '_NLS' as env_var_name
    from em_event_class_attrs
    where notif_order != 0
```

```
and event_class is not null
and is_translated = 1)
where event_type = 'metric_alert';
```

You can also obtain the description of notification attributes specific to an event type directly from the Enterprsie Manager console:

1. From the **Setup** menu, select **Notifications**, then select **Customize Email Formats**.
2. Select the event type.
3. Click **Customize**.
4. Click **Show Predefined Attributes**.

Environment variables, ending with the suffix _NLS, provide the translated value for given attribute. For example, METRIC_COLUMN_NLS environment variable will provide the translated value for the metric_column attribute. Translated values will be in the locale of the OMS.

Environment Variables Specific to Incident Notifications

Table 4–14 Incident-Specific Environment Variables

Environment Variable	Description
SEVERITY	Incident Severity, it is translated. Possible Values: Fatal, Critical, Warning, Informational, Clear
SEVERITY_CODE	Code for Severity. Possible values are the FATAL, CRITICAL, WARNING, MINOR_WARNING, INFORMATIONAL, and CLEAR
INCIDENT_REPORTED_TIME	Incident reported time
INCIDENT_ACKNOWLEDGED_BY_OWNER	Set yes, if incident is acknowledged by owner.
INCIDENT_ID	Incident ID
INCIDENT_OWNER	Incident Owner
ASSOC_EVENT_COUNT	The number events associated with this incident.
INCIDENT_STATUS	Incident status. There are two internal fixed resolution status. NEW CLOSED Users can define additional statuses.
ESCALATED	Is Incident escalated
ESCALATED_LEVEL	The escalated level of incident.
PRIORITY	Incident priority. It is the translated priority name. Possible Values: Urgent, Very High, Hight, Medium, Low,None

Table 4–14 (Cont.) Incident-Specific Environment Variables

Environment Variable	Description
PRIOTITY_CODE	Incident priority code It is the internal value defined in EM. PRIORITY_URGENT PRIORITY_VERY_HIGH PRIORITY_HIGH PRIORITY_MEDIUM PRIORITY_LOW PRIORITY_NONE
TICKET_STATUS	Status of external ticket, if it exists.
TICKET_ID	ID of external ticket, if it exists.
LAST_UPDATED_TIME	Incident last update time

Following lists the associated problem's environment variables, when the incident is associated with a problem.

Table 4–15 Associated Problem Environment Variables

Environment Variable	Description
ASSOC_PROBLEM_ ACKNOWLEDGED_BY_ OWNER	Set to yes, if this problem was acknowledged by owner
ASSOC_PROBLEM_ STATUS	Associated Problem Status
ASSOC_PROBLEM_ID	Associated Problem ID
ASSOC_PROBLEM_ PRIORITY	Associated Problem priority
ASSOC_PROBLEM_ OWNER	Associated Problem Owner if it is existed.
ASSOC_PROBLEM_ ESCALATION_LEVEL	Escalation level of the associated Problem has a value between 0 to 5.

Environment Variables Specific to Problem Notifications

Table 4–16 Problem-Specific Environment Variables

Environment Variable	Description
SEVERITY	Problem Severity, it is translated.
SEVERITY_CODE	Code for Severity. Possible values are : FATAL, CRITICAL, WARNING, MINOR_WARNING, INFORMATIONAL, and CLEAR
PROBLEM_REPORTED_ TIME	Problem reported time.

Table 4–16 (Cont.) Problem-Specific Environment Variables

Environment Variable	Description
PROBLEM_ ACKNOWLEDGED_BY_ OWNER	Set yes, if problem is acknowledged by owner.
PROBLEM_ID	Problem ID
PROBLEM_KEY	Problem Key
PROBLEM_OWNER	Problem Owner
ASSOC_INCIDENT_ COUNT	The number incident associated with this problem..
PROBLEM_STATUS	Incident status. They are STATUS_NEW STATUS_CLOSED Any other user defined status.
ESCALATED	Is Incident escalated. Yes if it is escalated, otherwise No.
ESCALATED_LEVEL	The escalated level of incident.
PRIORITY	Incident priority. It is the translated priority name..
PRIOTITY_CODE	Incident priority code It is the internal value defined in Enterprise Manager. PRIORITY_URGENT PRIORITY_VERY_HIGH PRIORITY_HIGH PRIORITY_MEDIUM PRIORITY_LOW PRIORITY_NONE
LAST_UPDATED_TIME	Last updated time
SR_ID	Oracle Service Request Id, if it exists.
BUD_ID	Oracle Bug ID, if an associated bug exists.

Environment Variables Common to Incident and Problem Notifications

An incident or problem may be associated with multiple event sources. An event source can be a Target, a Source Object, or both.

4.2.1.1.1 Environment Variables Related to Event Sources

Number of event sources will be set in EVENT_SOURCE_COUNT environment variable. Using the EVENT_SOURCE_COUNT information, a script can be written to loop through the relevant environment variables to fetch the information about multiple event sources. Environment variables for all event sources are prefixed with EVENT_SOURCE_. Environment variables for source objects are suffixed with SOURCE_<attribute_name> . For example, EVENT_SOURCE_1_SOURCE_TYPE provides the source object type of first event source. Environment variables for a target are suffixed with TARGET_<attribute_name>. For example, EVENT_SOURCE_1_ TARGET_NAME provides the target name of first event source.

The following table lists the environment variables for source object of x-th Event Source.

Table 4–17 Source Object of the x-th Event Source

Environment Variable	Description
EVENT_SOURCE_x_ SOURCE_GUID	Source Object GUID.
EVENT_SOURCE_x_ SOURCE_TYPE	Source Object Type
EVENT_SOURCE_x_ SOURCE_NAME	Source Object Name.
EVENT_SOURCE_x_ SOURCE_OWNER	Source Object Owner.
EVENT_SOURCE_x_ SOURCE_SUB_TYPE	Source Object Sub-Type.
EVENT_SOURCE_x_ SOURCE_URL	Source Object URL to EM console.

The following table lists the environment variables for target of *x*th Event Source.

Table 4–18 Target of x-th Event Source

Environment Variable	Description
EVENT_SOURCE_x_ TARGET_GUID	Target GUID
EVENT_SOURCE_x_ TARGET_NAME	Target name
EVENT_SOURCE_x_ TARGET_OWNER	Target Owner
EVENT_SOURCE_x_ TARGET_VERSION	Target version
EVENT_SOURCE_x_ TARGET_LIFE_CYCLE_ STATUS	Target life cycle status
EVENT_SOURCE_x_ TARGET_TYPE	Target Type
EVENT_SOURCE_x_ HOST_NAME	Target Host Name
EVENT_SOURCE_x_ TARGET_URL	Target URL to EM Console.

4.2.1.1.2 Script Examples

The sample OS script shown in Example 4–6 appends environment variable entries to a log file. In this example, the script logs a severity occurrence to a file server. If the file server is unreachable then an exit code of 100 is returned to force the Oracle Management Service Notification System to retry the notification

Example 4–6 Sample OS Command Script

```
#!/bin/ksh

LOG_FILE=/net/myhost/logs/event.log
if test -f $LOG_FILE
then
```

```
echo $TARGET_NAME $MESSAGE $EVENT_REPORTE_TIME >> $LOG_FILE
else
   exit 100
fi
```

Example 4–7 shows an OS script that logs alert information for both incidents and events to the file 'oscmdNotify.log'. The file is saved to the /net/myhost/logs directory.

Example 4–7 Alert Logging Scripts

```
#!/bin/sh
#
  LOG_FILE=/net/myhost/logs/oscmdNotify.log

  echo '-------------' >> $LOG_FILE

  echo 'issue_type=' $ISSUE_TYPE >> $LOG_FILE
  echo 'notif_type=' $NOTIF_TYPE >> $LOG_FILE
  echo 'message=' $MESSAGE >> $LOG_FILE
  echo 'message_url'  = $MESSAGE_URL >>$LOG_FILE
  echo 'severity=' $SEVERITY >> $LOG_FILE
  echo 'severity_code'  = $SEVERITY_CODE >>$LOG_FILE
  echo 'ruleset_name=' $RULESET_NAME >> $LOG_FILE
  echo 'rule_name=' $RULE_NAME >> $LOG_FILE
  echo 'rule_owner=' $RULE_OWNER >> $LOG_FILE
  echo 'repeat_count=' $REPEAT_COUNT >> $LOG_FILE
  echo 'categories_count'  = $CATEGORIES_COUNT >>$LOG_FILE
  echo 'category_1'  = $CATEGORY_1 >>$LOG_FILE
  echo 'category_2'  = $CATEGORY_2 >>$LOG_FILE
  echo 'category_code_1'  = $CATEGORY_CODE_1 >>$LOG_FILE
  echo 'category_code_2'  = $CATEGORY_CODE_2 >>$LOG_FILE
  echo 'category_codes_count'  = $CATEGORY_CODES_COUNT >>$LOG_FILE

# event
if [ $ISSUE_TYPE -eq 1 ]
then
  echo 'host_name=' $HOST_NAME >> $LOG_FILE
  echo 'event_type=' $EVENT_TYPE >> $LOG_FILE
  echo 'event_name=' $EVENT_NAME >> $LOG_FILE
  echo 'event_occurrence_time=' $EVENT_OCCURRENCE_TIME >> $LOG_FILE
  echo 'event_reported_time=' $EVENT_REPORTED_TIME >> $LOG_FILE
  echo 'sequence_id=' $SEQUENCE_ID >> $LOG_FILE
  echo 'event_type_attrs=' $EVENT_TYPE_ATTRS >> $LOG_FILE
  echo 'source_obj_name=' $SOURCE_OBJ_NAME >> $LOG_FILE
  echo 'source_obj_type=' $SOURCE_OBJ_TYPE >> $LOG_FILE
  echo 'source_obj_owner=' $SOURCE_OBJ_OWNER >> $LOG_FILE
  echo 'target_name'  = $TARGET_NAME >>$LOG_FILE
  echo 'target_url'  = $TARGET_URL >>$LOG_FILE
  echo 'target_owner=' $TARGET_OWNER >> $LOG_FILE
  echo 'target_type=' $TARGET_TYPE >> $LOG_FILE
  echo 'target_version=' $TARGET_VERSION >> $LOG_FILE
  echo 'lifecycle_status=' $TARGET_LIFECYCLE_STATUS >> $LOG_FILE
  echo 'assoc_incident_escalation_level'  = $ASSOC_INCIDENT_ESCALATION_LEVEL
>>$LOG_FILE
  echo 'assoc_incident_id'  = $ASSOC_INCIDENT_ID >>$LOG_FILE
  echo 'assoc_incident_owner'  = $ASSOC_INCIDENT_OWNER >>$LOG_FILE
  echo 'assoc_incident_acknowledged_by_owner'  = $ASSOC_INCIDENT_ACKNOWLEDGED_BY_
OWNER >>$LOG_FILE
  echo 'assoc_incident_acknowledged_details'  = $ASSOC_INCIDENT_ACKNOWLEDGED_
```

```
DETAILS >>$LOG_FILE
  echo 'assoc_incident_priority'  = $ASSOC_INCIDENT_PRIORITY >>$LOG_FILE
  echo 'assoc_incident_status'  = $ASSOC_INCIDENT_STATUS >>$LOG_FILE
  echo 'ca_job_status'  = $CA_JOB_STATUS >>$LOG_FILE
  echo 'event_context_attrs'  = $EVENT_CONTEXT_ATTRS >>$LOG_FILE
  echo 'last_updated_time'  = $LAST_UPDATED_TIME >>$LOG_FILE
  echo 'sequence_id'  = $SEQUENCE_ID >>$LOG_FILE
  echo 'test_date_attr_noref'  = $TEST_DATE_ATTR_NOREF >>$LOG_FILE
  echo 'test_raw_attr_noref'  = $TEST_RAW_ATTR_NOREF >>$LOG_FILE
  echo 'test_str_attr1'  = $TEST_STR_ATTR1 >>$LOG_FILE
  echo 'test_str_attr2'  = $TEST_STR_ATTR2 >>$LOG_FILE
  echo 'test_str_attr3'  = $TEST_STR_ATTR3 >>$LOG_FILE
  echo 'test_str_attr4'  = $TEST_STR_ATTR4 >>$LOG_FILE
  echo 'test_str_attr5'  = $TEST_STR_ATTR5 >>$LOG_FILE
  echo 'test_str_attr_ref'  = $TEST_STR_ATTR_REF >>$LOG_FILE
  echo 'total_occurrence_count'  = $TOTAL_OCCURRENCE_COUNT >>$LOG_FILE
fi

# incident
if [ $ISSUE_TYPE -eq 2 ]
then
  echo 'action_msg=' $ACTION_MSG >> $LOG_FILE
  echo 'incident_id=' $INCIDENT_ID >> $LOG_FILE
  echo 'incident_creation_time=' $INCIDENT_CREATION_TIME >> $LOG_FILE
  echo 'incident_owner=' $INCIDENT_OWNER >> $LOG_FILE
echo 'incident_acknowledged_by_owner'  = $INCIDENT_ACKNOWLEDGED_BY_OWNER >>$LOG_
FILE
  echo 'incident_status'  = $INCIDENT_STATUS >>$LOG_FILE
  echo 'last_modified_by=' $LAST_MODIFIED_BY >> $LOG_FILE
  echo 'last_updated_time=' $LAST_UPDATED_TIME >> $LOG_FILE
  echo 'assoc_event_count=' $ASSOC_EVENT_COUNT >> $LOG_FILE
  echo 'adr_incident_id=' $ADR_INCIDENT_ID >> $LOG_FILE
  echo 'occurrence_count=' $OCCURRENCE_COUNT >> $LOG_FILE
  echo 'escalated=' $ESCALATED >> $LOG_FILE
  echo 'escalated_level=' $ESCALATED_LEVEL >> $LOG_FILE
  echo 'priority=' $PRIORITY >> $LOG_FILE
  echo 'priority_code'  = $PRIORITY_CODE >>$LOG_FILE
  echo 'ticket_id=' $TICKET_ID >> $LOG_FILE
  echo 'ticket_status=' $TICKET_STATUS >> $LOG_FILE
  echo 'ticket_url=' $TICKET_ID_URL >> $LOG_FILE
  echo 'total_duplicate_count=' $TOTAL_DUPLICATE_COUNT >> $LOG_FILE
  echo 'source_count=' $EVENT_SOURCE_COUNT >> $LOG_FILE
  echo 'event_source_1_host_name'  = $EVENT_SOURCE_1_HOST_NAME >>$LOG_FILE
  echo 'event_source_1_target_guid'  = $EVENT_SOURCE_1_TARGET_GUID >>$LOG_FILE
  echo 'event_source_1_target_name'  = $EVENT_SOURCE_1_TARGET_NAME >>$LOG_FILE
  echo 'event_source_1_target_owner'  = $EVENT_SOURCE_1_TARGET_OWNER >>$LOG_FILE
  echo 'event_source_1_target_type'  = $EVENT_SOURCE_1_TARGET_TYPE >>$LOG_FILE
  echo 'event_source_1_target_url'  = $EVENT_SOURCE_1_TARGET_URL >>$LOG_FILE
  echo 'event_source_1_target_lifecycle_status'  = $EVENT_SOURCE_1_TARGET_
LIFECYCLE_STATUS >>$LOG_FILE
  echo 'event_source_1_target_version'  = $EVENT_SOURCE_1_TARGET_VERSION >>$LOG_
FILE
fi
exit 0
```

Example 4–8 shows a script that sends an alert to an HP OpenView console from Enterprise Manager Cloud Control. When a metric alert is triggered, the Enterprise Manager Cloud Control displays the alert. The HP OpenView script is then called, invoking opcmsg and forwarding the information to the HP OpenView management server.

Example 4–8 HP OpenView Script

```
/opt/OV/bin/OpC/opcmsg severity="$SEVERITY" app=OEM msg_grp=Oracle msg_
text="$MESSAGE" object="$TARGET_NAME"
```

4.2.1.1.3 Migrating pre-12c OS Command Scripts This section describes how to map pre-12c OS Command notification shell environment variables to 12c OS Command shell environment variables.

Please note that Policy Violations are no longer supported beginning with the 12c release.

Migrating Metric Alert Event Types

Following table is the mapping for the OS Command shell environment variables when the event_type is metric_alert.

Table 4–19 pre-12c/12c metric_alert Environment Variable Mapping

Pre-12c Environment Variable	Corresponding 12c Environment Variables
ACKNOWLEDGED	ASSOC_INCIDENT_ACKNOWLEDGED_BY_OWNER
ACKNOWLEDGED_BY	ASSOC_INCIDENT_OWNER
CYCLE_GUID	CYCLE_GUID
HOST	HOST_NAME
KEY_VALUE	Note: See detail description below.
KEY_VALUE_NAME	Note: See detail description below
MESSAGE	MESSAGE
METRIC	METRIC_COLUMN
NOTIF_TYPE	NOTIF_TYPE; use the map in section 2.3.5
REPEAT_COUNT	REPEAT_COUNT
RULE_NAME	RULESET_NAME
RULE_OWNER	RULE_OWNER
SEVERITY	SEVERITY
TARGET_NAME	TAGER_NAME
TARGET_TYPE	TARGET_TYPE
TIMESTAMP	EVENT_REPORTED_TIME
VIOLATION_CONTEXT	EVENT_CONTEXT_ATTRS
VIOLATION_GUID	SEVERITY_GUID
POLICY_RULE	No mapping, obsolete in NG release.

To get KEY_VALUE_NAME and KEY_VALUE from NG, perform the following steps.

- If $NUM_KEYS variable is null, then $KEY_VALUE_NAME and $KEY_VALUE are null.
- If $NUM_KEYS equals 1

  KEY_VALUE_NAME=$KEY_COLUMN_1

  KEY_VALUE=$KEY_VALUE_1_VALUE

- If $NUM_KEYS is greater than 1

  KEY_VALUE_NAME="$KEY_COLUMN_1;$KEY_COLUMN_2;..;KEY_COLUMN_x"

  KEY_VALUE="$KEY_COLUMN_1_VALUE;$KEY_COLUMN_2_VALUE;..;KEY_COLUMN_x_VALUE "

  Where x is the value of $NUM_KEYS and ";" is the separator.

Migrating Target Availability Event Types

Following table is the mapping for the OS Command shell environment variables when the event_type is 'target_availability'.

Table 4–20 pre-12c/12c target_availability Environment Variable Mappings

Pre-12c Environment Variable	Corresponding 12c Environment Variables
TARGET_NAME	TARGET_NAME
TARGET_TYPE	TARGET_TYPE
METRIC	Status
CYCLE_GUID	CYCLE_GUID
VIOLATION_CONTEXT	EVENT_CONTEXT_ATTRS
SEVERITY	TARGET_STATUS
HOST	HOST_NAME
MESSAGE	MESSAGE
NOTIF_TYPE	NOTIF_TYPE; use the map in section 2.3.5
TIMESTAMP	EVENT_REPORTED_TIME
RULE_NAME	RULESET_NAME
RULE_OWNER	RULE_OWNER
REPEAT_COUNT	REPEAT_COUNT
KEY_VALUE	""
KEY_VALUE_NAME	""

Migrating Job Status Change Event Types

Following table is the mapping for the OS Command shell environment variables when the event_type is 'job_status_change'.

Table 4–21 pre-12c/12c job_status_change Environment Variable Mappings

Pre-12c Environment Variable	Corresponding 12c Environment Variables
JOB_NAME	SOURCE_OBJ_NAME
JOB_OWNER	SOURCE_OBJ_OWNER
JOB_TYPE	SOURCE_OBJ_SUB_TYPE
JOB_STATUS	EXECUTION_STATUS
NUM_TARGETS	1 if $ TARGET_NAME is not null, 0 otherwise
TARGET_NAME1	TARGET_NAME

Table 4–21 (Cont.) pre-12c/12c job_status_change Environment Variable Mappings

Pre-12c Environment Variable	Corresponding 12c Environment Variables
TARGET_TYPE1	TARGET_TYPE
TIMESTAMP	EVENT_REPORTED_TIME
RULE_NAME	RULESET_NAME
RULE_OWNER	RULE_OWNER

Migrating Corrective Action-Related OS Scripts

Refer to section "Migrating Metric Alert Event Types" for mapping the following environment variables while receiving notifications for corrective actions.

KEY_VALUE

KEY_VALUE_NAME

METRIC

METRIC_VALUE

RULE_NAME

RULE_OWNER

SEVERITY

TIMESTAMP

VIOLATION_CONTEXT

Use the map below for mapping other environment variables.

Table 4–22 pre-12c/12c Corrective Action Environment Variable Mappings

Pre-12c Environment Variable	Corresponding 12c Environment Variables
NUM_TARGETS	1
TARGET_NAME1	TAGER_NAME
TARGET_TYPE1	TARGET_TYPE
JOB_NAME	CA_JOB_NAME
JOB_OWNER	CA_JOB_OWNER
JOB_STATUS	CA_JOB_STATUS
JOB_TYPE	CA_JOB_TYPE

Notification Type Mapping

Table 4–23 pre-12c/12c notif_type Mappings

notif_type (12c)	notif_type (Pre-12c)
NOTIF_NORMAL	1
NOTIF_REPEAT	4
NOTIF_DURATION	9
NOTIF_RETRY	2

4.2.1.2 Adding a Notification Method Based on a PL/SQL Procedure

A user-defined PL/SQL procedure can receive notifications for matching events, incidents and problem.

> **Note:** PL/SQL procedures used with pre-12c versions of Enterprise Manager will continue to work without modification. However, you should update the procedures to use the new signatures.

Complete the following four steps to define a notification method based on a PL/SQL procedure.

Step 1: Define the PL/SQL procedure.

The procedure must have one of the following signatures depending on the type of notification that will be received.

For Events:

PROCEDURE event_proc(event_msg IN gc$notif_event_msg)

For Incidents:

PROCEDURE incident_proc(incident_msg IN gc$notif_incident_msg)

For Problems:

PROCEDURE problem_proc(problem_msg IN gc$notif_problem_msg)

> **Note:** The notification method based on a PL/SQL procedure must be configured by an administrator with Super Administrator privileges before a user can select it while creating/editing a incident rule.

For more information on passing specific types of information to scripts or PL/SQL procedures, see the following sections:

"Passing Information to a PL/SQL Procedure" on page 4-34

"Passing Corrective Action Status Change Information" on page 4-49

"Passing Job Execution Status Information" on page 4-51

Step 2: Create the PL/SQL procedure on the Management Repository.

Create the PL/SQL procedure on the repository database using one of the following procedure specifications:

PROCEDURE event_proc(event_msg IN gc$notif_event_msg)

PROCEDURE incident_proc(incident_msg IN gc$notif_incident_msg)

PROCEDURE problem_proc(problem_msg IN gc$notif_problem_msg)

The PL/SQL procedure must be created on the repository database using the database account of the repository owner (such as SYSMAN)

If an error is encountered during the running of the procedure, the Notification System can be instructed to retry the sending of the notification to the procedure by raising a user-defined exception that uses the error code -20000. The procedure initially retried

after one minute, then two minutes, then three minutes and so on, until the notification is a day old, at which point it will be purged.

Step 3: Register your PL/SQL procedure as a new notification method.

Log in as a Super Administrator. From the **Setup** menu, choose **Notifications** and then **Notification Methods** to access the Notification Methods page. From this page, you can define a new notification based on 'PL/SQL Procedure'. See Section 4.2.1.2, "Adding a Notification Method Based on a PL/SQL Procedure".

Make sure to use a fully qualified name that includes the schema owner, package name and procedure name. The procedure will be executed by the repository owner and so the repository owner must have execute permission on the procedure.

Create a notification method based on your PL/SQL procedure. The following information is required when defining the method:

- Name
- Description
- PL/SQL Procedure

You must enter a fully qualified procedure name (for example, OWNER.PKGNAME.PROCNAME) and ensure that the owner of the Management Repository has execute privilege on the procedure.

An example of the required information is shown in Example 4–9.

Example 4–9 PL/SQL Procedure Required Information

```
Name Open trouble ticket
Description Notification method to open a trouble ticket in the event
PLSQL Procedure ticket_sys.ticket_ops.open_ticket
```

Step 4: Assign the notification method to an incident rule.

You can edit an existing rule (or create a new incident rule). From the **Setup** menu, select **Incidents** and then select **Incident Rules**. The Incident Rules page displays. From here, you can add an action to a rule specifying the new PL/SQL procedure found under **Advanced Notification Method.**

There can be more than one PL/SQL-based method configured for your Enterprise Manager environment.

Information about how incident, event, and problem information is passed to the PLSQL procedure is covered in the next section.

Passing Information to a PL/SQL Procedure

Passing event, incident, and problem information (payload) to PL/SQL procedures allows you to customize automated responses to these conditions. All 3 types of notification payloads have a common element - gc$notif_msg_info. It provides generic information that applies to all types of notifications. In addition, each of the 3 payloads have one specific element that provides the payload specific to the given issue type.

gc$notif_event_msg (payload for event notifications)

gc$notif_event_msg contains two objects - event payload object and message information object.

Table 4–24 Event Notification Payload

Attribute	Datatype	Additional Information
event_payload	gc$notif_event_payload	Event notification payload. See gc$notif_event_payload type definition for detail.
msg_info	gc$notif_msg_info	Notification message. See gc$notif_msg_info definition for detail.

gc$notif_incident_msg (payload for incident notifications)

gc$notif_incident_msg type contains two objects - incident payload and message information. This object represents the delivery payload for Incident notification message, contains all data associated with Incident notification, and can be accessed by user's custom PL/SQL procedures.

Table 4–25 Incident Notification Payload

Attribute	Datatype	Additional Information
incident_payload	gc$notif_incident_payload	Incident notification payload. See gc$notif_incident_payload type definition for detail.
msg_info	gc$notif_msg_info	Envelope level notification information. See gc$notif_msg_info type definition for detail.

gc$notif_problem_msg (payload for problem notifications)

This object represents the delivery payload for Problem notification message, contains all data associated with problem notification, and can be accessed by a user's custom PL/SQL procedures.

Table 4–26 Problem Notification Payload

Attribute	Datatype	Additional Information
problem_payload	gc$notif_problem_payload	Problem notification payload. See gc$notif_problem_payload type definition for detail.
msg_info	gc$notif_msg_info	Notification message. See gc$notif_msg_info type definition for detail.

gc$notif_msg_info (common for event/incident/problem payloads)

This object contains the generic notification information including notification_type, rule set and rule name, etc. for Event, Incident or Problem delivery payload.

Table 4–27 Event, Incident, Problem Common Payload

Attribute	Datatype	Description
notification_type	VARCHAR2(32)	Type of notification, can be one of the following values. GC$NOTIFICATION.NOTIF_NORMAL GC$NOTIFICATION.NOTIF_RETRY GC$NOTIFICATION.NOTIF_REPEAT GC$NOTIFICATION.NOTIF_DURATION GC$NOTIFICATION.NOTIF_CA GC$NOTIFICATION.NOTIF_RCA
repeat_count	NUMBER	Repeat notification count
ruleset_name	VARCHAR2(256)	Name of the rule set that triggered the notification
rule_name	VARCHAR2(256)	Name of the rule that triggered the notification
rule_owner	VARCAH2(256)	EM User who owns the rule set
message	VARCHAR2(4000)	Message about event/incident/problem.
message_url	VARCHAR2(4000)	Link to the Enterprise Manager console page that provides the details of the event/incident/problem.

gc$notif_event_payload (payload specific to event notifications)

This object represents the payload specific to event notifications.

Table 4–28 Common Payloads for Events, Incidents, and Problems

Attribute	Datatype	Additional Information
event_instance_guid	RAW(16)	Event instance global unique identifier.
event_sequence_guid	RAW(16)	Event sequence global unique identifier.
Target	gc$notif_target	Related Target Information object. See gc$notif_target type definition for detail.
Source	gc$notif_source	Related Source Information object, that is not a target. See gc$notif_source type definition for detail.
event_attrs	gc$notif_event_attr_array	The list of event specified attributes. See gc$notif_event_attr type definition for detail.
corrective_action	gc$notif_corrective_action_job	Corrective action information, optionally populated when corrective action job execution has completed.
event_type	VARCHAR2(20)	Event type - example: Metric Alert.
event_name	VARCHAR2(512)	Event name.
event_msg	VARCHAR2(4000)	Event message.
reported_date	DATE	Event reported date.
Occurrence_date	DATE	Event occurrence date.
Severity	VARCHAR2(128)	Event Severity. It is the translated severity name.
severity_code	VARCHAR2(32)	Event Severity code. It is the internal severity name used in Enterprise Manager.

Table 4–28 (Cont.) Common Payloads for Events, Incidents, and Problems

Attribute	Datatype	Additional Information
assoc_incident	gc$notif_issue_summary	Summary of associated incident. It is populated if the event is associated with an incident. See gc$notif_issue_summary type definition for detail
action_msg	VARCHAR2(4000)	Message describing the action to take for resolving the event.
rca_detail	VARCHAR2(4000)	Root cause analysis detail. The size of RCA details output is limited to 4000 characters long.
event_context_data	gc$notif_event_context_array	Event context data. See gc$notif_event_context type definition for detail.
categories	gc$category_string_array	List of categories that the event belongs to. Category is translated based on locale defined in OMS server. Notification system sends up to 10 categories.
category_codes	gc$category_string_array	Codes for the categories. The size of array is up to 10.

gc$notif_incident_payload (payload specific to incident notifications)

It contains the incident specific attributes, associated problem and ticket information.

Table 4–29 Incident Notification Payloads

Attribute	Datatype	Additional Information
incident_attrs	gc$notif_issue_attrs	Incident specific attributes. See gc$notif_issue_attrs type definition for detail.
assoc_event_count	NUMBER	The total number of events associated with this incident.
ticket_status	VARCHAR2(64)	The status of external Ticket, if it exists.
ticket_id	VARCHAR2(128)	The ID of external Ticket, if it exists.
ticket_url	VARCHAR2(4000)	The URL for external Ticket, if it exists.
assoc_problem	gc$notif_issue_summary	Summary of the problem, if it has an associated problem. See gc$notif_issue_summary type definition for detail.

gc$notif_problem_payload (payload specific to problems)

It contains problem specific attributes, key, Service Request(SR) and Bug information.

Table 4–30 Problem Payload

Attribute	Datatype	Additional Information
problem_attrs	gc$notif_issue_attrs	Problem specific attributes. See gc$notif_issue_attrs type definition for detail.
problem_key	VARCHAR2(850)	Problem key if it is generated.
assoc_incident_count	NUMBER	Number of incidents associated with this problem.

Table 4–30 (Cont.) Problem Payload

Attribute	Datatype	Additional Information
sr_id	VARCHAR2(64)	Oracle Service Request Id, if it exists.
sr_url	VARCHAR2(4000)	URL for Oracle Service Request, if it exists.
bug_id	VARCHAR2(64)	Oracle Bug ID, if an associated bug exists.

gc$notif_issue_attrs (payload common to incidents and problems)

It provides common details for incident and problem. It contains details such as id, severity, priority, status, categories, acknowledged by owner, and source information associated with.

Table 4–31 Payload Common to Incidents and Problems

Attribute	Datatype	Additional Information
Id	NUMBER(16)	ID of the incident or problem.
Severity	VARCHAR2(128)	Issue Severity. It is the translated.
severity_code	VARCHAR2(32)	Issue Severity Code.The possible values are defined in descending order of severity: GC$EVENT.FATAL GC$EVENT.CRITICAL GC$EVENT.WARNING GC$EVENT.MINOR_WARNING GC$EVENT.INFORMATIONAL GC$EVENT.CLEAR
priority	VARCHAR2(128)	Issue Priority. It is the translated priority name.
priority_code	VARCHAR2(32)	Issue Priority. It is the internal value defined in EM. The possible values are defined in descending order of priority: GC$EVENT.PRIORITY_URGENT GC$EVENT.PRIORITY_VERY_HIGH GC$EVENT.PRIORITY_HIGH GC$EVENT.PRIORITY_MEDIUM GC$EVENT.PRIORITY_LOW GC$EVENT.PRIORITY_NONE
status	VARCHAR2(32)	Status of Issue. The possible values are GC$EVENT.STATUS_NEW GC$EVENT.STATUS_CLOSED Any other user defined status.
escalation_level	NUMBER(1)	Escalation level of the issue, has a value between 0 to 5.
owner	VARCHAR(256)	Issue Owner. Set to NULL if no owner exists.
acknowledged_by_owner	NUMBER(1)	Set to 1, if this issue was acknowledged by owner.
creation_date	DATE	Issue creation date.

Table 4–31 (Cont.) Payload Common to Incidents and Problems

Attribute	Datatype	Additional Information
closed_date	DATE	Issue closed date, null if not closed.
categories	gc$category_string_ array	List of categories that the event belongs to. Category is translated based on locale defined in OMS server. Notification system sends up to 10 categories.
category_codes	gc$category_string_ array	Codes for the categories. Notification system sends up to 10 category codes.
source_info_arr	gc$notif_source_ info_array	Array of source information associated with this issue. See $gcnotif_source_info type definition for detail.
last_modified_by	VARCHAR2(256)	Last modified by user.
last_updated_date	DATE	Last updated date.

gc$notif_issue_summary (common to event and incident payloads)

This object represents the associated incident summary in Event payload, or associated problem summary in Incident payload, respectively.

Table 4–32 Payload

Attribute	Datatype	Additional Information
id	NUMBER	Issue Id, either Incident Id or Problem Id.
severity	VARCHAR(128)	The severity level of an issue. It is translated severity name.
severity_code	VARCHAR2(32)	Issue Severity Code, has one of the following values. GC$EVENT.FATAL GC$EVENT.CRITICAL GC$EVENT.WARNING GC$EVENT.MINOR_WARNING GC$EVENT.INFORMATIONAL GC$EVENT.CLEAR
priority	VARCHAR2(128)	Current priority. It is the translated priority name.
priority_code	VARCHAR2(32)	Issue priority code, has one of the following values. GC$EVENT.PRIORITY_URGENT GC$EVENT.PRIORITY_VERY_HIGH GC$EVENT.PRIORITY_HIGH GC$EVENT.PRIORITY_MEDIUM GC$EVENT.PRIORITY_LOW GC$EVENT.PRIORITY_NONE

Table 4–32 (Cont.) Payload

Attribute	Datatype	Additional Information
status	VARCHAR2(64)	Status of issue. The possible values are GC$EVENT.STATUS_NEW GC$EVENT.STATUS_CLOSED GC$EVENT.WIP (work in progress) GC$EVENT.RESOLVED any other user defined status
escalation_level	VARCHAR2(2)	Issue escalation level range from 0 to 5, default 0.
owner	VARCHAR2(256)	Issue Owner. Set to NULL if no owner exists.
acknowledged_by_owner	NUMBER(1)	Set to 1, if this issue was acknowledged by owner.

gc$category_string_array

gc$category_string_array is an array of string containing the categories which event, incident or problem is associated with. Note that notification system delivers up to 10 categories.

gc$notif_event_context_array

gc$notif_event_context_array provides information about the additional diagnostic data that was captured at event detection time. Note that notification system delivers up to 200 elements from the captured event context.

CREATE OR REPLACE TYPE gc$notif_event_context_array AS VARRAY(200) OF gc$notif_event_context;

gc$notif_event_context

This object represents the detail of event context data which is additional contextual information that is captured by the source system at the point of event generation that could be useful for diagnosis. The context for an event should consist of set of keys and values along with data type (Number or String only).

Table 4–33 Event Context Type

Attribute	Datatype	Additional Information
Name	VARCHAR2(256)	The event context name.
Type	NUMBER(1)	The data type of the value, which is stored (0) - for numeric data (1) - for string data.
Value	NUMBER	The numerical value.
string_value	VARCHAR2(4000)	The string value.

gc$notif_corrective_action_job

This object provides information about the execution of a corrective action job. Note that the corrective actions are supported for metric alert and target availability events only.

Table 4–34 Corrective Action Job-Specific Attributes

Attribute	Datatype	Additional Information
job_guid	RAW(16)	Corrective Action Job global unique identifier.
job_name	VARCHAR2(128)	The value will be the name of the Corrective Action. It applies to Metric Alert and Target Availability Events.
job_owner	VARCHAR2(256)	Corrective action job owner.
job_type	VARCHAR2(256)	Corrective action job type.
job_status	VARCHAR2(64)	Corrective action job execution status.
job_status_code	NUMBER	Corrective action job execution status code. It is the internal value defined in EM.
job_step_output	VARCHAR2(4000)	The value will be the text output from the Corrective Action execution. This will be truncated to last 4000 characters.
job_execution_guid	RAW(16)	Corrective Action Job execution global unique identifier.
job_state_change_guid	RAW(16)	Corrective Action Job change global unique identifier.
occurred_date	DATE	Corrective action job occurred date.

gc$notif_source_info_array

It is used providing access to the multiple sources that an incident or a problem could be related to. NOTE: The notification system delivers up to 200 sources associated with an incident or a problem.

CREATE OR REPLACE TYPE gc$notif_source_info_array AS VARRAY(200) OF gc$notif_source_info;

gc$notif_source_info

Notification Source Information which is used for referencing Source information containing either target or source, or both.

Table 4–35 Source Information Type

Attribute	Datatype	Additional Information
target	gc$notif_target	It is populated when the event is related to a target. See gc$notif_target type definition for detail.
Source	gc$notif_source	It is populated when the event is related to a (non-target) source. See gc$notif_source type definition for detail.

gc$notif_source

Source object is used for referencing source objects other than a job target.

Table 4–36 Payload

Attribute	Datatype	Additional Information
source_guid	RAW(16)	Source's global unique identifier.
source_type	VARCHAR2(120)	Type of the Source object, e.g., TARGET, JOB, TEMPLATE, etc.
source_name	VARCHAR2(256)	Source Object Name.
source_owner	VARCHAR2(256)	Owner of the Source object.
source_sub_type	VARCHAR2(256)	Sub-type of the Source object, for example, within the TARGET these would be the target types like Host, Database etc.
source_url	VARCHAR2(4000)	Source's event console URL.

gc$notif_target

Target information object is used for providing target information.

Table 4–37 Target Information

Attribute	Datatype	Additional Information
target_guid	RAW(16)	Target's global unique identifier.
target_name	VARCHAR2(256)	Name of target.
target_owner	VARCHAR2(256)	Owner of target.
target_lifecycle_status	VARCHAR2(1024)	Life Cycle Status of the target.
target_version	VARCHAR2(64)	Target Version of the target.
target_type	VARCHAR2(128)	Type of a target.
target_timezone	VARCHAR2(64)	Target's regional time zone.
host_name	VARCHAR2(256)	The name of the host on which the target is deployed upon.
target_url	VARCHAR2(4000)	Target's EM Console url.
udtp_array	gc$notif_udtp_array	The list of user defined target properties. It is populated for events that are associated with a target. It is populated for incidents and problems, when they are associated with a single source (gc$notif_source_info).

gc$notif_udtp_array

It is array of gc$notif_udtp type and size is up to 20.

CREATE OR REPLACE TYPE gc$notif_udtp_array AS VARRAY(20) OF gc$notif_udtp;

gc$notif_udtp

User defined Target Properties (UDTP) is used for referencing User defined target properties. UDTP should consist of set of property key name and property value.

Table 4–38 Payload

Attribute	Datatype	Additional Information
name	VARCHAR2(64),	The name of property.

Table 4–38 (Cont.) Payload

Attribute	Datatype	Additional Information
value	VARCHAR2(1024)	Property value.
label	VARCHAR(256)	Property label.
nls_id	VARCHAR(64)	Property nls id

gc$notif_event_attr_array

Array of gc$notif_event_attr is used for referencing Event specific attributes, and its size is up to 35.

CREATE OR REPLACE TYPE gc$notif_event_attr_array AS VARRAY(35) OF gc$notif_event_attr;

gc$notif_event_attr

It is used for referencing Event type specific attributes.

Table 4–39 Event Attribute Type

Attribute	Datatype	Additional Information
name	VARCHAR2(64)	The internal name of event type specific attribute.
value	VARCHAR2(4000)	value.
nls_value	VARCHAR2(4000)	Translated value for the attribute.

4.2.1.2.1 Migrating Pre-12c PL/SQL Advanced Notification Methods

Pre-12c notifications map to event notifications in Enterprise Manager 12c. Event types metric_alert, target_availability and job_status_alert correspond to the pre-12c notification functionality. Note that policy violations functionality is removed for this release.

This section describes the mapping between Enterprise Manager 12c PL/SQL notification payload to the pre-12c PL/SQL notification payload. You can use this information for updating the existing pre-12c PL/SQL user callback procedures to use the 12c PL/SQL notification payload.

Please note that Policy Violations are no longer supported in the 12c release.

Mapping for MGMT_NOTIFY_SEVERITY

When event type is metric_alert

Use the following map when gc$notif_event_payload .event_type='metric_alert'.

Table 4–40 Metric Alert Mapping

MGMT_NOTIFY_SEVERITY	12c Notification Payload
TARGET_NAME	gc$notif_target.target_name
TARGET_TYPE	gc$notif_target.target_type
TIMEZONE	gc$notif_target.target_timezone

Table 4–40 (Cont.) Metric Alert Mapping

MGMT_NOTIFY_SEVERITY	12c Notification Payload
HOST_NAME	gc$notif_target.host_name
MERTIC_NAME	gc$notif_event_attr.value where its name=' metric_group' in gc$notif_event_attr_array.
METRIC_DESCRIPTION	gc$notif_event_attr.value where its name=' metric_description' in gc$notif_event_attr_ array.
METRIC_COLUMN	gc$notif_event_attr.value where its name=' metric_column' in gc$notif_event_attr_array.
METRIC_VALUE	gc$notif_event_attr.value where its name=' value' in gc$notif_event_attr_array.
KEY_VALUE	It is applied for multiple keys based metric when value of gc$notif_event_ attr.name='num_keys' is not null and is greater than 0 in gc$notif_event_attr_array. See detail descriptions below.
KEY_VALUE_NAME	It is applied for multiple keys based metric when value of gc$notif_event_ attr.name='num_keys' is not null and is greater than 0 in gc$notif_event_attr_array. See detail descriptions below.
KEY_VALUE_GUID	gc$notif_event_attr.value where its name='key_ value' in gc$notif_event_attr_ array.
CTXT_LIST	gc$notif_event_context_array
COLLECTION_ TIMESTAMP	gc$notif_event_payload. reported_date
SEVERITY_CODE	Derive from gc$notif_event_ payload.severity_code, see section 1.2.1.1.2 for the mapping over the value.
MESSAGE	gc$notif_msg_info.message
SEVERITY_GUID	gc$notif_event_attr.value where its name=' severity_guid' in gc$notif_event_attr_array.
METRIC_GUID	gc$notif_event_attr.value where its name=' metric_guid_id' in gc$notif_event_attr_array.
TARGET_GUID	gc$notif_target.target_guid
RULE_OWNER	gc$notif_msg_info.rule_owner
RULE_NAME	gc$notif_msg_info.ruleset_name

The following example shows you how to obtain similar pre-12c KEY_VALUE and KEY_VALUE_NAME from an Enterprise Manager 12c notification payload.

First, check its gc$notif_event_attr.value where its name='num_keys' in gc$notif_ event_attr_array.

If it is null or its value is equal to 0, then KEY_VALUE and KEY_VALUE_NAME are null for this event.

If it is not null and value is equal to 1, then it is single key metric.

KEY_VALUE_NAME= value of gc$notif_event_attr where name='key_column_1' in gc$notif_event_attr_array.

KEY_VALUE = value of gc$notif_event_attr where name='key_value' in gc$notif_event_attr_array.

For example: METRIC= Filesystem Space Available (%) num_key=1

KEY_VALUE_NAME= Mount Point

KEY_VALUE= /

If it is not null and value is greater than 1, then it multiple keys metric.

KEY_VALUE_NAME = value of gc$notif_event_attr where name='key_column_1' + ";" +

value of gc$notif_event_attr where name='key_column_2' + ";" +

. …. (up to where name ='key_column_num_key')

KEY_VALUE = value of gc$notif_event_attr where name='key_column_1_value' + ";" +

value of gc$notif_event_attr where name='key_column_2_value' + ";" +

. …. (up to where name='key_column_num_key_value')

The ";" is separator between names or values.

For example: METRIC= Program's Max CPU Utilization (%) which num_key=2

KEY_VALUE_NAME= Program Name;Owner

KEY_VALUE= loadcpu;userId

Severity Code mapping from 12c to pre-12c when the event type is metric_alert

Table 4–41 Severity Code Mapping

12c Severity Code	Pre-12c Severity Code
GC_EVENT_RECEIVER.FATAL 32	MGMT_GLOBAL.G_SEVERITY_CRITICAL 25
GC_EVENT_RECEIVER.CRITICAL 16	MGMT_GLOBAL.G_SEVERITY_CRITICAL 25
GC_EVENT_RECEIVER.WARNING 8	MGMT_GLOBAL.G_SEVERITY_WARNING 20
GC_EVENT_RECEIVER.CLEAR 0	MGMT_GLOBAL.G_SEVERITY_CLEAR 15

When event type is target_availability

Use the following map when gc$notif_event_payload .event_type='target_availability'.

Table 4–42 Target Availability Mapping

MGMT_NOTIFY_SEVERITY	12c Notification Payload
TARGET_NAME	gc$notif_target.target_name
TARGET_TYPE	gc$notif_target.target_type
TIMEZONE	gc$notif_target.target_timezone
HOST_NAME	gc$notif_target.host_name
MERTIC_NAME	Use fixed value "Response".

Table 4–42 (Cont.) Target Availability Mapping

MGMT_NOTIFY_SEVERITY	12c Notification Payload
METRIC_DESCRIPTION	NULL
METRIC_COLUMN	Use fixed value "Status".
METRIC_VALUE	gc$notif_event_attr.value where its name='target_status' in gc$notif_event_attr_ array.
KEY_VALUE	NULL
KEY_VALUE_NAME	NULL
KEY_VALUE_GUID	NULL
CTXT_LIST	gc$notif_event_context_array
COLLECTION_TIMESTAMP	gc$notif_event_payload. reported_date
SEVERITY_CODE	gc$notif_event_attr.value where its name=' avail_severity' in gc$notif_event_attr_array.
MESSAGE	gc$notif_msg_info.message
SEVERITY_GUID	gc$notif_event_attr.value where its name=' severity_guid' in gc$notif_event_attr_array.
METRIC_GUID	gc$notif_event_attr.value where its name=' metric_guid_id' in gc$notif_event_attr_array.
TARGET_GUID	gc$notif_target.target_guid
RULE_OWNER	gc$notif_msg_info.rule_owner
RULE_NAME	gc$notif_msg_info.ruleset_name

Mapping for MGMT_NOTIFY_JOB

Use the following map when gc$notif_event_payload .event_type=job_status_change'.

Table 4–43 Job Status Change Mapping

MGMT_NOTIFY_JOB	12c Notification Payload
JOB_NAME	gc$notif_source.source_name
JOB_OWNER	gc$notif_source.source_owner
JOB_TYPE	gc$notif_source.source_sub_type
JOB_STATUS	gc$notif_event_attr.value where its name=' execution_status_code' in gc$notif_event_ attr_array.
STATE_CHANGE_GUID	gc$notif_event_attr.value where its name=' state_change_guid' in gc$notif_event_attr_ array.
JOB_GUID	gc$notif_source.source_guid
EXECUTION_ID	gc$notif_event_attr.value where its name=' execution_id' in gc$notif_event_attr_array.
TARGETS	gc$notif_target.target_name, gc$notif_ target.target_type
RULE_OWNER	gc$notif_msg_info.rule_owner

Table 4–43 (Cont.) Job Status Change Mapping

MGMT_NOTIFY_JOB	12c Notification Payload
RULE_NAME	gc$notif_msg_info.ruleset_name
OCCURRED_DATE	gc$notif_event_payload. reported_date

Mapping for MGMT_NOTIFY_CORRECTIVE_ACTION

Note that corrective action related payload is populated when gc$notif_msg_info.notification_type is set to NOTIF_CA.

For mapping the following attributes, use the mapping information provided for MGMT_NOTIFY_SEVERITY object Table 4–40, " Metric Alert Mapping"

MERTIC_NAME

METRIC_COLUMN

METRIC_VALUE

KEY_VALUE

KEY_VALUE_NAME

KEY_VALUE_GUID

CTXT_LIST

RULE_OWNER

RULE_NAME

OCCURRED_DATE

For mapping the job related attributes in MGMT_NOTIFY_CORRECTIVE_ACTION object, use the following map.

Table 4–44 Corrective Action Mapping

MGMT_NOTIFY_CORRECTIVE_ACTION	12c Notification Payload
JOB_NAME	gc$ notif_corrective_action_job.job_name
JOB_OWNER	gc$ notif_corrective_action_job.job_owner
JOB_TYPE	gc$ notif_corrective_action_job.job_type
JOB_STATUS	gc$ notif_corrective_action_job.status_code
STATE_CHANGE_GUID	gc$ notif_corrective_action_job. job_state_change_guid
JOB_GUID	gc$ notif_corrective_action_job. job _guid
EXECUTION_ID	gc$ notif_corrective_action_job. job_execution_guid
OCCURRED_DATE	gc$ notif_corrective_action_job.occurred_date
TARGETS	There can be at most one target. Use the values from gc$notif_target.target_name, gc$notif_target.target_type for the associated target.

4.2.1.3 Adding a Notification Method Based on an SNMP Trap

Enterprise Manager supports integration with third-party management tools through the SNMP. For example, you can use SNMP to notify a third-party application that a selected metric has exceeded its threshold.

The trap is an SNMP Version 1 trap and is described by the MIB definition shown at the end of this chapter. See "Management Information Base (MIB)".

For more comprehensive configuration information, see the documentation specific to your platform; SNMP configuration differs from platform to platform.

> **Note:** Notification methods based on SNMP traps must be configured by an administrator with Super Administrator privileges before any user can then choose to select one or more of these SNMP trap methods while creating/editing a incident rule.

Step 1: Define a new notification method based on an SNMP trap.

Log in to Enterprise Manager as a Super Administrator. From the **Setup** menu, select **Notifications** and then select **Notification Method** to access the Notification Methods page. From this page you can add a new method based on an SNMP trap.

You must provide the name of the host (machine) on which the SNMP master agent is running and other details as shown in the following example. In Example 4–10, the SNMP host will receive your SNMP traps.

Example 4–10 SNMP Trap Required Information

```
Name HP OpenView Console
Description Notification method to send trap to HP OpenView console
SNMP Trap Host Name machine1.oracle.com
SNMP Host Port 162
SNMP Community public
This SNMP host will receive your SNMP traps.
```

> **Note:** A Test Trap button exists for you to test your setup.

Metric severity information will be passed as a series of variables in the SNMP trap.

An example SNMP Trap is shown in Example 4–11. Each piece of information is sent as a variable embedded in the SNMP Trap.

Example 4–11 SNMP Trap

```
*************V1 TRAP***[3]*****************
Community : public
Enterprise :1.3.6.1.4.1.111.15.2
Generic :6
Specific :3
TimeStamp :48960
Agent adress :10.228.163.210
1.3.6.1.4.1.111.15.3.1.1.2.1: NOTIF_NORMAL
1.3.6.1.4.1.111.15.3.1.1.3.1: Memory Utilization is 98.343%, crossed warning (90)
or critical (95) threshold.
1.3.6.1.4.1.111.15.3.1.1.4.1:
https://machine6140830.oracle.com:15430/em/redirect?pageType=sdk-core-event-consol
e-detailEvent&
issueID=AADFAE01AD3A95E3E040E40AD2A32041
```

```
1.3.6.1.4.1.111.15.3.1.1.5.1: Critical
1.3.6.1.4.1.111.15.3.1.1.6.1: CRITICAL
1.3.6.1.4.1.111.15.3.1.1.7.1: 0
1.3.6.1.4.1.111.15.3.1.1.10.1: Aug 19, 2011 4:50:18 PM PDT
1.3.6.1.4.1.111.15.3.1.1.11.1: Capacity
1.3.6.1.4.1.111.15.3.1.1.12.1: Capacity
1.3.6.1.4.1.111.15.3.1.1.13.1: Metric Alert
1.3.6.1.4.1.111.15.3.1.1.14.1: Load:memUsedPct
1.3.6.1.4.1.111.15.3.1.1.15.1: 15
1.3.6.1.4.1.111.15.3.1.1.16.1:
1.3.6.1.4.1.111.15.3.1.1.17.1: No
1.3.6.1.4.1.111.15.3.1.1.18.1: New
1.3.6.1.4.1.111.15.3.1.1.19.1: None
1.3.6.1.4.1.111.15.3.1.1.20.1: 0
1.3.6.1.4.1.111.15.3.1.1.21.1: machine6140830.oracle.com
1.3.6.1.4.1.111.15.3.1.1.22.1:
https://machine6140830.oracle.com:15430/em/redirect?pageType=TARGET_
HOMEPAGE&targetName=machine6140
830.oracle.com&targetType=host
1.3.6.1.4.1.111.15.3.1.1.23.1: Host
1.3.6.1.4.1.111.15.3.1.1.24.1: machine6140830.oracle.com
1.3.6.1.4.1.111.15.3.1.1.25.1: SYSMAN
1.3.6.1.4.1.111.15.3.1.1.27.1: 4.8.0.0.0
1.3.6.1.4.1.111.15.3.1.1.28.1:
1.3.6.1.4.1.111.15.3.1.1.39.1: snmp ruleset
1.3.6.1.4.1.111.15.3.1.1.40.1: snmp ruleset,snmp rule
1.3.6.1.4.1.111.15.3.1.1.41.1: SYSMAN
1.3.6.1.4.1.111.15.3.1.1.42.1: AADFAE01AD3A95E3E040E40AD2A32041
1.3.6.1.4.1.111.15.3.1.1.61.1: METRIC_GUID=86821B5F0CE858D6E4A7F7390E88B73C
1.3.6.1.4.1.111.15.3.1.1.62.1: SEVERITY_GUID=AADFADD572DE08E2E040E40AD2A3202C
1.3.6.1.4.1.111.15.3.1.1.63.1: CYCLE_GUID=AADFAE01AD3695E3E040E40AD2A32041
1.3.6.1.4.1.111.15.3.1.1.64.1: COLL_NAME=LoadLinux
1.3.6.1.4.1.111.15.3.1.1.65.1: METRIC_GROUP=Load
1.3.6.1.4.1.111.15.3.1.1.66.1: METRIC_COLUMN=Memory Utilization (%)
1.3.6.1.4.1.111.15.3.1.1.67.1: METRIC_DESCRIPTION=
1.3.6.1.4.1.111.15.3.1.1.68.1: VALUE=98.343
1.3.6.1.4.1.111.15.3.1.1.69.1: KEY_VALUE=
 1.3.6.1.4.1.111.15.3.1.1.84.1: NUM_KEYS=0
```

Step 2: Assign the notification method to a rule.

You can edit an existing rule (or create a new incident rule), then add an action to the rule that subscribes to the advanced notification method.

4.3 Passing Corrective Action Status Change Information

Passing corrective action status change attributes (such as new status, job name, job type, or rule owner) to PL/SQL procedures or OS commands/scripts allows you to customize automated responses to status changes. For example, you many want to call an OS script to open a trouble ticket for an in-house support trouble ticket system if a critical corrective action fails to run. In this case, you will want to pass status (for example, Problems or Aborted) to the script to open a trouble ticket and escalate the problem.

4.3.1 Passing Corrective Action Execution Status to an OS Command or Script

The notification system passes information to an OS script or executable via system environment variables. Conventions used to access environmental variables vary depending on the operating system:

- UNIX: $ENV_VARIABLE
- MS Windows: %ENV_VARIABLE%

The notification system sets the following environment variables before calling the script. The notification system will set the environment variable $NOTIF_TYPE = NOTIF_CA for Corrective Action Execution. The script can then use any or all of these variables within the logic of the script.

Following table lists the environment variables for corrective action, they are populated when a corrective action is completed for an event.

Table 4–45 Corrective Action Environment Variables

Environment Variable	Description
CA_JOB_STATUS	Corrective action job execution status.
CA_JOB_NAME	Name of the Corrective Action.
CA_JOB_OWNER	Owner of Corrective Action.
CA_JOB_STEP_OUTPUT	The value will be the text output from the Corrective Action execution.
CA_JOB_TYPE	Corrective Action Job type

4.3.2 Passing Corrective Action Execution Status to a PLSQL Procedure

The notification system passes corrective action status change information to PL/SQL procedure - PROCEDURE p(event_msg IN gc$notif_event_msg). The instance gc$notif_corrective_action_job object is defined in event_msg.event_payload. corrective_action if event_msg. msg_info. notification_type is equal to GC$NOTIFICATIONNOTIF_CA. When a corrective action executes, the notification system calls the PL/SQL procedure associated with the incident rule and passes the populated object to the procedure. The procedure is then able to access the fields of the object that has been passed to it. See Table 4–34, " Corrective Action Job-Specific Attributes" for details.

The following status codes are possible values for the job_status field of the MGMT_ NOTIFY_CORRECTIVE_ACTION object.

Table 4–46 Corrective Action Status Codes

Name	Datatype	Value
SCHEDULED_STATUS	NUMBER(2)	1
EXECUTING_STATUS	NUMBER(2)	2
ABORTED_STATUS	NUMBER(2)	3
FAILED_STATUS	NUMBER(2)	4
COMPLETED_STATUS	NUMBER(2)	5
SUSPENDED_STATUS	NUMBER(2)	6
AGENTDOWN_STATUS	NUMBER(2)	7
STOPPED_STATUS	NUMBER(2)	8

Table 4–46 (Cont.) Corrective Action Status Codes

Name	Datatype	Value
SUSPENDED_LOCK_STATUS	NUMBER(2)	9
SUSPENDED_EVENT_STATUS	NUMBER(2)	10
SUSPENDED_BLACKOUT_STATUS	NUMBER(2)	11
STOP_PENDING_STATUS	NUMBER(2)	12
SUSPEND_PENDING_STATUS	NUMBER(2)	13
INACTIVE_STATUS	NUMBER(2)	14
QUEUED_STATUS	NUMBER(2)	15
FAILED_RETRIED_STATUS	NUMBER(2)	16
WAITING_STATUS	NUMBER(2)	17
SKIPPED_STATUS	NUMBER(2)	18
REASSIGNED_STATUS	NUMBER(2)	20

4.4 Passing Job Execution Status Information

Passing job status change attributes (such as new status, job name, job type, or rule owner) to PL/SQL procedures or OS commands/scripts allows you to customize automated responses to status changes. For example, you many want to call an OS script to open a trouble ticket for an in-house support trouble ticket system if a critical job fails to run. In this case you will want to pass status (for example, Problems or Aborted) to the script to open a trouble ticket and escalate the problem. The job execution status information is one of event type - job_status_change event, and its content is in OS command and PL/SQL payload as described in Section 4.2.1.1, "Adding a Notification Method based on an OS Command or Script" and Section 4.2.1.2, "Adding a Notification Method Based on a PL/SQL Procedure".

4.4.1 Passing Job Execution Status to a PL/SQL Procedure

The notification system passes job status change information to a PL/SQL procedure via the event_msg.event_payload object where event_type is equal to job_status_change. An instance of this object is created for every status change. When a job changes status, the notification system calls the PL/SQL p(event_msg IN `gc$notif_event_msg)` procedure associated with the incident rule and passes the populated object to the procedure. The procedure is then able to access the fields of the event_msg.event_payload object that has been passed to it.

Table 4–47 lists all corrective action status change attributes that can be passed:

Table 4–47 Job Status Attributes

Attribute	Datatype	Additional Information
event_msg.event_payload.source.source_name	VARCHAR2(128)	The job name.
event_msg.event_payload.source.source_owner	VARCHAR2(256)	The owner of the job.
event_msg.event_payload.source.source_sub_type	VARCHAR2(32)	The type of the job.

Table 4–47 (Cont.) Job Status Attributes

Attribute	Datatype	Additional Information
event_msg.event_ payload. event_ attrs(i).value where event_attrs(i).name=' execution_status'	NUMBER	The new status of the job.
event_msg.event_ payload. event_ attrs(i).value where event_ attrs(i).name='state_ change_guid'	RAW(16)	The GUID of the state change record.
event_msg.event_ payload.source.source_ guid	RAW(16)	The unique id of the job.
event_msg target.event_ payload. event_ attrs(i).value where event_attrs(i).name=' execution_id'	RAW(16)	The unique id of the execution.
event_msg.event_ payload.target	gc$notif_target	Target Information object..
event_msg.msg_ info.rule_owner	VARCHAR2(64)	The name of the notification rule that cause the notification to be sent.
event_msg.msg_ info.rule_name	VARCHAR2(132)	The owner of the notification rule that cause the notification to be sent.
event_msg.event_ payload. reported_date	DATE	The time and date when the status change happened.

When a job status change occurs for the job, the notification system creates an instance of the event_msg.event_payload. event_attrs(i).value where event_attrs(i).name=' execution_status' object and populates it with values from the status change. The following status codes have been defined as constants in the MGMT_JOBS package and can be used to determine the type of status in the job_status field of the event_ msg.event_payload. event_attrs(i).value where event_attrs(i).name=' execution_status' object.

Table 4–48 Job Status Codes

Name	Datatype	Value
SCHEDULED_STATUS	NUMBER(2)	1
EXECUTING_STATUS	NUMBER(2)	2
ABORTED_STATUS	NUMBER(2)	3
FAILED_STATUS	NUMBER(2)	4
COMPLETED_STATUS	NUMBER(2)	5
SUSPENDED_STATUS	NUMBER(2)	6
AGENTDOWN_STATUS	NUMBER(2)	7
STOPPED_STATUS	NUMBER(2)	8
SUSPENDED_LOCK_STATUS	NUMBER(2)	9

Table 4–48 (Cont.) Job Status Codes

Name	Datatype	Value
SUSPENDED_EVENT_STATUS	NUMBER(2)	10
SUSPENDED_BLACKOUT_STATUS	NUMBER(2)	11
STOP_PENDING_STATUS	NUMBER(2)	12
SUSPEND_PENDING_STATUS	NUMBER(2)	13
INACTIVE_STATUS	NUMBER(2)	14
QUEUED_STATUS	NUMBER(2)	15
FAILED_RETRIED_STATUS	NUMBER(2)	16
WAITING_STATUS	NUMBER(2)	17
SKIPPED_STATUS	NUMBER(2)	18
REASSIGNED_STATUS	NUMBER(2)	20

Example 4–12 PL/SQL Procedure Using a Status Code (Job)

```
CREATE  TABLE job_log (jobid RAW(16), status_code NUMBER(2), occured DATE);

CREATE OR REPLACE PROCEDURE LOG_JOB_STATUS_CHANGE(event_msg IN GC$NOTIF_EVENT_MSG)
IS
  l_attrs gc$notif_event_attr_array;
  exec_status_code NUMBER(2) := NULL;
  occured_date DATE := NULL;
  job_guid RAW(16) := NULL;

BEGIN
  IF event_msg.event_payload.event_type = 'job_status_change'
  THEN
    l_attrs := event_msg.event_payload.event_attrs;
    IF l_attrs IS NOT NULL
    THEN
      FOR i IN 1..l_attrs.COUNT
      LOOP
        IF l_attrs(i).name = 'exec_status_code'
        THEN
          exec_status_code := TO_NUMBER(l_attrs(i).value);
        END IF;
      END LOOP;
    END IF;

    occured_date := event_msg.event_payload.reported_date;
    job_guid := event_msg.event_payload.source.source_guid;
    -- Log all jobs' status
    BEGIN
      INSERT INTO job_log (jobid, status_code, occured)
      VALUES (job_guid, exec_status_code, occured_date);
    EXCEPTION
    WHEN OTHERS
    THEN
      -- If there are any problems then get the notification retried
      RAISE_APPLICATION_ERROR(-20000, 'Please retry');
    END;
    COMMIT;

  ELSE
```

```
    null; -- it is not a job_status_change event, ignore
  END IF;
END LOG_JOB_STATUS_CHANGE;
/
```

4.4.2 Passing Job Execution Status to an OS Command or Script

The notification system passes job execution status information to an OS script or executable via system environment variables. Conventions used to access environmental variables vary depending on the operating system:

- UNIX: $ENV_VARIABLE
- MS Windows: %ENV_VARIABLE%

The notification system sets the following environment variables before calling the script. The script can then use any or all of these variables within the logic of the script.

Table 4–49 Environment Variables

Environment Variable	Description
SOURCE_OBJ_NAME	The name of the job.
SOURCE_OBJ_OWNE	The owner of the job.
SOURCE_OBJ_SUB_TYPE	The type of job.
EXEC_STATUS_CODE	The job status.
EVENT_REPORTED_ TIME	Time when the severity occurred.
TARGET_NAME	The name of the target.
TARGET_TYPE	The type of the target.
RULE_NAME	Name of the notification rule that resulted in the severity.
RULE_OWNER	Name of the Enterprise Manager administrator who owns the notification rule.

4.5 Passing User-Defined Target Properties to Notification Methods

Enterprise Manager allows you to define target properties (accessed from the target home page) that can be used to store environmental or usage context information specific to that target. Target property values are passed to custom notification methods where they can be processed using conditional logic or simply passed as additional alert information to third-party devices, such as ticketing systems. By default, Enterprise Manager passes all defined target properties to notification methods.

Note: Target properties are not passed to notification methods when short e-mail format is used.

Figure 4–4 Host Target Properties

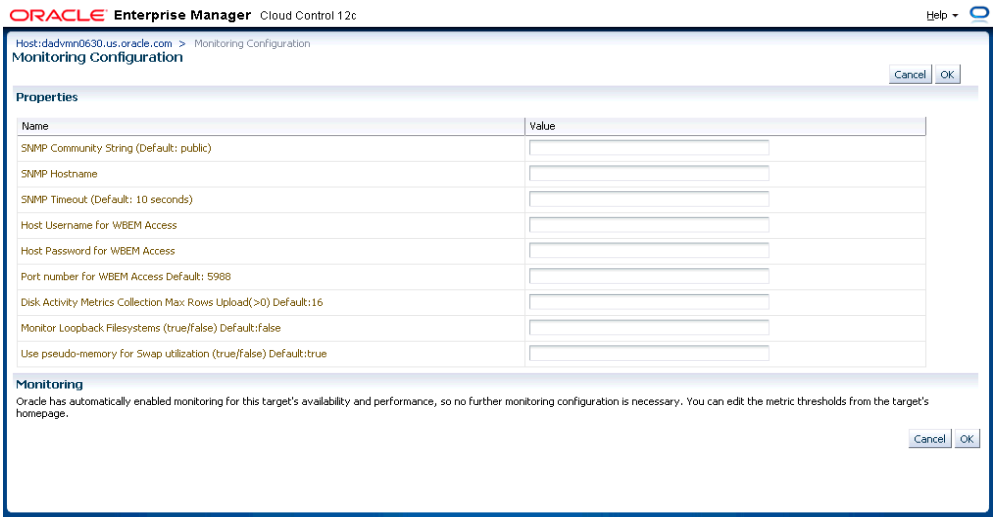

4.6 Management Information Base (MIB)

Enterprise Manager Cloud Control can send SNMP Traps to third-party, SNMP-enabled applications. Details of the trap contents can be obtained from the management information base (MIB) variables. The following sections discuss Enterprise Manager MIB variables in detail.

4.6.1 About MIBs

A MIB is a text file, written in ASN.1 notation, which describes the variables containing the information that SNMP can access. The variables described in a MIB, which are also called MIB objects, are the items that can be monitored using SNMP. There is one MIB for each element being monitored. Each monolithic or subagent consults its respective MIB in order to learn the variables it can retrieve and their characteristics. The encapsulation of this information in the MIB is what enables master agents to register new subagents dynamically — everything the master agent needs to know about the subagent is contained in its MIB. The management framework and management applications also consult these MIBs for the same purpose. MIBs can be either standard (also called public) or proprietary (also called private or vendor).

The actual values of the variables are not part of the MIB, but are retrieved through a platform-dependent process called "instrumentation". The concept of the MIB is very important because all SNMP communications refer to one or more MIB objects. What is transmitted to the framework is, essentially, MIB variables and their current values.

4.6.2 Reading the MIB Variable Descriptions

This section covers the format used to describe MIB variables. Note that the STATUS element of SNMP MIB definition, Version 2, is not included in these MIB variable descriptions. Since Oracle has implemented all MIB variables as CURRENT, this value does not vary.

4.6.2.1 Variable Name

Syntax
Maps to the SYNTAX element of SNMP MIB definition, Version 2.

Max-Access
Maps to the MAX-ACCESS element of SNMP MIB definition, Version 2.

Status
Maps to the STATUS element of SNMP MIB definition, Version 2.

Explanation
Describes the function, use and precise derivation of the variable. (For example, a variable might be derived from a particular configuration file parameter or performance table field.) When appropriate, incorporates the DESCRIPTION part of the MIB definition, Version 2.

Typical Range
Describes the typical, rather than theoretical, range of the variable. For example, while integer values for many MIB variables can theoretically range up to 4294967295, a typical range in an actual installation will vary to a lesser extent. On the other hand, some variable values for a large database can actually exceed this "theoretical" limit (a "wraparound"). Specifying that a variable value typically ranges from 0 to 1,000 or 1,000 to 3 billion will help the third-party developer to develop the most useful graphical display for the variable.

Significance
Describes the significance of the variable when monitoring a typical installation. Alternative ratings are Very Important, Important, Less Important, or Not Normally Used. Clearly, the DBA will want to monitor some variables more closely than others. However, which variables fall into this category can vary from installation to installation, depending on the application, the size of the database, and on the DBA's objectives. Nevertheless, assessing a variable's significance relative to the other variables in the MIB can help third-party developers focus their efforts on those variables of most interest to the most DBAs.

Related Variables
Lists other variables in this MIB, or other MIBs implemented by Oracle, that relate in some way to this variable. For example, the value of this variable might derive from that of another MIB variable. Or perhaps the value of this variable varies inversely to that of another variable. Knowing this information, third-party developers can develop useful graphic displays of related MIB variables.

Suggested Presentation
Suggests how this variable can be presented most usefully to the DBA using the management application: as a simple value, as a gauge, or as an alarm, for example.

4.6.2.2 MIB Definition

You can find the SNMP MIB file at the following location:

$ORACLE_HOME/network/doc/omstrap.v1

The file omstrap.v1 is the OMS MIB.

4.7 Troubleshooting Notifications

To function properly, the notification system relies on various components of Enterprise Manager and your IT infrastructure. For this reason, there can be many causes of notification failure. The following guidelines and suggestions can help you isolate potential problems with the notification system.

4.7.1 General Setup

The first step in diagnosing notification issues is to ensure that you have properly configured and defined your notification environment.

OS Command, PL/SQL and SNMP Trap Notifications

Make sure all OS Command, PLSQL and SNMP Trap Notification Methods are valid by clicking the Test button. This will send a test notification and show any problems the OMS has in contacting the method. Make sure that your method was called, for example, if the OS Command notification is supposed to write information to a log file, check that it has written information to its log file.

E-mail Notifications

- Make sure an e-mail gateway is set up under the Notification Methods page of Setup. The Sender's e-mail address should be valid. Clicking the Test button will send an e-mail to the Sender's e-mail address. Make sure this e-mail is received. Note that the Test button ignores any Notification Schedule.
- Make sure an e-mail address is set up. Clicking the Test button will send an e-mail to specified address and you should make sure this e-mail is received. Note that the Test button ignores any Notification Schedule.
- Make sure an e-mail schedule is defined. No e-mails will be sent unless a Notification Schedule has been defined.
- Make sure a incident rule is defined that matches the states you are interested and make sure e-mail and notification methods are assigned to the rule.

4.7.2 Notification System Errors

For any alerts involving problems with notifications, check the following for notification errors.

- Any serious errors in the Notification System are logged as system errors in the MGMT_SYSTEM_ERROR_LOG table. From the **Setup** menu, select **Management Services and Repository** to view these errors.
- Check for any delivery errors. You can view them from Incident Manager. From the **Enterprise** menu, select **Monitoring**, then select **Incident Manager**. The details will give the reason why the notification was not delivered. Delivery errors are stored in MGMT_NOTIFICATION_LOG with the DELIVERED column set to 'N'.
- Severities will not be displayed in the Enterprise Manager console if no metric values have been loaded for the metric associated with the severity.

4.7.3 Notification System Trace Messages

The Notification System can produce trace messages in sysman/log/emoms.trc file.

Tracing is configured by setting the *log4j.em.notification* property flag using the `emctl set property` command. You can set the trace level to INFO, WARN, DEBUG. For example,

```
./emctl set property -sysman_pwd your_sysman_password -name log4j.em.notification
-value DEBUG
```

Trace messages contain the string "em.notification". If you are working in a UNIX environment, you can search for messages in the emoms.trc and emoms_pbs.trc files using the `grep` command. For example,

```
grep em.notification emoms.trc emoms_pbs.trc
```

What to look for in the trace file.

The following entries in the emoms.trc file are relevant to notifications.

Normal Startup Messages

When the OMS starts, you should see these types of messages.

```
2011-08-17 13:50:29,458 [EventInitializer] INFO  em.notification init.167 - Short
format maximum length is 155
2011-08-17 13:50:29,460 [EventInitializer] INFO  em.notification init.185 - Short
format is set to both subject and body
2011-08-17 13:50:29,460 [EventInitializer] INFO  em.notification init.194 -
Content-Transfer-Encoding is 8-bit
2011-08-17 13:50:29,460 [EventInitializer] DEBUG em.notification
registerAdminMsgCallBack.272 - Registering notification system message call back
2011-08-17 13:50:29,461 [EventInitializer] DEBUG em.notification
registerAdminMsgCallBack.276 - Notification system message callback is registered
successfully
2011-08-17 13:50:29,713 [EventInitializer] DEBUG em.notification
upgradeEmailTemplates.2629 - Enter upgradeEmailTemplates
2011-08-17 13:50:29,735 [EventInitializer] INFO  em.notification
upgradeEmailTemplates.2687 - Email template upgrade is not required since no
customized templates exist.
2011-08-17 13:49:28,739 [EventCoordinator] INFO  events.EventCoordinator logp.251
- Creating event worker thread pool: min = 4 max = 15
2011-08-17 13:49:28,791 [[STANDBY] ExecuteThread: '2' for queue:
'weblogic.kernel.Default (self-tuning)'] INFO emdrep.pingHBRecorder
initReversePingThreadPool.937 - Creating thread pool for reverse ping : min = 10
max = 50
2011-08-17 13:49:28,797 [[STANDBY] ExecuteThread: '2' for queue:
'weblogic.kernel.Default (self-tuning)'] DEBUG emdrep.HostPingCoordinator logp.251
- Creating thread pool of worker thread  for host ping: min = 1 max = 10
2011-08-17 13:49:28,799 [[STANDBY] ExecuteThread: '2' for queue:
'weblogic.kernel.Default (self-tuning)'] DEBUG emdrep.HostPingCoordinator logp.251
- Creating thread pool for output of worker's  output for host ping: min = 2 max =
20
2011-08-17 13:49:30,327 [ConnectorCoordinator] INFO
connector.ConnectorPoolManager logp.251 - Creating Event thread pool: min = 3 max
= 10
2011-08-17 13:51:48,152 [NotificationMgrThread] INFO  notification.pbs logp.251 -
Creating thread pool: min = 6 max = 24
2011-08-17 13:51:48,152 [NotificationMgrThread] INFO  em.rca logp.251 - Creating
RCA thread pool: min = 3 max = 20
```

Notification Delivery Messages

```
2006-11-08 03:18:45,387 [NotificationMgrThread] INFO  em.notification run.682 -
Notification ready on EMAIL1

2006-11-08 03:18:46,006 [DeliveryThread-EMAIL1] INFO  em.notification run.114 -
Deliver to SYSMAN/admin@oracle.com

2006-11-08 03:18:47,006 [DeliveryThread-EMAIL1] INFO  em.notification run.227 -
Notification handled for SYSMAN/admin@oracle.com
```

Notification System Error Messages

```
2011-08-17 14:02:23,905 [NotificationMgrThread] DEBUG notification.pbs logp.251 -
Notification ready on EMAIL1
2011-08-17 14:02:23,911 [NotificationMgrThread] DEBUG notification.pbs logp.251 -
Notification ready on PLSQL4
2011-08-17 14:02:23,915 [NotificationMgrThread] DEBUG notification.pbs logp.251 -
Notification ready on OSCMD14
2011-08-17 14:02:19,057 [DeliveryThread-EMAIL1] INFO  notification.pbs logp.251 -
Deliver to To: my.admin@oracle.com; issue type: 1; notification type: 1
2011-08-17 14:02:19,120 [DeliveryThread-OSCMD14] INFO  notification.pbs logp.251 -
Deliver to SYSMAN, OSCMD, 8; issue type: 1; notification type: 1
2011-08-17 14:02:19,346 [DeliveryThread-PLSQL4] INFO  notification.pbs logp.251 -
Deliver to SYSMAN, LOG_JOB_STATUS_CHANGE, 9; issue type: 1; notification type: 1
2011-08-17 14:02:19,977 [DeliveryThread-PLSQL4] DEBUG notification.pbs logp.251 -
Notification handled for SYSMAN, LOG_JOB_STATUS_CHANGE, 9
2011-08-17 14:02:20,464 [DeliveryThread-EMAIL1] DEBUG notification.pbs logp.251 -
Notification handled for To: my.admin@oracle.com
2011-08-17 14:02:20,921 [DeliveryThread-OSCMD14] DEBUG notification.pbs logp.251 -
Notification handled for SYSMAN, OSCMD, 8
```

4.7.4 E-mail Errors

The SMTP gateway is not set up correctly:

```
Failed to send e-mail to my.admin@oracle.com: For e-mail notifications to be sent,
your Super Administrator must configure an Outgoing Mail (SMTP) Server within
Enterprise Manager. (SYSMAN, myrule)
```

Invalid host name:

```
Failed to connect to gateway: badhost.oracle.com: Sending failed;
nested exception is:
javax.mail.MessagingException: Unknown SMTP host: badhost.example.com;
```

Invalid e-mail address:

```
Failed to connect to gateway: rgmemeasmtp.oraclecorp.com: Sending failed;
nested exception is:
javax.mail.MessagingException: 550 5.7.1 <smpemailtest_ie@example.com>... Access
denied
```

Always use the Test button to make sure the e-mail gateway configuration is valid. Check that an e-mail is received at the sender's e-mail address

4.7.5 OS Command Errors

When attempting to execute an OS command or script, the following errors may occur. Use the Test button to make sure OS Command configuration is valid. If there are any errors, they will appear in the console.

Invalid path or no read permissions on file:

```
Could not find /bin/myscript (machineb10.oracle.com_Management_Service) (SYSMAN,
myrule )
```

No execute permission on executable:

```
Error calling /bin/myscript: java.io.IOException: /bin/myscript: cannot execute
(machineb10.oracle.com_Management_Service) (SYSMAN, myrule )
```

Timeout because OS Command ran too long:

```
Timeout occurred running /bin/myscript (machineb10.oracle.com_Management_Service)
(SYSMAN, myrule )
```

Any errors such as out of memory or too many processes running on OMS machine will be logged as appropriate.

Always use the Test button to make sure OS Command configuration is valid.

4.7.6 SNMP Trap Errors

Use the Test button to make sure SNMP Trap configuration is valid.

Other possible SNMP trap problems include: invalid host name, port, or community for a machine running an SNMP Console.

4.7.7 PL/SQL Errors

When attempting to execute an PL/SQL procedure, the following errors may occur. Use the Test button to make sure the procedure is valid. If there are any errors, they will appear in the console.

Procedure name is invalid or is not fully qualified. Example: SCOTT.PKG.PROC

```
Error calling PL/SQL procedure plsql_proc: ORA-06576: not a valid function or
procedure name (SYSMAN, myrule)
```

Procedure is not the correct signature. Example: PROCEDURE event_proc(s IN GC$NOTIF_EVENT_MSG)

```
Error calling PL/SQL procedure plsql_proc: ORA-06553: PLS-306: wrong number or
types of arguments in call to 'PLSQL_PROC' (SYSMAN, myrule)
```

Procedure has bug and is raising an exception.

```
Error calling PL/SQL procedure plsql_proc: ORA-06531: Reference to uninitialized
collection (SYSMAN, myrule)
```

Care should be taken to avoid leaking cursors in your PL/SQL. Any exception due to this condition will result in delivery failure with the message being displayed in the Details section of the alert in the Cloud Control console.

Always use the Test button to make sure the PL/SQL configuration is valid.

5

Managing with Groups

This chapter introduces the concept of group management and contains the following sections:

- Introduction to Groups
- Managing Groups
- About Out-of-Box Reports
- About Redundancy Systems
- About Privilege Propagating Groups

5.1 Introduction to Groups

Today's IT operations can be responsible for managing a great number of components, such as databases, application servers, hosts, or other components, which can be time consuming and impossible to manage individually. The Enterprise Manager Cloud Control group management system lets you combine components (called targets in Enterprise Manager) into logical sets, called groups. This enables you to organize, manage, and effectively monitor the potentially large number of targets in your enterprise.

Enterprise Manager Groups can include:

- Targets of the same type, such as:
  - All hosts in your data center
  - All of your production databases
- Targets of different types, such as:
  - The database, application server, listener, and host that are used in your application environment
  - Targets operating within a particular data center region

> **Note:** An Enterprise Manager "System," used specifically to group the components on which a service runs, is a special kind of Enterprise Manager group. Many of the functions and capabilities for groups and systems are similar.

Typically you can gather together targets that you want to manage as a group. If you use the target properties (for example, *Line of Business* or *Deployment Type*) to put operational information about your targets in Enterprise Manager, you can use these

properties when creating groups to locate targets. For example, you could search for all databases of *Deployment Type = Production* and belonging to *Line of Business 'HCM'*. You can also create a group hierarchy and use nested groups.

5.2 Managing Groups

By combining targets in a group, Enterprise Manager offers a wealth of management features that enable you to efficiently manage these targets as one group. Using the Group pages, you can:

- View a summary status of the targets within the group.
- Monitor outstanding alerts and incidents for the group collectively, rather than individually.
- Monitor the overall performance of the group.
- Perform administrative tasks, such as scheduling jobs for the entire group, or blacking out the group for maintenance periods.

You can also customize the console to provide direct access to group management pages.

5.2.1 Using the Groups Page

When you choose Groups from the Targets menu on the Enterprise Manager menu bar, the Groups page appears. From the page you can view the currently available groups and perform the following tasks:

- View a list of all the defined groups.
- Search for existing groups and save search criteria for future searches.
- View a roll-up of the outstanding alerts and incidents for members in a group.
- Create administration groups, associate template collections, and disassociate template collections
- Add groups or privilege propagating groups, remove groups, and change the configuration of currently defined groups.
- Drill down from a specific group to collectively monitor and manage its member targets.

Redundancy systems and special high availability groups are not accessed from this Groups page. You can access them from the All Targets page.

5.2.2 About the Group Home Page

The Group Home page enables you to quickly view key information about members of a group, eliminating the need to navigate to individual member targets to check on availability and performance. You can view the entire group on a single screen and drill down to obtain further details. The rolled up numbers include alerts and incidents for all members including those in nested groups. The Group Home page provides the following sections:

- A General section that shows the Owner, Group Type, and Privilege Propagation status.
- A Status section that shows how many member targets are in up, down, and unknown states. For nested groups, this segment shows how many targets are in up, down, and unknown states across all its sub-groups. The status roll up count is

based on the unique member targets across all sub-groups. Consequently, even if a target appears more than once in sub-groups, it is counted only once in status roll ups. Click member names to go to the member Status page.

- An Overview of Incidents and Problems section that displays the number of outstanding critical, warning, and error alerts associated with the current group. For nested groups, this segment shows how many targets are in an alert state across all its sub-groups.

  The rolled up information is shown for all the member targets regardless of their status. The status roll up count is based on the unique member targets across all sub-groups. Consequently, even if a target appears more than once in sub-groups, its alerts are counted only once in alert roll ups.

  Click the number in the Problems column to go to the Incident Manager page to search, view, and manage exceptions and issues in your environment. By using Incident Manager, you can track outstanding incidents and problems.

- A Compliance Summary section that shows how many of your group members do not comply with Enterprise Manager policy rules. Non-compliant members are indicated with the number of critical, warning, and informational incidents along the Average Compliance Score (as a percentage) for each compliance rule. You can click the Members tab to see the Member Targets and their types along with any violations and an average score for each member target.

  The numbers include the group-level incidents (if any group-level policy is defined), as well as group members violations. The rolled-up information for all policy categories, including security, is shown for all the member targets regardless of their status.

- A Job Activity section that displays the status for jobs that have started within the previous 7 days. The embedded table shows you the number of executions submitted to the group or any group members listed by status type, such as Problem Executions, Suspended Executions, and so on.

- A Blackouts section that allows you to create blackout periods and view the status of existing blackouts. The table displays the Scheduled and Active blackouts for each group or group member. Click Create to define primary blackout identification information and assign targets to be blacked out.

- A Patch Recommendations section that shows the total number of Oracle critical patch advisories (including one or more critical patches) that are applicable to your enterprise, and the number of Oracle homes in your enterprise to which those patches should be applied. You can view the information by Classification or by Target Type.

  Click the Current number link to go to the Group Critical Patch Advisories page. If your Oracle MetaLink Credentials are not configured, click Not Configured to go to the Patching Setup page. After you configure this page, Enterprise Manager collects information about Oracle critical patch advisories that are relevant to your enterprise.

- An Inventory and Usage section where you can view inventory summaries for deployments such as hosts, database installations, and fusion middleware installations on an enterprise basis or for specific targets. You can select an option such as Platform or Version to roll up inventory. Optionally, you can click See Details to navigate to the Inventory and Usage Details page where you can perform more detailed tasks such as viewing trends in inventory counts charted across a time line, revising selections to refresh chart and details based on new selections, or export deployment and details tables to CSV files.

- A Configuration Changes for Last 7 Days section that displays the number for configuration changes and Relationship changes incurred over the previous 7 days. Configuration history is a log of changes to a target, such as a group or beacon, recorded over a period of time. The recorded history includes changes both to configurations and to relationships. Relationships are the associations that exist among managed entities. You can click the number of changes in either column to view more detailed information about the change.

5.2.3 About the Group Charts Page

The Group Charts page enables you to monitor the collective performance of the group. Out-of-box performance charts are provided based on the type of members in the group. For example, when databases are part of the group, a Wait Time (%) chart is provided that shows the top databases with the highest wait time percentage values. You can view this performance information over the last 24 hours, last 7 days, or last 31 days. You can also add your own custom charts to the page.

You can access the Charts page by choosing Charts from the Monitoring sub-menu of the Group menu.

5.2.4 About the Group Members Page

The Group Members page summarizes information about the member targets in the group. It includes information on their current availability status, roll-up of open alerts and incidents, and key performance metrics based on the type of targets in the group.

You can visually assess availability and relative performance across all member targets. You can sort on any of the columns to rank members by a certain criterion (for example, database targets in order of decreasing wait time percentage). Default key performance metrics are displayed based on the targets you select, but you can customize these to include additional metrics that are important for managing your group.

5.2.5 Viewing Group Status History

You can use the Group Status History page to view the historical availability of a member during a specified time period, view the current status of all group members, or access the home pages for members.

Bar graphs provide a historical presentation of the availability of group members during a time period you select from the View Data drop-down list. The color-coded graphs can show statuses of Up, Down, Under Blackout, Agent Down, Metric Collection Error, and Status Pending. You can select time periods of 24 hours, 7 days, or 31 days.

To view the current status or a member, you can click a Status icon to go to the Availability page, which shows the member's current and past availability status within the last 24 hours, 7 days, or 31 days. Click a member Name to go to the member's Home page. You can use this page as a starting point when evaluating the performance of the selected member.

You can access the Group Status History page by choosing Status History from the Monitoring section of the Group menu.

5.2.6 About the System Dashboard

The System Dashboard enables you to pro-actively monitor the status and alerts in the group as they occur. The color-coded interface is designed to highlight problem areas

using the universal colors of alarm—targets that are down are highlighted in red, metrics in critical severity are shown as red dots, metrics in warning severity are shown as yellow dots, and metrics operating within normal boundary conditions are shown as green dots.

Using these colors, you can easily spot the problem areas for any target and drill down for details as needed. An alert table is also included to provide a summary for all open alerts in the group. The alerts in the table are presented in reverse chronological order to show the most recent alerts first, but you can also click any column in the table to change the sort order.

The Dashboard allows you to drill down for more detailed information. You can click the following items in the Dashboard for more information:

- A target name to access the target home page
- A group or system name to access the System Dashboard
- Status icon corresponding to specific metric columns to access the metric detail page
- Alerts icon for a group to access the Alerts page for that aggregate object
- Status icon for a metric with key values to access the metric page with a list of all key values
- Status icon for a metric with a specific key value to access the metric detail page with the specified key
- Dashboard header to access the group home page
- Status icon for down, critical or warning alerts to access the Alerts page
- Alerts messages to access the metric detail page containing the alert history for the target

Click **Customize** to access the Edit Group pages. By default, Enterprise Manager takes you to the Edit Group Dashboard page where you can change the target display and data refresh frequency. However, you can also modify any other group properties that affect the content of the System Dashboard. Columns that appear in the Dashboard target area mirror the columns that appear in the Edit Group Columns page. To display additional columns, click Modify on the Edit Group Columns page and add the desired metric columns.

In the "Group by Target Type" mode, the Dashboard displays information of the targets based on the specific target types present in the group or system. The statuses and alerts displayed are rolled up for the targets in that specific target type.

Columns that appear in the Dashboard target area mirror the columns that appear in the Edit Group Columns page. To display additional columns, click Modify on the Edit Group Columns page and add the desired metric columns.

If you minimize the dashboard window, pertinent alert information associated with the group or system is still displayed in the Microsoft Windows toolbar. For example, (#1 X3 !5) denotes there is 1 Target Down Alert, 3 Critical Alerts and 5 Warning Alerts associated with this group or system.

5.3 About Out-of-Box Reports

Enterprise Manager provides several out-of-box reports for groups as part of the reporting framework, called Information Publisher. These reports display important administrative information, such as hardware and operating system summaries across

all hosts within a group, and monitoring information, such as outstanding alerts and incidents for a group.

You can access these reports from the **Information Publisher Reports** menu item on the Groups menu.

See Also: Chapter 11, "Using Information Publisher"

5.4 About Redundancy Systems

A redundancy system is a group that contains members of the same type that function collectively as a unit. A type of redundancy system functions like a single logical target that supports a status (availability) metric. A redundancy system is considered up (available) if at least one of the member targets is up.

You can create and administer a redundancy system from the All Targets page. Redundancy systems support all group management features previously discussed.

When you define the Redundancy System, you must choose the member type for the members in the Redundancy System.

You can define the options for how availability of the redundancy system is calculated by selecting either Number or Percentage:

- Number - When you choose Number, you can specify either the number of member targets that should be up in order for the group to be considered up, or the number of member targets that should be down in order for the system to be considered down.
- Percentage - When you choose Percentage, you can specify either the minimum percentage of member targets that must be up in order for the system to be considered up, or the minimum percentage of member targets that are down in order to consider the system to be down. If you choose Percentage, the required number of member targets will be rounded off to the next integer. For example if you define the Percentage as 50% and the total number of member targets is 5, then the value used for calculating the availability will be 3.

Do not use redundancy systems if the group you want to model is an Oracle Real Application Clusters database, host cluster, HTTP server high availability group, or OC4J high availability group. Instead, you can use the following specialized target types for this purpose:

- Cluster
- Cluster Database
- HTTP HA Group
- OC4J HA Group

5.5 About Privilege Propagating Groups

Privilege propagating groups enable administrators to grant privileges to other administrators in a manner in which new administrators get the same privileges as its member targets. For example, granting *operator* privilege on a group to an Administrator grants him the *operator* privilege on its member targets and also to any members that will be added in the future. Privilege propagating groups can contain individual targets or other privilege propagating groups.

Privileges on the group can be granted to an Enterprise Manager user or a role. Use a role if the privileges you want to grant are to be granted to a group of Enterprise Manager users.

For example, suppose you create a large privilege propagating group and grant a privilege to a role which is then granted to administrators. If new targets are later added to the privilege propagating group, then the administrators receive the privileges on the target automatically. Additionally, when a new administrator is hired, you only need to grant the role to the administrator for the administrator to receive all the privileges on the targets automatically.

5.5.1 Creating Privilege Propagating Groups

The privilege propagating group creation function is a privileged activity. The privilege propagating group feature contains two privileges:

- Create Privilege Propagating Group

  This privileged activity allows the administrators to create the privilege propagating groups. Administrators with this privilege can create propagating groups and delegate the group administration activity to other users.

- Group Administration

  This privilege can be granted to administrators on specific group targets and is used to delegate the group administration activities to other administrators. It is granted to both conventional and privilege propagating groups.

5.5.2 Using the Group Administration Privilege

The Group Administration Privilege is available for both Privilege Propagating Groups and conventional groups. If you are granted this privilege, you can grant access to the group to other Enterprise Manager users without having to be the SuperAdministrator to grant the privilege.

5.5.3 Adding Members to Privilege Propagating Groups

The target privileges granted on a propagating group are propagated to member targets. The administrator grants target objects scoped to another administrator, and the grantee maintains the same privileges on member targets. The propagating groups maintain the following features:

- The administrator with a Create Privilege Propagating Group privilege will be able to create a propagating group
- To add a target as a member of a propagating group, the administrator must have *Full* target privileges on the target

You can add any non-aggregated target as the member of a privilege propagating group. For aggregated targets in Cloud Control version 12*c*, cluster and RAC databases and other propagating groups can be added as members (cluster and RAC databases must be added via the *EM CLI* verb). There is no support for this through the Enterprise Manager interface in version 10.2.0.5. Cloud Control version 12*c*, however, supports more aggregated target types, such as redundancy systems, systems and services. These, along with cluster and RAC databases, can be added in version 12*c* via the Cloud Control console.

If you are not the group creator, you must have at least the *Full* target privilege on the group to add a target to the group.

5.5.4 Converting Conventional Groups to Privilege Propagating Groups

In Enterprise Manager release 12*c* you can convert conventional groups to privilege propagating groups (and vice-versa) through the use of the specified EM CLI verb. Two new parameters have been added in the *modify_group* EM CLI verb:

- *privilege_propagation*

  This parameter is used to modify the privilege propagation behavior of the group. The possible value of this parameter is either true or false.

- *drop_existing_grants*

  This parameter indicates whether existing privilege grants on that group are to be revoked at the time of converting a group from privilege propagation to normal (or vice versa). The possible values of this parameter are yes or no. The default value of this parameter is yes.

These same enhancements have been implemented on the following EM CLI verbs: *modify_system*, *modify_redundancy_group*, and *modify_aggregrate_service*.

The EM CLI verb is listed below:

```
emcli modify_group
   -name="name"
   [-type=<group>]
   [-add_targets="name1:type1;name2:type2;..."]...
   [-delete_targets="name1:type1;name2:type2;..."]...
   [-privilege_propagation = true/false]
   [-drop_existing_grants = Yes/No]
```

For more information about this verb and other EM CLI verbs, see the *EM CLI Reference Manual*.

6

Administration Groups

Administration Groups greatly simplify the process of setting up targets for management in Enterprise Manager by automating the application of management settings such as monitoring settings or compliance standards. Typically, these settings are manually applied to individual target, or perhaps semi-automatically using custom scripts. However, by defining Administration Groups, Enterprise Manager uses specific target properties to direct the target to the appropriate Administration Group and then automatically apply the requisite monitoring and management settings. This level of automation simplifies the target setup process and also enables a datacenter to easily scale as new targets are added to Enterprise Manager for management.

This chapter covers the following topics:

- What is an Administration Group?
- Planning
- Implementing Administration Groups and Template Collections
- Removing Administration Groups

6.1 What is an Administration Group?

Administration groups are a special type of group used to fully automate application of monitoring and other management settings targets upon joining the group. When a target is added to the group, Enterprise Manager applies these settings using a Template Collection consisting of Monitoring Templates, compliance standards, and cloud policies. This completely eliminates the need for administrator intervention. The following illustration demonstrates the typical Administration Group workflow:

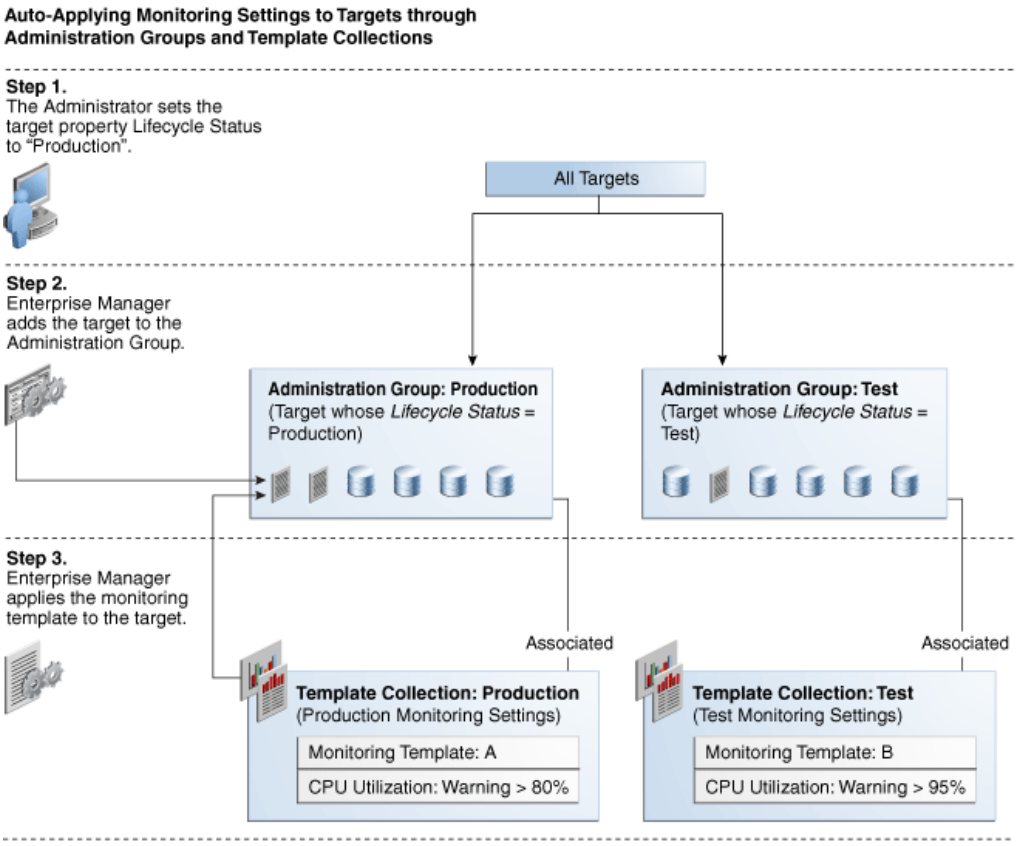

The first step involves setting a target's Lifecycle Status property when a target is first added to Enterprise Manager for monitoring. At that time, you determine where in the prioritization hierarchy that target belongs; the highest level being "mission critical" and the lowest being "development."

Target Lifecycle Status prioritization consists of the following levels:

- Mission Critical (highest priority)
- Production
- Stage
- Test
- Development (lowest priority)

As shown in step two of the illustration, once Lifecycle Status is set, Enterprise Manger uses it to determine which Administration Group the target belongs.

In order to prevent different monitoring settings to be applied to the same target, Administration Groups were designed to be mutually exclusive with other Administration Groups in terms of group membership. Administration groups can also be used for hierarchically classifying targets in an organization. For example, in the previous illustration, you have an Administration Group hierarchy consisting of two subgroups: *Production* targets and *Test* targets, with each subgroup having its own Template Collections. In this example, the Production group inherits monitoring settings from Monitoring Template A while targets in the Test subgroup inherit monitoring settings from Monitoring Template B.

6.1.1 Developing an Administration Group

In order to create an Administration Group, you must have both *Full Any Target* and *Create Privilege Propagating Group* target privileges.

Developing an Administration Group is performed in two phases:

- **Planning**
  - Plan your Administration Group hierarchy by creating a group hierarchy based on how you manage your targets.
  - Plan the management settings associated with the Administration Groups in the hierarchy.
    * Management settings: Monitoring settings, Compliance standard settings, Cloud policy settings
    * For Monitoring settings, you can have additional metric settings or override metric settings lower in your hierarchy
    * For Compliance standards or Cloud policies, additional rules/policies lower in the hierarchy are additive
- **Implementation**
  - Enter the group hierarchy definition and management settings in Enterprise Manager.
    * Create the Administration Group hierarchy.
    * Create the Monitoring Templates, compliance standards, cloud policies and add these to Template Collections.
    * Associate Template Collections with Administration Groups.
    * Add targets to the Administration Group by assigning the appropriate values to the target properties such that Enterprise Manager automatically adds them to the appropriate Administration Group.

6.2 Planning

As with any management decision, the key to effective implementation is planning and preparation. The same holds true for Administration Groups.

Step 1: Plan Your Group Hierarchy

You can only have one Administration Group hierarchy in your Enterprise Manager deployment, thus ensuring that Administration Group member targets can only directly belong to one Administration Group. This prevents monitoring conflicts from occurring as a result of having a target join multiple Administration Groups with different associated monitoring settings.

To define the hierarchy, you start with the highest (root) level consisting of all targets that have been added to Enterprise Manager. On the next level of the hierarchy, you organize the targets into groups such that all targets that are monitored and managed in the same way are part of the same group.

The attributes used to define Administration Group membership criteria are based on *target properties*, which are attributes of every target that specify operational information and within the organization. For example, *location*, *line of business* to which it belongs, and *lifecycle status*. Target properties that can be used in the creation of Administration Groups are:

- Lifecycle Status

> **Note:** Lifecycle Status target property is of particular importance when creating Administration Groups as it denotes a target's operational status. Lifecycle Status can be any of the following: Mission Critical, Production, Staging, Test, or Development.

- Location
- Line of Business
- Department
- Cost Center
- Contact
- Target Version
- Customer Support Identifier

You cannot manually add targets to an Administration Group. Instead, you set the target properties of the target (prospective group member) to match the membership criteria defined for the Administration Group. Once the target properties are set, Enterprise Manager automatically adds the target to the appropriate Administration Group.

In the following illustration, two Administration Groups are created, *Production* and *Test*, because monitoring settings for production targets will differ from the monitoring settings for test targets.

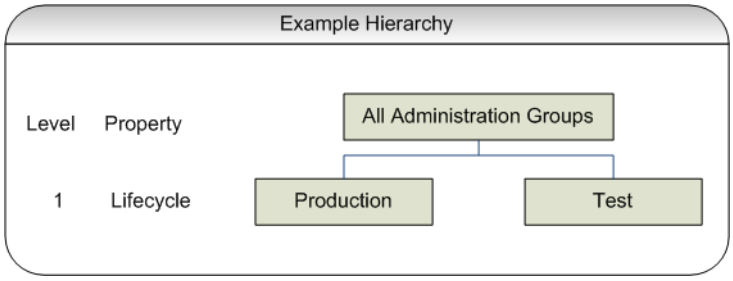

In this example, the group membership criteria are based on the *Lifecycle Status* target property. Targets whose *Lifecycle Status* is 'Production' join the Production group and targets whose *Lifecycle Status* is 'Test' join the Test group. For this reason, *Lifecycle Status* is the target property that determines the first level in the Administration Group hierarchy, and each value of the Lifecycle Status property determines the membership criteria of each Administration Group in the first level.

Additional levels in the Administration Group hierarchy can be added based on other target properties. Typically, additional levels are added if there are additional monitoring (or management) settings that need to be applied and these could be different for different subsets of targets in the Administration Group. For example, in the *Production* group, there could be additional monitoring settings for targets in *Finance* line of business that are different from targets in *Sales* line of business. In this case, an additional level based on *Line of Business* target property level would be added.

The end result of this hierarchy planning exercise is summarized in the following table.

Tree Root Level	Level 1 Target Property Lifecycle Status	Level 2 Target Property Line of Business
All Targets	Production or Mission Critical	Finance
		Sales
	Staging, Test, or Development	Finance
		Sales

Each cell of the table represents a group. The values in each cell represent the values of the target property that define membership criteria for the group.

It is possible to have the group membership criteria be based on more than one target property value. In that case, any target whose target property matches any of the values will be added to the group. For example, in the case of the Production group, if the *Lifecycle Status* of a target is either *Production* or *Mission Critical*, then it will be added to the Production group.

It is also important to remember that group membership criteria is cumulative. For example, for the *Finance* group under *Production or Mission Critical* group, a target must have its *Lifecycle Status* set to *Production or Mission Critical* **AND** its *Line of Business* set to *Finance* before it can join the group.

For this planning example, the resulting Administration Group hierarchy would appear as shown in the following graphic.

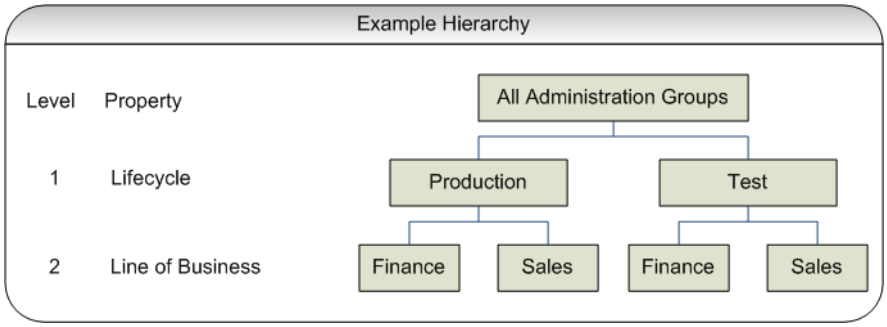

It is important to note that a target can become part of hierarchy if and only if its property values match criteria at both the levels. A target possessing matching values for *lifecycle status* cannot become member of the Administration Group at the first level.

Step 2: Assign Target Properties

After establishing the desired Administration Group hierarchy, you must make sure properties are set correctly for each target to ensure they join the correct Administration Group. Using target properties, Enterprise Manager automatically places targets into the appropriate Administration Group without user intervention. You can set target properties at the same time you add targets to Enterprise Manager (either using the Enterprise Manager console or using the EM CLI verb `add_target`). For targets that have already been added to Enterprise Manager, you can

also set the target properties via the console or using EM CLI. The benefit of setting the target properties prior to the creation of the Administration Group hierarchy is that once the group hierarchy is created, the targets whose properties are set will automatically join their appropriate Administration Groups. However, target properties can be set after the Administration Group hierarchy is created.

For small numbers of targets, you can change target properties directly from the Enterprise Manager console.

1. From an Enterprise Manager target's option menu, select **Target Setup**, then select **Properties**.

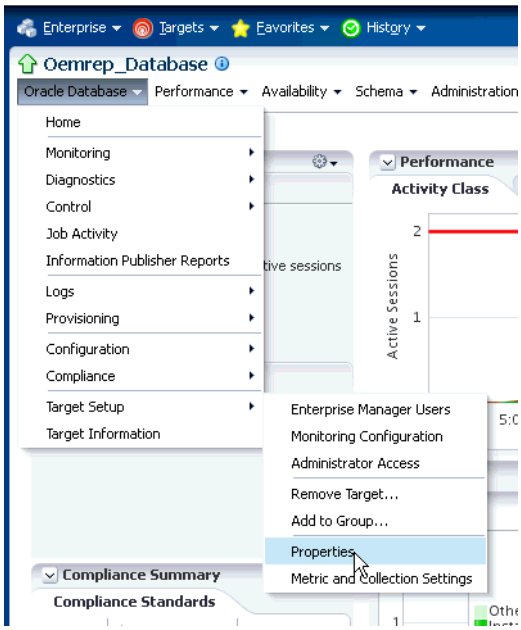

2. On the **Target Properties** page, click **Edit** to change the property values.

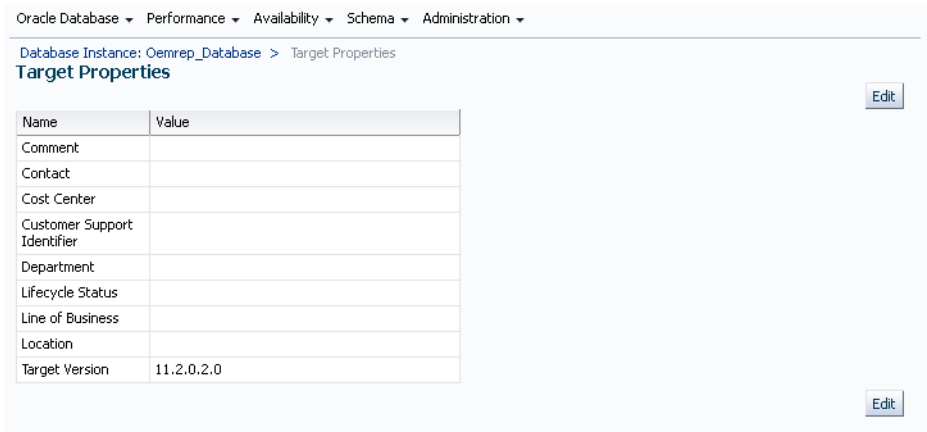

3. Once you have set the target properties, click **OK**.

For large numbers of targets, it is best to use the Enterprise Manager Command Line Interface (EM CLI) `set_target_property_value` verb to perform a mass update. For more information about this EM CLI verb, see the Enterprise Manager Command Line Interface guide.

At any time, you can use the **All Targets** page to view properties across all targets. To view target properties:

1. From the **Targets** menu, select **All Targets** to display the All Targets page.
2. From the **View** menu, select **Columns**, then select **Show All**.
3. Alternatively, if you are interested in specific target properties, choose **Columns** and then select **Show More Columns** to display column selector, as shown in following graphic.

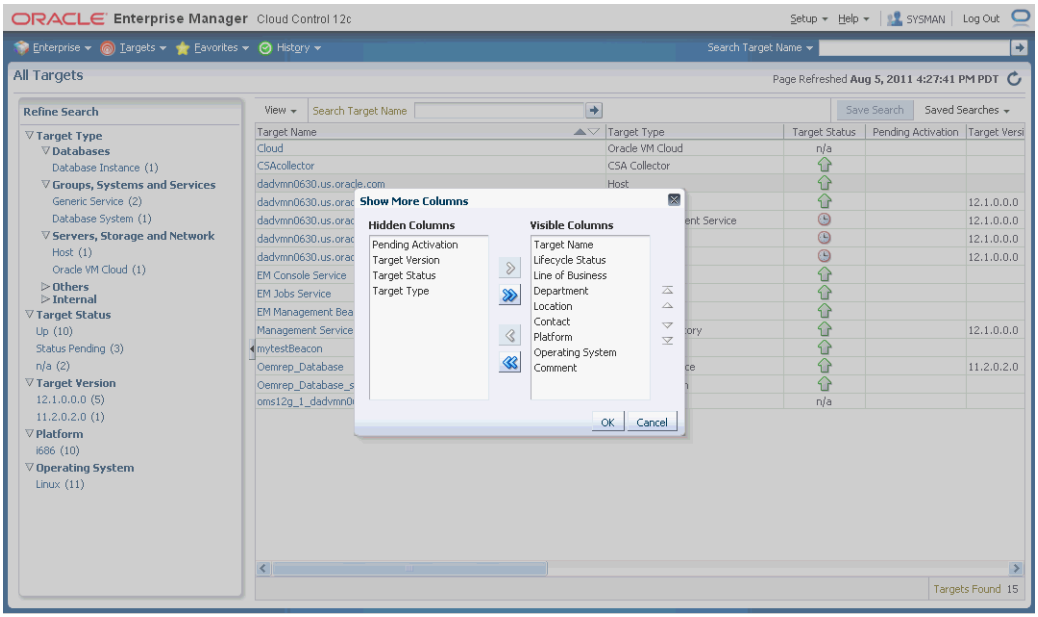

Step 3: Prepare for Creating Template Collections

Template collections contain the monitoring settings and other management settings that are meant to be applied to targets as they join the Administration Group. Monitoring settings for targets are defined in Monitoring Templates. Monitoring templates are defined on a per target type basis, so you will need to create Monitoring Templates for each of the different target types in your Administration Group. You will most likely create multiple Monitoring Templates to define the appropriate monitoring settings for an Administration Group. For example, you might create a database Monitoring Template containing the metric settings for your production databases and a separate Monitoring Template containing the settings for your non-production databases. Other management settings that can be added to a Template Collection include Compliance Standards and Cloud Policies. Ensure all of these entities that you want to add to your Template Collection are correctly defined in Enterprise Manager before adding them to Template Collections.

If you have an Administration Group hierarchy defined with more than two levels, such as the hierarchy shown in the following figure, it is important to understand how management settings are applied to the targets in the Administration Group.

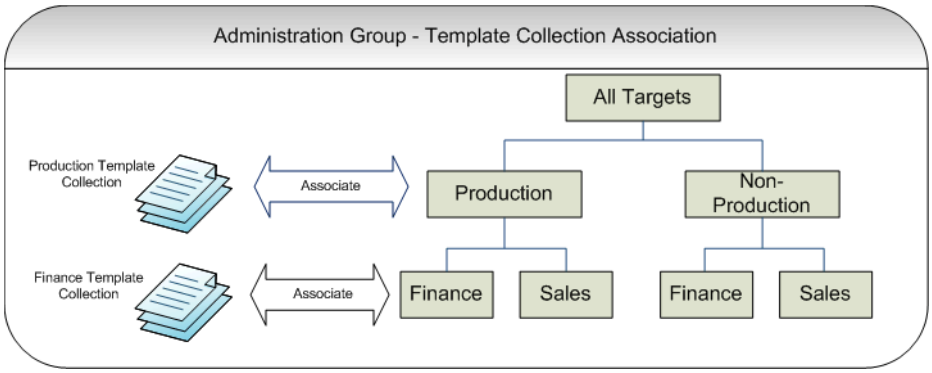

Each group in the Administration Group hierarchy can be associated with a Template Collection (containing Monitoring Templates, compliance standards, and cloud policies). If you associate a Template Collection containing monitoring settings with the *Production* group, then the monitoring settings will apply to the *Finance* and *Sales* subgroup under *Production*. If the *Finance* group under *Production* has additional monitoring settings, then you can create a Monitoring Template with only those additional monitoring settings. (Later, this Monitoring Template should be added to another Template Collection and associated with the *Finance* group). The monitoring settings from the *Finance Template Collection* will be logically combined with the monitoring settings from the *Production Template Collection*. In case there are duplicate metric settings in both Template Collections, then the metric settings from the *Finance Template Collection* takes precedence and will be applied to the targets in the *Finance* group. This precedence rule only applies to the case of metric settings. In the case of compliance standard rules and cloud policies, even if there are duplicate compliance standard rules and cloud policies in both Template Collections, they will be all applied to the targets in the *Finance* group.

Once you have completed all the planning and preparation steps, you are ready to begin creating an Administration Group.

6.3 Implementing Administration Groups and Template Collections

With the preparatory work complete, you are ready to begin the four step process of creating an Administration Group hierarchy and Template Collections. The Administration Group user interface is organized to guide you through the creation process, with each tab containing the requisite operations to perform each step.

This process involves:

1. Creating the Administration Group hierarchy.
2. Creating Template Collections.
3. Associating Template Collections to Administration Group.
4. Synchronizing the targets with the selected items.

The following graphic shows a completed Administration Group hierarchy with associated Template Collections. It illustrates how Enterprise Manager uses this to automate the application of target monitoring settings.

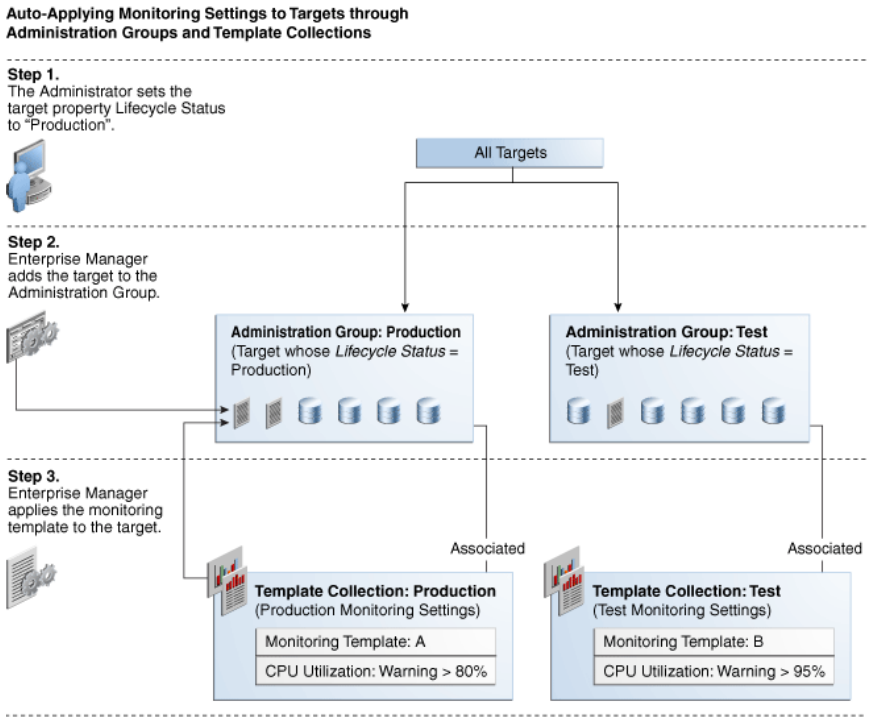

6.3.1 Creating an Administration Group

The following four primary tasks summarize the Administration Group creation process. These tasks are conveniently arranged in sequence via tabbed pages.

> **Important:** In order to create an Administration Group, you must have both **Full Any Target** and **Create Privilege Propagating Group** target privileges.

Task 1: Access the Administration Group and Template Collections page.

Task 2: Define the hierarchy.

From the **Hierarchy** tab, you define the Administration Group hierarchy that matches the way you manage your targets. See Section 6.3.1.2, "Defining a Hierarchy".

Task 3: Define the Template Collections.

From the **Template Collections** tab, you define the monitoring and management settings you want applied to targets. See Section 6.3.1.3, "Defining Template Collections".

Task 4: Associate the Template Collections with the Administration Groups

From the **Associations** tab, you tie the monitoring and management settings to the appropriate Administration Group. See Section 6.3.1.4, "Associating Template Collections with Administration Groups".

6.3.1.1 Accessing the Administration Group Home Page

All Administration Group operations are performed from the Administration Groups home page.

From the **Setup** menu, select **Add Target** and then select **Administration Groups**. The Administration Groups homepage displays.

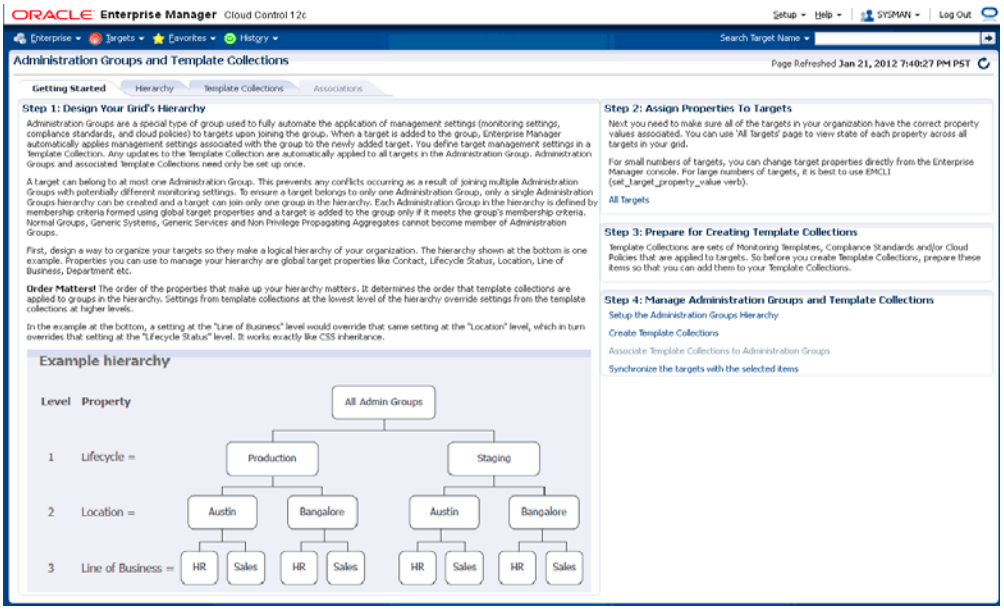

Read the relevant information on the **Getting Started** page. The information contained in this page summarizes the steps outlined in this chapter. For your convenience, links are provided that take you to appropriate Administration Group functions, as well as the Enterprise Manager **All Targets** page where you can view target properties.

6.3.1.2 Defining a Hierarchy

On this page you define the Administration Group hierarchy that reflects the organizational hierarchy you planned earlier and which target properties are associated with a particular hierarchy level.

On the left side of the page are two tables: Hierarchy Levels and Hierarchy Nodes.

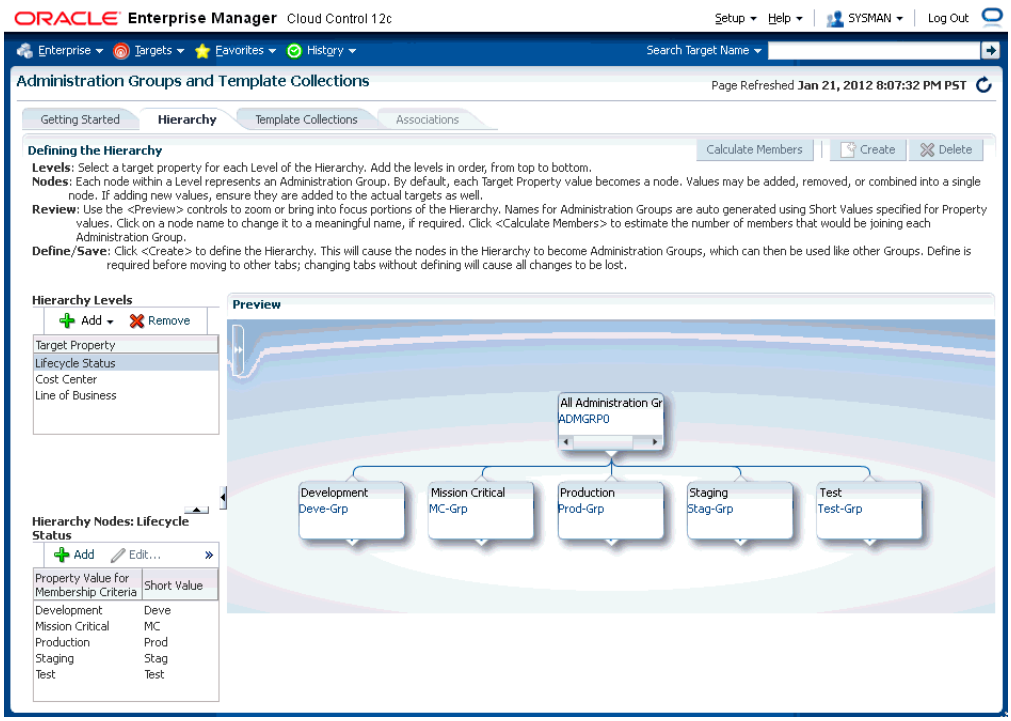

The **Hierarchy Levels** table allows you to add the target properties that define Administration Group hierarchy. The **Hierarchy Nodes** table allows you to define the values associated with the target properties in the **Hierarchy Levels** table. When you select a target property, the related property values are made available in the **Hierarchy Nodes** table, where you can add/remove/merge/split the values. In the **Hierarchy Nodes** table, each row corresponds to a single Administration Group.

Adding a Hierarchy Level

1. On the **Administration Group** page, click the **Hierarchy** tab.
2. From the **Hierarchy Levels** table, click **Add** and choose one of the available target properties. Repeat this step until you have added all target properties of interest.
3. Click on one of the newly added target properties in the **Hierarchy Levels** list. The membership values for property are displayed in the **Hierarchy Nodes** table.

   Enterprise Manager finds all existing values of the target property across all targets and displays them in the **Hierarchy Nodes** table. For some target properties, such as Lifecycle Status, predefined property values already exist and are automatically displayed in the Hierarchy Nodes table. However, property values that are not yet available, will need to be added.

   The next step shows you how to add property values.
4. From the **Hierarchy Nodes** table, click **Add**. The associated property value add dialog containing existing values from various targets displays. Add the requisite value(s). Multiple values can be specified using a comma separated list. For example, to add multiple locations such as San Francisco and Zurich, add the **Location** target property to the **Hierarchy Level** table. Select **Location** and then click **Add** in the **Hierarchy Nodes** table. The **Values for Hierarchy Nodes** dialog displays. Enter "San Francisco,Zurich" as shown in the following graphic.

Note: You cannot specify additional values for properties with predefined values (such as *Lifecycle Status*) or read-only properties (such as *Target Version*). Also, you cannot specify more than 25 property values (the **Hierarchy Nodes** table can have only 25 rows). If more than 25 property values need to be accommodated, values can be merged.

Click **OK**. The two locations "San Francisco" and "Zurich" appear as nodes in the **Preview** pane as shown in the following graphic.

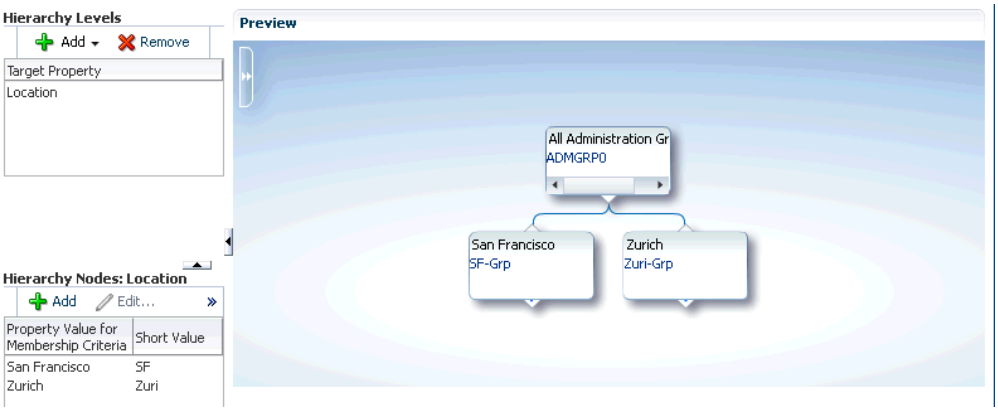

Sometimes it is useful to treat multiple property values as one. For example, if a combination of values is needed, such as *Production* or *Mission Critical* for the *Lifecycle Status* property, they need to be merged (combined into a single node).

To merge property values:

1. Select a target property from the list of chosen properties in the **Hierarchy Levels** table. The associated property values are displayed.
2. Select two or more property values by holding down the *Shift* key and clicking on the desired values.
3. Click **Merge**. If **Merge** is not visible, click **Overflow** on the **Hierarchy Nodes** table.

5. Continue adding hierarchy levels until the group hierarchy is complete. The **Preview** pane dynamically displays any changes you make to your Administration Group hierarchy.

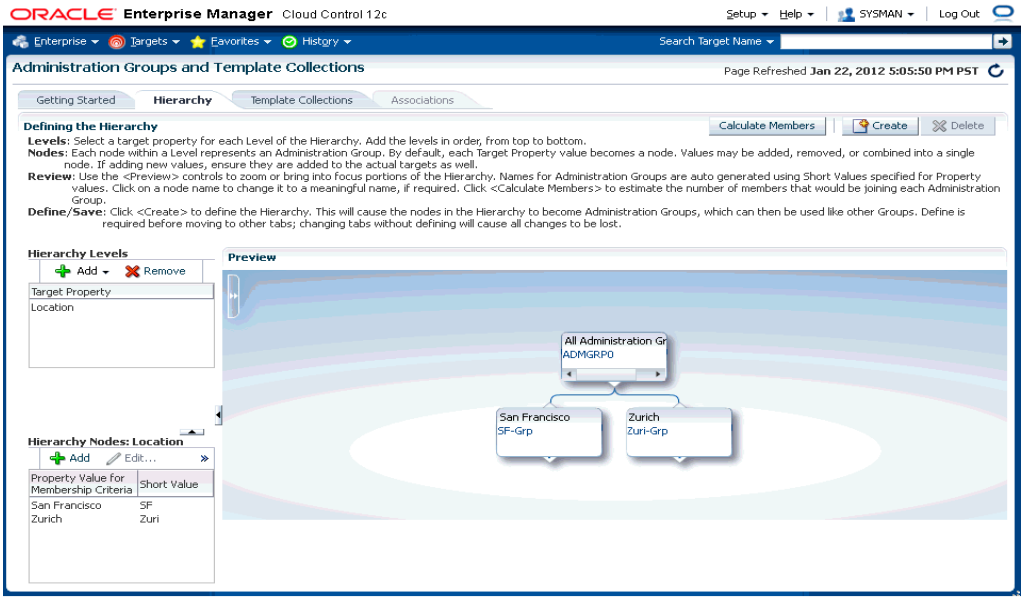

6. Click **Create** to define the hierarchy.

> **IMPORTANT:** Review and define the complete hierarchy before clicking **Create**. Once the Administration Group hierarchy is created, you cannot add additional levels to the hierarchy without deleting and recreating the entire hierarchy.

Once your Administration Group hierarchy has been created, the only change you can make is to add/remove group membership criteria property values, which equates to creating/deleting additional Administration Groups for a given level. Using the previous example, if in addition to San Francisco and Zurich you add more locations, say New York and Bangalore, you can click **Add** in the **Hierarchy Node** table to add additional locations, as shown in the following graphic.

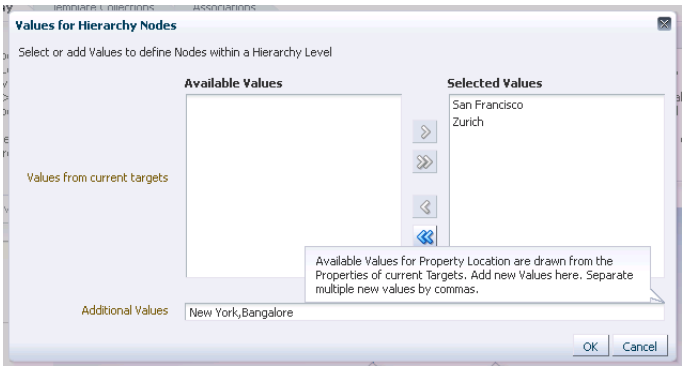

Click **Update** to save your changes.

6.3.1.3 Defining Template Collections

A Template Collection is a assemblage monitoring/management settings to be applied to targets in the Administration Group . Multiple Monitoring Templates can be added to a Template Collection that in turn is associated with an Administration Group.

However, you can only have one Monitoring Template of a particular target type in the Template Collection. The Monitoring Template should contain the complete set of metric settings for the target in the Administration Group. You should create one Monitoring Template for each type of target in the Administration Group. For example, you can have a Template Collection containing a template for database and a template for listener, but you cannot have a Template Collection containing two templates for databases. When members targets are added to an Administration Group, the template monitoring and management settings are automatically applied.

You create Template Collections when you define Administration Groups. Template collections may consist of three types of monitoring/management setting categories:

- Monitoring Templates (monitoring settings)
- Compliance Standards (compliance policy rules)
- Cloud Policies (cloud policies such as determining when to start virtual machines or scale out clusters).

When creating a Template Collection, you can use the default Monitoring Templates, compliance standards, or cloud templates supplied with Enterprise Manager or you can create your own.

To create a Template Collection:

1. Click the **Template Collections** tab. The Template Collection page displays.

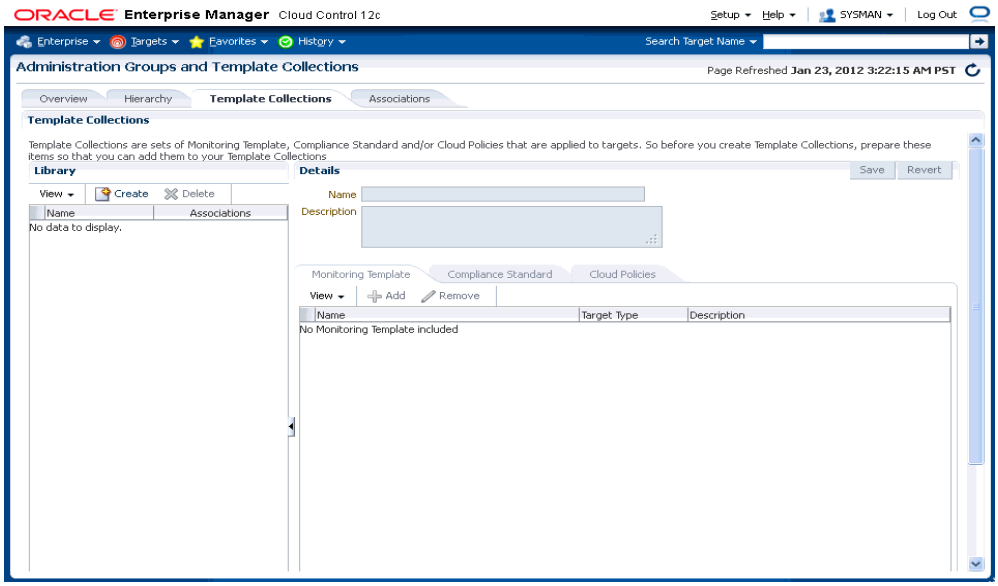

2. Click **Create**.
3. In the **Name** field, specify the Template Collection name.
4. Click the Template Collection member type you want to add (Monitoring Template, Compliance Standard, Cloud Polices). The requisite definition page appears.
5. Click **Add**. A list of available template entities appears.

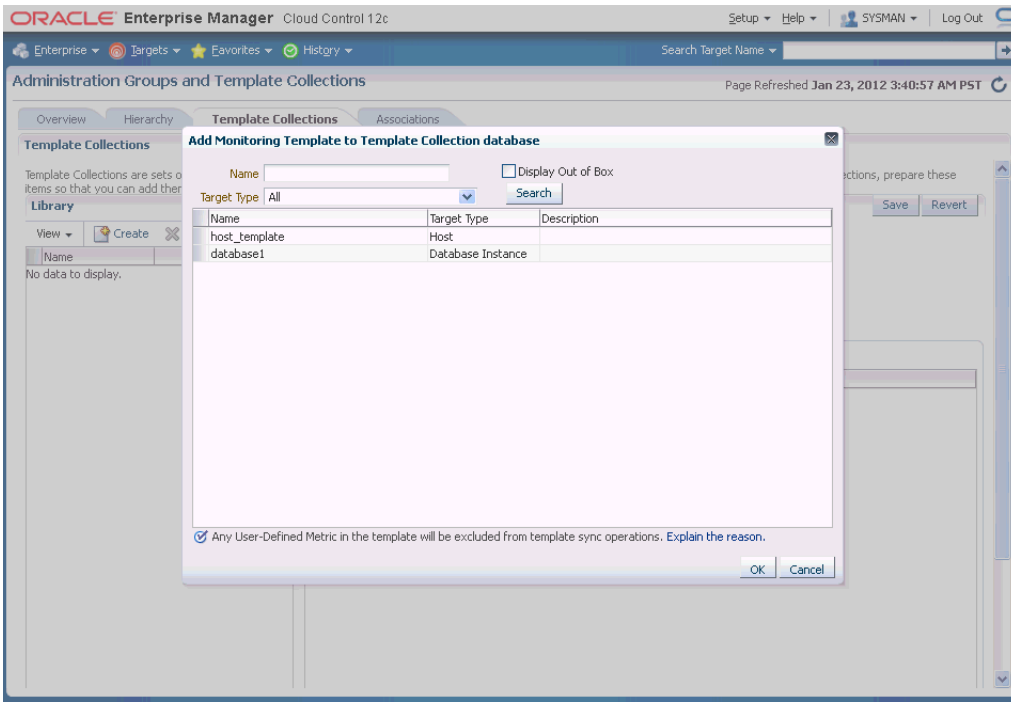

6. Select the desired template entities you want added to the Template Collection.
7. Click **OK**.
8. Continue adding template entities (Monitoring Template, Compliance Standard, Cloud Policies) as required.
9. Click **Save**. The newly defined collection appears in the **Template Collections Library**.
10. To create another Template Collection, click **Create** from the **Library** region and create and repeat steps two through eight. Repeat this process until you have created all required Template Collections.

> **Note:** When editing existing Template Collections, you can back out of any changes made during the editing session by clicking **Revert**. This restores the Template Collection to its state when it was last saved.

Required Privileges

To create a Template Collection, *Create Template Collection* resource privilege. To include a Monitoring Template into a Template Collection, you need at least *View* privilege on the specific Monitoring Template or *View Any Monitoring Template* privilege, which allows you to view any Monitoring Template and add it to the Template Collection. The following table summarizes privilege requirements for all Enterprise Manager operations related to Template Collection creation.

Enterprise Manager Operation	Minimum Privilege Requirement
Create Administration Group hierarchy.	Full Any Target Create Privilege Propagating Group

Enterprise Manager Operation	Minimum Privilege Requirement
Create Monitoring Templates.	No privileges required
Create Template Collection.	Create Template Collection (resource privilege VIEW on the Monitoring Template to be added to the Template Collection or View any Monitoring Template (resource privilege)
Create compliance standards.	Create Compliance Entity No privileges are required to view compliance standards.
Create cloud policies.	Create Any Policy View Cloud Policy
Associate Template Collection with Administration Group.	VIEW on the specific Template Collection. OPERATOR on the group
Perform on-demand synchronization.	OPERATOR on the group or MANAGE_TC_OPERATION (subset of OPERATOR)
Define global synchronization schedule.	Enterprise Manager Super Administrator privileges.
Set the value of target properties for a target (allows the target to "join" an Administration Group).	Configure Target on the specific target
Delete an Administration Group hierarchy.	Full Any Target

6.3.1.4 Associating Template Collections with Administration Groups

Once you have defined one or more Template Collections, you need to associate them to Administration Groups in the hierarchy. You can associate a Template Collection with one or more Administration Groups. As a rule, you should associate the Template Collection with the applicable Administration Group residing at the highest level in the hierarchy as the Template Collection will also be applied to targets joining any subgroup.

The **Associations** page displays the current Administration Group hierarchy diagram. Each Administration Group in the hierarchy can only be associated with one Template Collection.

Associating a Template Collection with an Administration Group

1. Click the **Template/Group Associations** tab. The Template/Group Associations page displays.

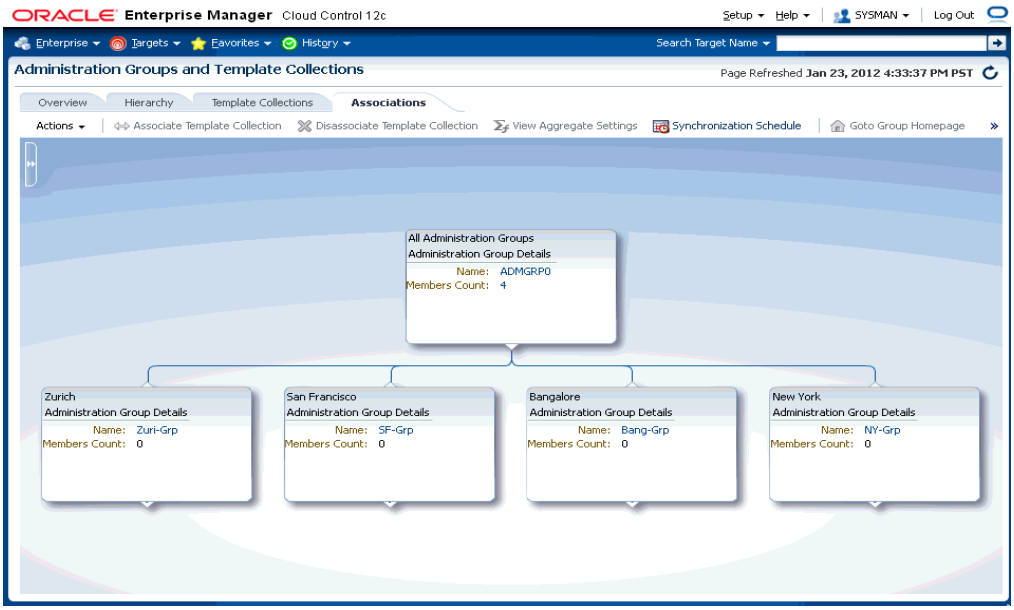

2. Select the desired Administration Group in the hierarchy.
3. Click **Associate Template Collection**. The Choose a Template Collection dialog displays.
4. Choose the desired Template Collection and click **Select**. The Template Collection details are displayed. If too much information is displayed, you can use the Zoom and Layout controls located at the upper-left corner of the hierarchy display region.

> **Note:** All sub-nodes in the hierarchy will inherit the selected Template Collection.

5. Repeat steps 1-3 until Template Collections have been associated with the desired groups.

> **Note:** The target privileges of the administrator who performs the association will be used when Enterprise Manager applies the template to the group.

> **Note:** Settings from Monitoring Templates applied at lower levels in the hierarchy override settings inherited from higher levels. This does not apply to compliance standards or cloud policies.

Setting the Global Synchronization Schedule

In order to apply the Template Collection/Administration Group association, you must set up a global sychronization schedule. This schedule is used to perform synchronization operations, such as applying templates to targets in Administration Groups. If no synchronization schedule is set up, then most synchronization changes will remain in pending mode. When there are any pending synchronization operations, they will be scheduled on the next available date based on the

synchronization schedule. Thus, it important to set the synchronization schedule as there is no default setting.

To set up the synchronization schedule:

1. Click **Synchronization Schedule**. The Synchronization Schedule dialog displays.

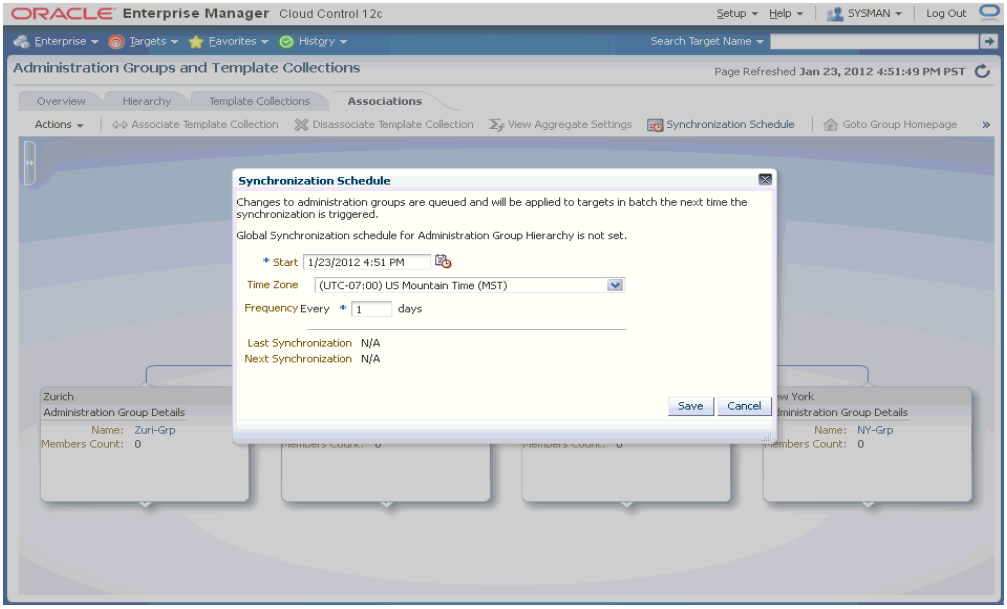

2. Click **Edit** and then choose a date and time you want the Administration Group-Template synchronization to take place. By default, the current date and time is shown.

> **Note:** You can specify a start date for synchronization operations and interval in days. Whenever there are any pending sync operations, then they will be scheduled on the next available date based on the this schedule.

3. Click **Save**.

Synchronization Sequence of Events

The following table summarizes when template synchronization operations (such as apply operations) occur on targets in Administration Groups.

Action	When Synchronization Occurs
Target is added to an Administration Group (by setting its target properties)	Immediate upon joining the Administration Group.
Template collection is associated with the Administration Group.	Targets in an Administration Group will be synchronized based on next scheduled date in global Synchronization Schedule.
Changes are made to any of the templates in the Template Collection.	Targets in an Administration Group will be synchronized based on the next scheduled date in global Synchronization Schedule.

Action	When Synchronization Occurs
Target is removed from an Administration Group (by changing its target properties).	No change in target's monitoring settings. Compliance Standards and Cloud Policies will be disassociated with the target.
Template collection is disassociated with Administration Group.	No change in target's monitoring settings. Compliance Standards and Cloud Policies will be disassociated with the target.
User performs an on-demand synchronization by clicking on the **Start Synchronization** button in the **Synchronization Status** region in the Administration Group's homepage.	Immediate synchronization operation occurs.

Viewing Synchronization Status

You can check the current synchronization status for a specific Administration Group directly from the group's homepage.

1. Select an Administration Group in the hierarchy.
2. Click **Goto Group Homepage**.
3. From the **Synchronization Status** region, you can view the status of the Monitoring Template, compliance standard, and/or cloud policies synchronization (In Sync, Pending, or Failed).

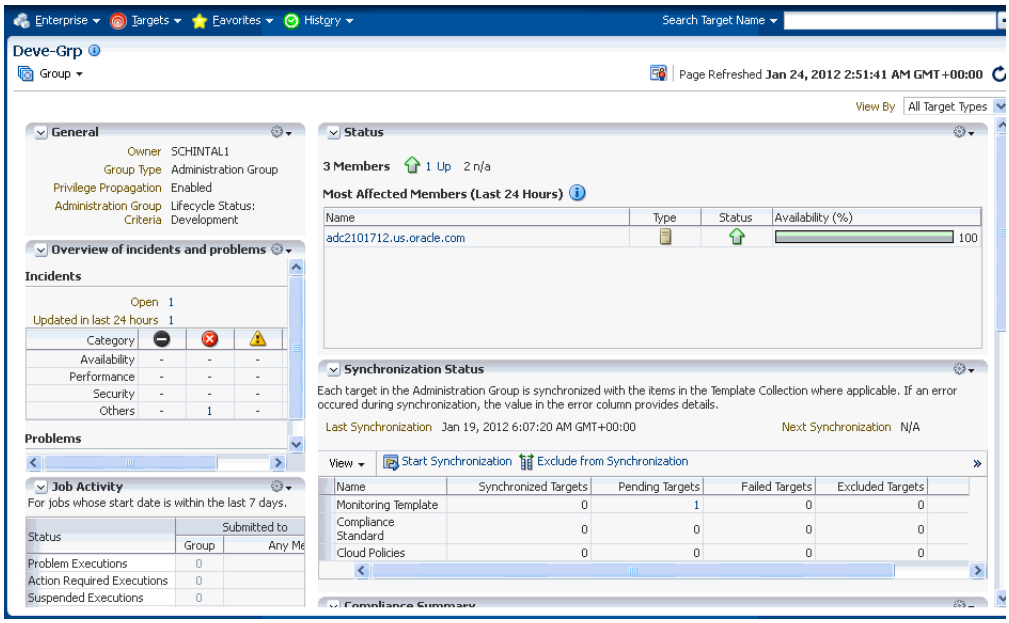

You can initiate an immediate synchronization by clicking **Start Synchronization**.

Disassociating a Template Collection from a Group

To remove a Template Collection from an Administration Group.

1. From the hierarchy diagram, select the Administration Group with the Template Collection you wish to remove.
2. Click **Disassociate Template Collection**.

The Template Collection is immediately removed. See Section , "Synchronization Sequence of Events" for more information.

Viewing Aggregate Settings

For any Administration Group, you can easily view what Template Collection components (Monitoring Templates, compliance standards, and/or cloud policies) are associated with individual group members.

> **Note:** For Monitoring Templates, the settings for a target could be a union of two or more Monitoring Templates from different Template Collections.

1. From the hierarchy diagram, select the desired Administration Group.
2. Click **View Aggregate Settings**.

   The **Aggregate Settings** page appears. This page displays all Monitoring Templates, Compliance Standards and Cloud Policies associated with the selected Administration Group (listed by target type).

Viewing the Administration Group Homepage

Like regular groups, each Administration Group has an associated group homepage providing a comprehensive overview of group member status and/or activity such as synchronization status, job activity, or critical patch advisories. To view Administration Group homepages:

1. From the hierarchy diagram, select an Administration Group.
2. Click **Goto Group Homepage.** The homepage for that particular Administration Group displays.

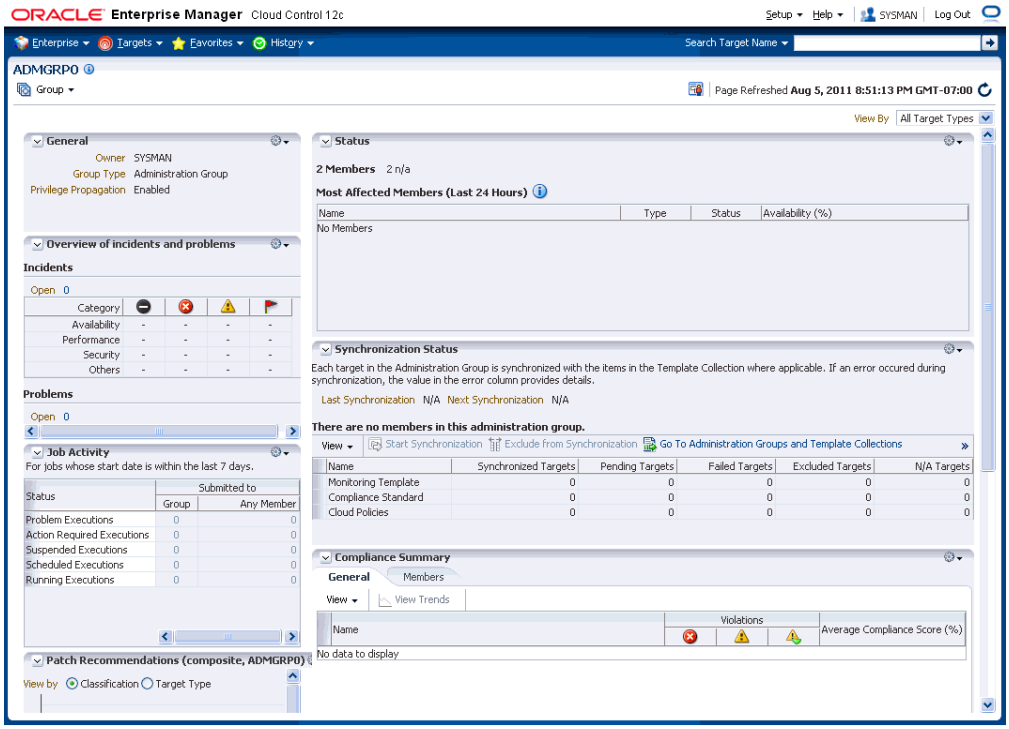

Alternatively, from the Enterprise Manger **Targets** menu, choose **Groups**. From the table, you can expand the group hierarchy.

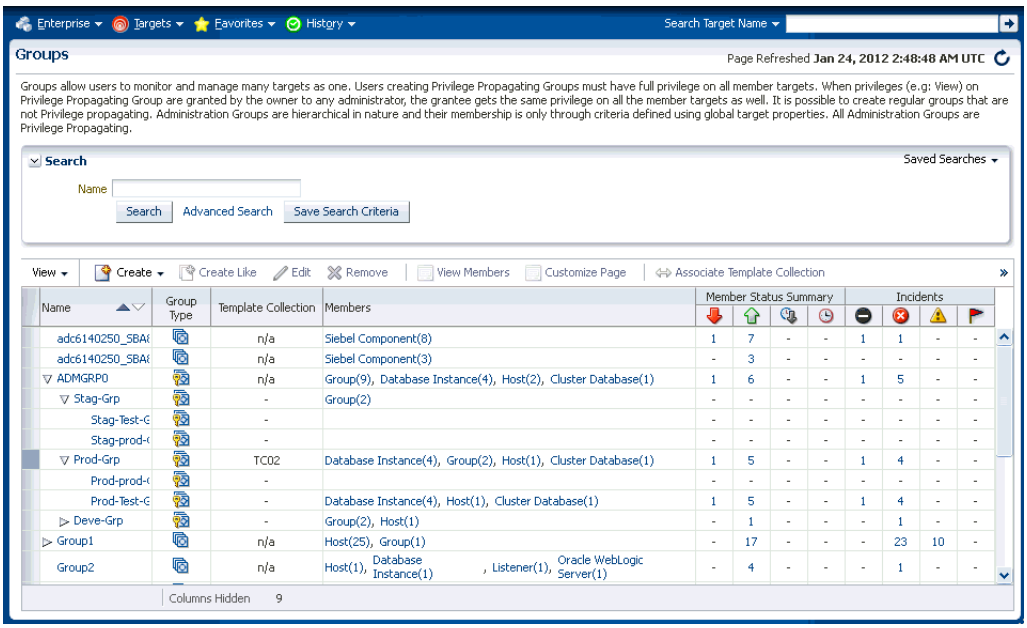

Identifying Targets Not Part of Any Administration Group

From the **Associations** page, you can determine which targets do not belong to any Administration Group by generating an *Unassigned Targets Report*.

1. From the **Actions** menu, select **Unassigned Targets Report**. The report lists all the targets that are not part of any Administration Group. The values for the target properties defining the Administration Groups hierarchy are shown.

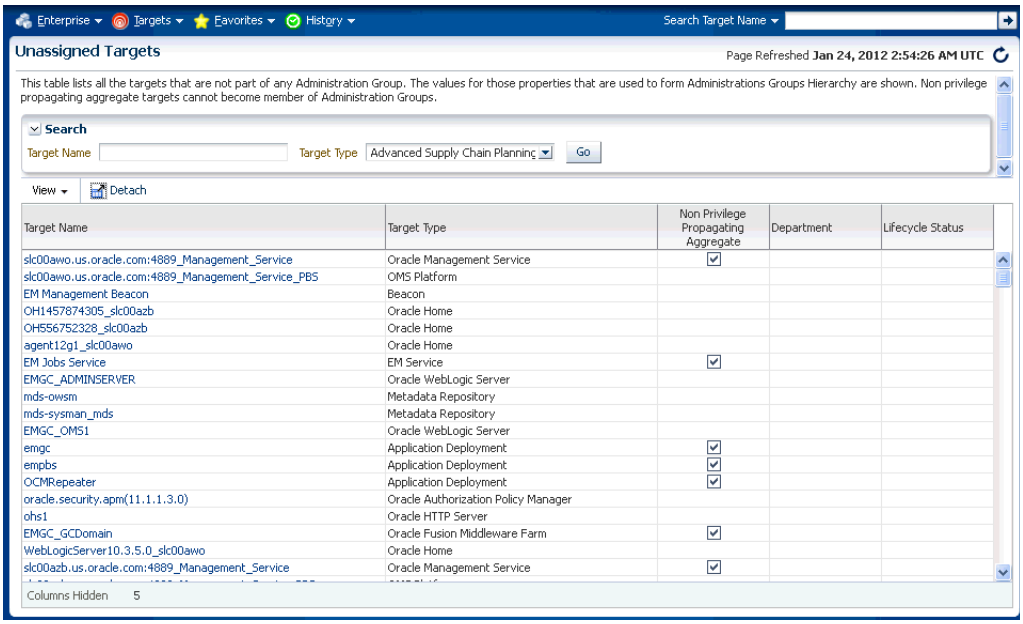

2. From the **View** menu, choose the customization options to display only the desired information.

> **Note:** The **Non-Privilege Propagating Aggregate** column indicates whether a target is a non-privilege propagating aggregate. This type of target cannot be added to an Administration Group, which are by design privilege propagating. For this reason, any aggregate target added to Administration Group must also be privilege propagating.

On this page, you can review the list to see if there any targets that need to be added to the Administration Group. Click on the target names shown in this page to access the target's **Edit Target Properties** page where you can change the target property values. After making the requisite changes and clicking OK, you are returned to the **Unassigned Targets** page

If so, you can set the target properties of these targets to add them to the appropriate Administration Groups. For information on changing target properties, see "Planning" on page 6-3.

3. Click your browser *back* button to return to the **Administration Groups and Template Collections** homepage.

6.4 Removing Administration Groups

You can completely remove an Administration Group hierarchy or just individual Administration Groups from the hierarchy. Deleting an Administration Group will not delete targets or Template Collections, but it will remove associations. Any stored membership criteria is removed. When you delete an Administration Group, any stored membership criteria is removed.

To remove the entire Administration Group hierarchy:

1. From the **Setup** menu, select **Add Target**, then select **Administration Groups**.
2. Click on the **Hierarchy Definition** tab.

3. Click **Delete**.

To remove individual Administration Groups from the hierarchy:

1. From the **Setup** menu, choose **Add Target**, then select **Administration Groups**.
2. Click on the **Hierarchy Definition** tab.
3. From the **Hierarchy Levels** table, choose the target property that corresponds to the hierarchy level containing the Administration Group to be removed.
4. From the **Hierarchy Nodes** table, select the Administration Group (**Property Value for Membership Criteria**) to be removed.
5. Choose **Remove** from the drop-down menu.

7

Metric Extensions

Metric extensions provides you with the ability to extend Oracle's monitoring capabilities to monitor conditions specific to your IT environment. This provides you with a comprehensive view of your environment. Furthermore, metric extensions allow you to simplify your IT organization's operational processes by leveraging Enterprise Manager as the single central monitoring tool for your entire datacenter instead of relying on other monitoring tools to provide this supplementary monitoring.

This chapter covers the following:

- What are Metric Extensions?
- Metric Extension Lifecycle
- Working with Metric Extensions
- Adapters
- Converting User-defined Metrics to Metric Extensions
- Metric Extension Command Line Verbs

7.1 What are Metric Extensions?

Metric extensions also allow you to create metrics on any target type and customize metric thresholds and collections. Unlike user-defined metrics (used to extend monitoring in previous Enterprise Manager releases), metric extensions allow you to create full-fledged metrics for a multitude of target types, such as:

- Hosts
- Databases
- Fusion Applications
- IBM Websphere
- Oracle Exadata databases and storage servers
- Siebel components
- Oracle Business Intelligence components

You manage metric extensions from the Metric Extensions page. This page lists all metric extensions in addition to allowing you to create, edit, import/export, and deploy metric extensions.

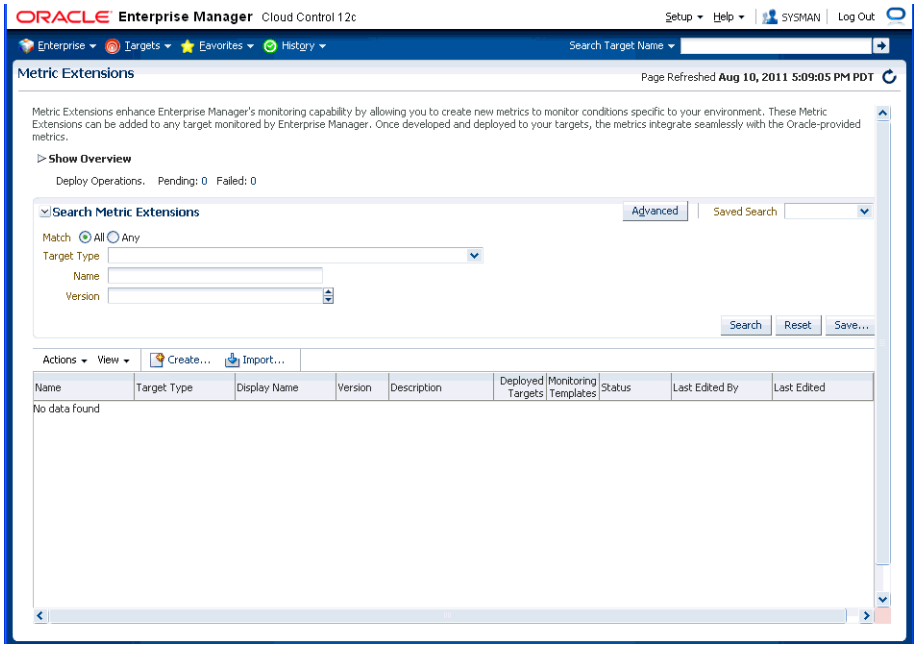

The cornerstone of the metric extension is the Oracle Integration Adapter. Adapters provide a means to gather data about targets using specific protocols. Adapter availability depends on the target type your metric extension monitors.

How Do Metric Extensions Differ from User-defined Metrics?

In previous releases of Enterprise Manager, user-defined metrics were used to extend monitoring capability in a limited fashion: user-defined metrics could be used to collect point values through execution of OS scripts and a somewhat more complex set of values (one per object) through SQL. Unlike metric extensions, user-defined metrics have several limitations:

- **Limited Integration**: If the OS or SQL user-defined metric executed custom scripts, or required atonal dependent files, the user needed to manually transfer these files to the target's file system.
- **Limited Application of Query Protocols**: OS user-defined metrics cannot model child objects of servers by returning multiple rows from a metric (this capability only exists for SQL user-defined metrics).
- **Limited Data Collection**: Full-fledged Enterprise Manager metrics can collect multiple pieces of data with a single query and reflect the associated data in alert context. In the case of user-defined metrics, multiple pieces of data can be collected by creating multiple user-defined metrics, however, it is not possible to refer to the related data when alerts are generated because they are collected separately.
- **Limited Query Protocols**: User-defined metrics can only use the "OS" and "SQL" protocols, unlike metric extensions which can use additional protocols such as SNMP and JMX.
- **Limited Target Application**: You can only create OS user-defined metrics against host targets and SQL user-defined metrics against database targets. No other target types are permitted. User-defined metrics only allow OS user-defined metrics against host targets and SQL user-defined metrics against database targets. If, for example, you want to deploy a user-defined metric against Weblogic instances in your environment, you will not be able unable to do so, making it

impossible to associate suspending of monitoring (blackouts) on these targets when servers are undergoing maintenance periods.

Most importantly, the primary difference between metric extensions and user-defined metrics is that, unlike user-defined metrics, metric dissensions are full-fledged metrics similar to Enterprise Manager out-of-box metrics. They are handled and exposed in all Enterprise Manager monitoring features as any Enterprise Manager-provided metric and will automatically apply to any new features introduced.

7.2 Metric Extension Lifecycle

Developing a metric extension involves the same three phases you would expect from any programmatic customization:

- Developing Your Metric Extension
- Testing Your Metric Extension
- Deploying and Publishing Your Metric Extension

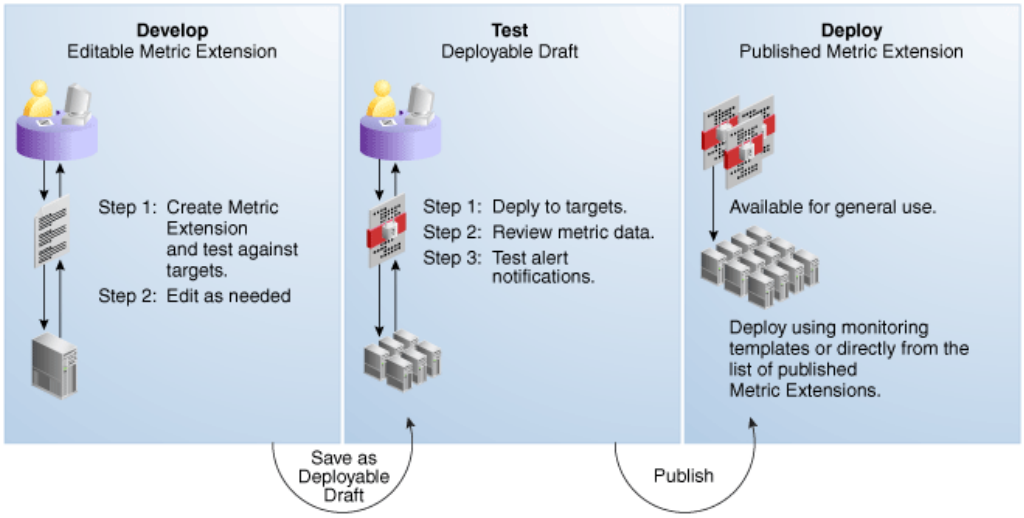

Developing Your Metric Extension

The first step is to define your monitoring requirements. This includes deciding the target type, what data needs to be collected, what mechanism (adapter) can be used to collect that data, and if elevated credentials are required. After making these decisions, you are ready to begin developing your metric extension. Enterprise Manager provides an intuitive user interface to guide you through the creation process.

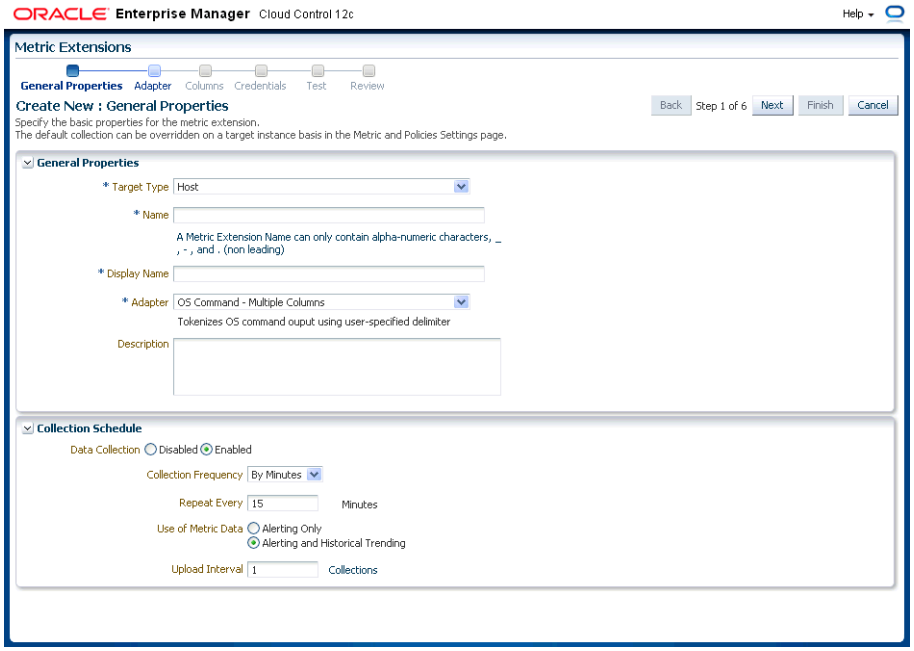

The metric extension wizard allows you to develop and refine your metric extension in a completely editable format. And more importantly, allows you to interactively test your metric extension against selected targets without having first to deploy the extension to a dedicated test environment. The **Test** page allows you to run real-time metric evaluations to ensure there are no syntactical errors in your script or metric extension definition.

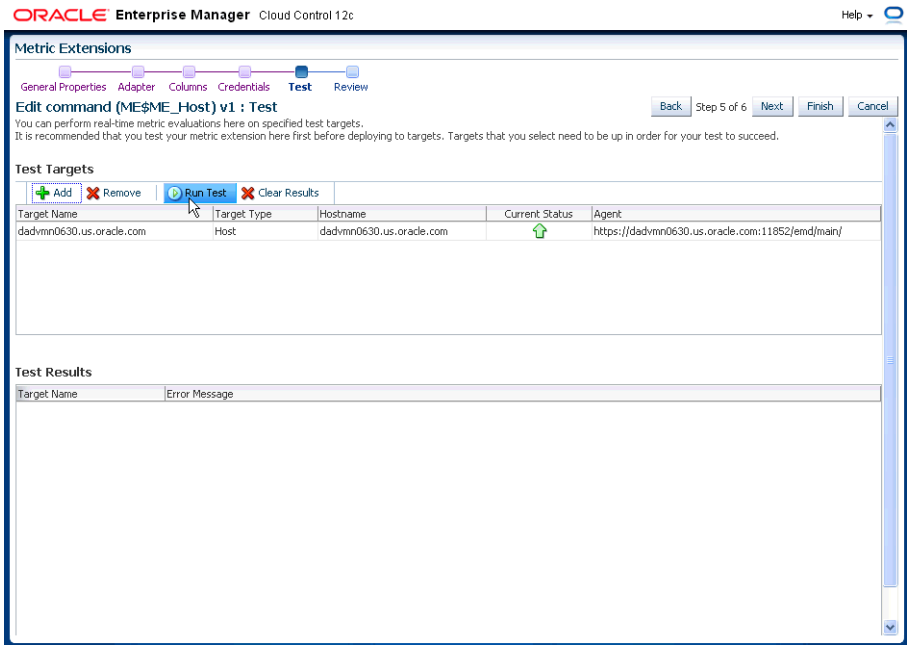

When you have completed working on your metric extension, you can click Finish to exit the wizard. The newly created metric extension appears in the Metric Extension Library where you can edit can be opened for further editing or saved as a deployable draft that can be tested against multiple targets.

Note: You can edit a metric extension only if its status is *editable*. Once it is saved as a deployable draft, you must create a new version to implement further edits.

Testing Your Metric Extension

Once your metric extension returns the expected data during real-time target testing, you are ready to test its robustness and actual behavior in Enterprise Manager by deploying it against targets and start collecting data. At this point, the metric extension is still private (only the developer can deploy to targets), but is identical to Oracle out-of-box metrics behavior wise.This step involves selecting your editable metric extension in the library and generating a deployable draft.

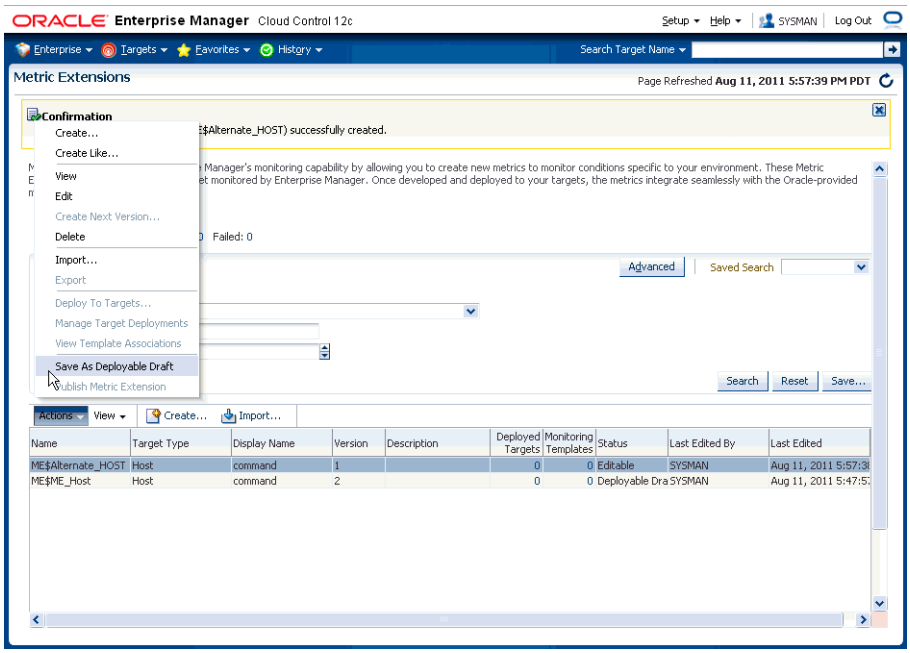

You can now deploy the metric extension to actual targets by going through the "Deploy To Targets…" action. After target deployment, you can review the metric data returned and test alert notifications. As mentioned previously, you will not be able to edit the metric extension once a deployable draft is created: You must create a new version of the metric extension.

Deploying Your Metric Extension

After rigorous testing through multiple metric extension versions and target deployments, your metric extension is ready for deployment to your production environment. Until this point, your metric extension is only viewable by you, the metric extension creator. To make it accessible to all Enterprise Manager administrators, it must be published.

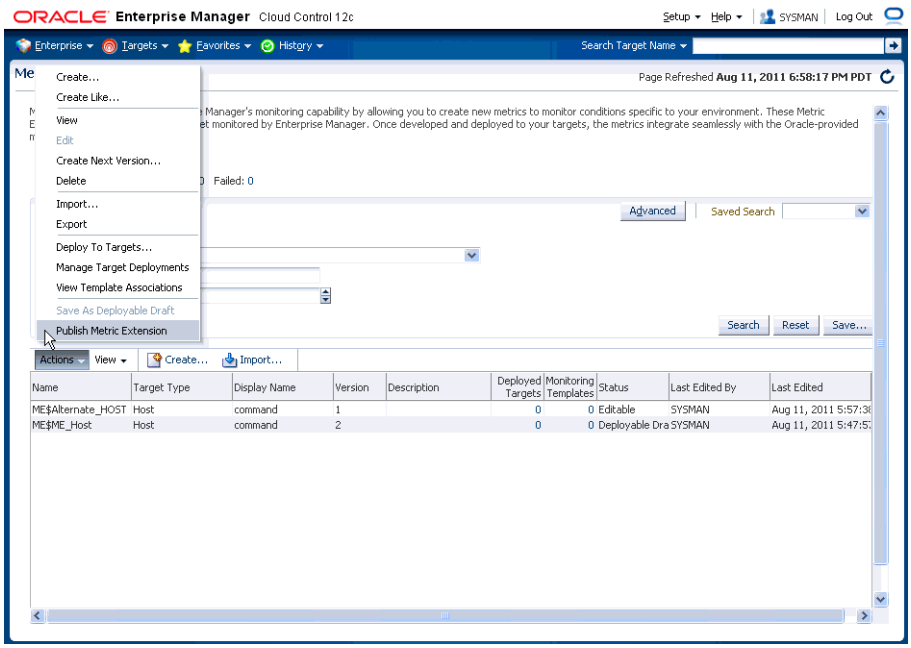

Now that your metric extension has been made public, your metric extension can be deployed to intended production targets. If you are monitoring a small number of targets, you can select the **Deploy To Targets** menu option and add targets one at a time. For large numbers of targets, you deploy metric extensions to targets using monitoring templates. An extension is added to a monitoring template in the same way a full-fledged metric is added. The monitoring template is then deployed to the targets.

Note: You cannot add metric extensions to monitoring templates before publishing the extension. If you attempt to do so, the monitoring template page will warn you about it, and will not proceed until you remove the metric extension.

7.3 Working with Metric Extensions

Most all metric extension operations can be carried out from the Metric Extension home page. If you need to perform operations on published extensions outside of the UI, Enterprise Manger also provides EM CLI verbs to handle such operations as importing/exporting metric extensions to archive files and migrating legacy user-defined metrics to metric extensions. This section covers metric extension operations carried out from the UI.

7.3.1 Administrator Privilege Requirements

In order to create, edit, view, deploy or undeploy metric extensions, you must have the requisite administrator privileges. Enterprise Manager administrators must have the following privileges:

- **Create Metric Extension**: System level access that:

  Lets administrators view and deploy metric extensions

  Allows administrators to edit and delete extensions.

- **Edit Metric Extension**: Lets users with "Create Metric Extension" privilege edit and create next versions of a particular metric extensions. The metric extension creator has this privilege by default.

  Note: This privilege must be granted on a per-metric extension basis.

- **Full Metric Extension**: In addition to the Edit Metric Extension privileges, allows deletion of a particular metric extension.

- **Manage Metrics**: Lets users deploy and un-deploy extensions on targets

  Note: The Manage Metrics privilege must be granted on a per-target basis.

7.3.2 Granting Create Metric Extension Privilege

To grant create metric extension privileges to another administrator:

1. From the **Setup** menu, select **Security**, then select **Administrators**.
2. Choose the Administrator you would like to grant the privilege to.
3. Click **Edit**.
4. Go to the Resource Privileges tab, and click **Manage Privilege Grants** for the Metric Extension resource type.
5. Under Resource Type Privileges, click the **Create Metric Extension** check box.
6. Click **Continue**, review changes, and click **Finish** in the Review tab.

7.3.3 Creating a New Metric Extension

To create a new metric extension:

1. From the **Enterprise** menu, select **Monitoring,** then select **Metric Extensions.**
2. Click **Create New**. Enterprise Manager will determine whether you have the Create Extension privilege and guide you through the creation process.
3. Decide on a metric extension name. Be aware that the name (and Display Name) must be unique across a target type.
4. Enter the general parameters.

   The selected Adapter type defines the properties you must specify in the next step of the metric extension wizard. The following adapter types are available:

   - OS Command Adapter - Single Column

     Executes the specified OS command and returns the command output as a single value. The metric result is a 1 row, 1 column table.

   - OS Command Adapter- Multiple Values

     Executes the specified OS command and returns each command output line as a separate value. The metric result is a multi-row, 1 column table.

   - OS Command Adapter - Multiple Columns

     Executes the specified OS command and parses each command output line (delimited by a user-specified string) into multiple values. The metric result is a mult-row, multi-column table.

   - SQL Adapter

     Executes custom SQL queries or function calls against single instance databases and instances on Real Application Clusters (RAC).

- SNMP (Simple Network Management Protocol) Adapter

  Allow Enterprise Manager Management Agents to query SNMP agents for Management Information Base (MIB) variable information to be used as metric data.

- JMX (Java Management Extensions) Adapter

  Retrieves JMX attributes from JMX-enabled servers and returns these attributes as a metric table.

Refer to the Adapters section for specific information on the selected adapter needed in the Adapter page (step 2) of the wizard.

Note: Be aware that if you change the metric extension Adapter, all your previous adapter properties (in Step 2) will be cleared.

5. From the Columns page, add metric columns defining the data returned from the adapter. Note that the column order should match the order the adapter returns the data in.

   - **Column Type**

     A column is either a Key column, or Data column. A Key column uniquely identifies a row in the table. For example, employee ID is a unique identifier of a table of employees. A Data column is any non-unique data in a row. For example, the first and last names of an employee.

   - **Value Type**

     A value type is Number or String. This determines the alert comparison operators that are available, and how Enterprise Manager renders collection data for this metric column.

   - **Alert Thresholds**

     The Comparison Operation, Warning, and Critical fields define an alert threshold.

   - **Alert Thresholds By Key**

     The Comparison Operation, Warning Thresholds By Key, and Critical Thresholds By Key fields allow you to specify distinct alert thresholds for different rows in a table. This option becomes available if there are any Key columns defined. For example, if your metric is monitoring CPU Usage, you can specify a different alert threshold for each distinct CPU. The syntax is to specify the key column values in a comma separated list, the "=" symbol, followed by the alert threshold. Multiple thresholds for different rows can be separated by the semi-colon symbol ";". For example, if the key columns of the CPU Usage metric are cpu_id and core_id, and you want to add a warning threshold of 50% for procecessor1, core1, and a threshold of 60% for processor2, core2, you would specify:
     procecessor1,core1=50;processor2,core2=60

   - **Manually Clearable Alert**

     If this option is set to true, then the alert will not automatically clear when the alert threshold is no longer satisfied. For example, if your metric is counting the number of errors in the system log files, and you set an alert threshold of 50, if an alert is raised once the threshold is met, the alert will not automatically clear once the error count falls back below 50. The alert will

need to be manually cleared in the Alerts UI in the target home page or Incident Manager.

- **Number of Occurrences Before Alert**

  The number of consecutive metric collections where the alert threshold is met, before an alert is raised.

- **Alert Message / Clear Message**

  The message that is sent when the alert is raised / cleared. Variables that are available for use are: %columnName%, %keyValue%, %value%, %warning_threshold%, %critical_threshold%

  You can also retrieve the value of another column by surrounding the desired column name with "%". For example, if you are creating a alert for the cpu_usage column, you can get the value of the core_temperature column by using %core_temperature%. Note that the same alert / clear message is used for warning or critical alerts.

  Note: Think carefully and make sure all Key columns are added, because you cannot create additional Key columns in newer versions of the metric extension. Once you click **Save As Deployable Draft**, the Key columns are final (edits to column display name, alert thresholds are still allowed). You can still add new Data columns in newer versions. Also be aware that some properties of existing Data columns cannot be changed later, including Column Type, Value Type, Comparison Operator (you can add a new operator, but not change an existing operator), and Manually Clearable Alert.

- **Metric Category**

  The metric category this column belongs to.

6. From the Credentials page, you can override the default monitoring credentials by using custom monitoring credential sets. By default, the metric extension wizard chooses the existing credentials used by Oracle out-of-box metrics for the particular target type. For example, metric extensions will use the dbsnmp user for database targets. You have the option to override the default credentials, by creating a custom monitoring credential set through the "emcli create_credential_set" command. Refer to the *Enterprise Manager Command Line Interface Guide* for additional details. Some adapters may use additional credentials, refer to the Adapters section for specific information.
7. From the Test page, add available test targets.
8. Click **Run Test** to validate the metric extension. The extension is deployed to the test targets specified by the user and a real-time collection is executed. Afterwards, the metric extension is automatically undeployed. The results and any errors are added to the **Test Results** region.
9. Repeat the edit /test cycle until the metric extension returns data as expected.
10. Click **Finish**.

7.3.4 Creating a New Metric Extension (Create Like)

To create a new metric extension based on an existing metric extension:

1. From the **Enterprise** menu, select **Monitoring**, then select **Metric Extensions.**

2. From the **Metric Extensions** page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select an existing metric extension.
4. From the **Actions** menu, select **Create Like**. Enterprise Manager will determine whether you have the Create Extension privilege and guide you through the creation process.
5. Make desired modifications.
6. From the **Test** page, add available test targets.
7. Click **Run Test** to validate the metric extension. The extension is deployed to the test targets specified by the user and a real-time collection is executed. Afterwards, the metric extension is automatically undeployed. The results and any errors are added to the **Test Results** region.
8. Repeat the edit /test cycle until the metric extension returns data as expected.
9. Click **Finish**.

7.3.5 Editing a Metric Extension

Before editing an existing metric extension, you must have Edit privileges on the extension you are editing or be the extension creator. Note: Once a metric extension is saved as a deployable draft, it cannot be edited, you can only create a new version.

To edit an existing metric extension:

1. From the **Enterprise** menu, select **Monitoring**, then select **Metric Extensions.**
2. From the **Metric Extensions** page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select the metric extension to be edited.
4. From the **Actions** menu, select **Edit**.
5. Update the metric extension as needed.
6. From the **Test** page, add available test targets.
7. Click **Run Test** to validate the metric extension. The extension is deployed to the test targets specified by the user and a real-time collection is executed. Afterwards, the metric extension is automatically undeployed. The results and any errors are added to the **Test Results** region.
8. Repeat the edit /test cycle until the metric extension returns data as expected.
9. Click **Finish**.

7.3.6 Creating the Next Version of an Existing Metric Extension

Before creating the next version of an existing metric extension, you must have Edit privileges on the extension you are versioning or be the extension creator.

To create next version of an existing metric extension:

1. From the **Enterprise** menu, select **Monitoring**, then select **Metric Extensions**.

2. From the Metric Extensions page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select the metric extension to be versioned.
4. From the **Actions** menu, select **Create Next Version**.
5. Update the metric extension as needed. The target type, and extension name cannot be edited, but all other general properties can be modified. There are also restrictions on metric columns modifications. See Note in Creating a New Metric Extension section for more details.
6. From the Test page, add available test targets.
7. Click **Run Test** to validate the metric extension. The extension is deployed to the test targets specified by the user and a real-time collection is executed. Afterwards, the metric extension is automatically undeployed. The results and any errors are added to the **Test Results** region.
8. Repeat the edit /test cycle until the metric extension returns data as expected.
9. Click **Finish**.

7.3.7 Importing a Metric Extension

Metric extensions can be converted to portable, self-contained packages that allow you to move the metric extension to other Enterprise Manager installations, or for storage/backup. These packages are called Metric Extension Archives (MEA) files.

MEA files are zip files containing all components that make up the metric extension: metric metadata, collections, and associated scripts/jar files. Each MEA file can contain only one metric extension. To add the metric extension back to your Enterprise Manager installation, you must import the metric extension from the MEA.

To import a metric extension from an MEA file:

1. From the **Enterprise** menu, select **Monitoring**, then select **Metric Extensions.**
2. Click **Import**.
3. Browse to file location, and select the MEA file. Enterprise Manager checks if the target type and metric extension name combination is already used in the system. If not, the system will create a new metric extension. If the extension name is already in use, the system will attempt to create a new version of the existing extension using the MEA contents. This will require the MEA to contain a superset of all the existing metric extension's metric columns. You also have the option to rename the metric extension.
4. Clicking on OK creates the new metric extension or the new version of an existing metric extension.
5. From the **Actions** menu, select **Edit** to verify the entries.
6. From the **Test** page, add available test targets.
7. Click **Run Test** to validate the metric extension. The extension is deployed to the test targets specified by the user and a real-time collection is executed. Afterwards, the metric extension is automatically undeployed. The results and any errors are added to the **Test Results** region.
8. Repeat the edit /test cycle until the metric extension returns data as expected.
9. Click **Finish**.

7.3.8 Exporting a Metric Extension

Existing metric extensions can be package as self-contained zip files (exported) for portability and/or backup and storage.

To export an existing metric extension:

1. From the **Enterprise** menu, select **Monitoring**, then select **Metric Extensions.**
2. From the **Metric Extensions** page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select the metric extension to be exported.
4. From the **Actions** menu, select **Export**. Enterprise Manager prompts you to enter the name and location of the MEA file that is to be created.
5. Enter the name and location of the package. Enterprise Manager displays the confirmation page after the export is complete. Note: You can only export the production version. Note: You can only export Deployable Draft and Published metric extension versions.
6. Confirm the export file is downloaded.

7.3.9 Deleting a Metric Extension

Initiating the deletion of a metric extension is simple. However, the actual deletion triggers a cascade of activity by Enterprise Manager to completely purge the metric extension from the system. This includes closing open metric alerts, and purging collected metric data (if the latest metric extension version is deleted).

Before a metric extension version can be deleted, it must be undeployed from all targets, and removed from all monitoring templates (including templates in pending apply status).

To delete a metric extension:

1. From the **Enterprise** menu, select **Monitoring**, then select **Metric Extensions.**
2. From the **Metric Extensions** page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select the metric extension that is to be deleted.
4. From the **Actions** menu, select **Delete**. Enterprise Manager prompts you to confirm the deletion.
5. Confirm the deletion.

7.3.10 Granting Edit/Full Access to Metric Extensions

Before an Enterprise Manager administrator can be edit, or delete a metric extension created by another administrator, they must have been granted requisite access privileges. Edit privilege allows editing and creating next versions of the extension, and Full privilege allows the above operations and deletion of the extension.

To grant edit/full access to an existing metric extension to another administrator:

1. From the **Setup** menu, select **Security**, then select **Administrators**.
2. Choose the Administrator you would like to grant access to.
3. Click **Edit**.

4. Go to the Resource Privileges tab, and click **Manage Privilege Grants for the Metric Extension** resource type.
5. Under **Resource Privileges**, you can search for and add existing metric extensions. Add the metric extensions you would like to grant privileges to. This allows the user to edit and create next versions of the metric extension.
6. If you would additionally like to allow delete operations, then click the pencil icon in the **Manage Resource Privilege Grants** column, and select **Full Metric Extension** privilege in the page that shows up.
7. Click **Continue**, review changes, and click **Finish** in the review tab.

7.3.11 Deploying Metric Extensions to a Group of Targets

A metric extension must be deployed to a target in order for it to begin collecting data.

To deploy a metric extension to one or more targets:

1. From the **Enterprise** menu, select **Monitoring**, then select **Metric Extensions.**
2. From the **Metric Extensions** page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select the metric extension that is to be deployed.
4. From the **Actions** menu, select **Manage Target Deployments**. The **Manage Target Deployments** page appears showing you on which target(s) the selected metric extension is already deployed.
5. Return to the **Metric Extensions** page.
6. Select the metric extension.
7. From the **Actions** menu, select **Deploy to Targets**. Enterprise Manager determines whether you have "Manage Target Metrics" privilege, and only those targets where you do show up in the target selector.
8. Add the targets where the metric extension is to be deployed and click Submit. Enterprise Manager submits a job deploying the metric extension to each of the targets. A single job is submitted per deployment request.
9. You are automatically redirected to the Pending Operations page, which shows a list of currently scheduled, executing, or failed metric extension deploy operations. Once the deploy operation completes, the entry is removed from the pending operations table.

7.3.12 Updating Older Versions of Metric Extensions Already deployed to a Group of Targets

When a newer metric extension version is published, you may want to update any older deployed instances of the metric extension.

To update old versions of the metric extension already deployed to targets:

1. From the **Enterprise** menu, select **Monitoring**, then select **Metric Extensions.**
2. From the **Metric Extensions** page, determine which extensions are accessible. The page displays the list of metric extensions along with target type, owner, production version and deployment information.
3. Select the metric extension to be upgraded.

4. From the **Actions** menu, select **Manage Target Deployments**. The **Manage Target Deployments** page appears showing a list of targets where the metric extension is already deployed.
5. Select the list of targets where the extension is to be upgraded and click **Upgrade**. Enterprise Manager submits a job for the deployment of the newest Published metric extension to the selected targets. A single job is submitted per deployment request.
6. You are automatically redirected to the Pending Operations page, which shows a list of currently scheduled, executing, or failed metric extension deploy operations. Once the deploy operation completes, the entry is removed from the pending operations table.

7.4 Adapters

Oracle Integration Adapters provide comprehensive, easy-to-use monitoring connectivity with a variety of target types. The adapter enables communication with an enterprise application and translates the application data to standards-compliant XML and back.

The metric extension target type determines which adapters are made available from the UI. A complete list of all adapters is shown below.

- OS Command Adapter - Single Column
- OS Command Adapter- Multiple Values
- OS Command Adapter - Multiple Columns
- SQL Adapter
- SNMP (Simple Network Management Protocol) Adapter
- JMX Adapter

7.4.1 OS Command Adapter - Single Column

Executes the specified OS command and returns the command output as a single value. The metric result is a 1 row, 1 column table.

Basic Properties

The complete command line will be constructed as: Command + Script + Arguments.

- **Command** - The command to execute. For example, `%perlBin%/perl`. The complete command line will be constructed as: Command + Script + Arguments.
- **Script** - A script to pass to the command. For example, `%scriptsDir%/myscript.pl`. You can upload custom files to the agent, which will be accessible under the `%scriptsDir%` directory.
- **Arguments** - Additional arguments to be appended to the Command.

Advance Properties

- **Input Properties** - Additional properties can be passed to the command through its standard input stream. This is usually used for secure content, such as username or passwords, that you don't want to be visible to other users. For example, you can add the following Input Property:

`Name=targetName,Value=%NAME%`

which the command can read through it's standard input stream as "STDINtargetName=<target name>".

- **Environment Variables** - Additional properties can be accessible to the command from environment variables. For example, you can add Environment Variable: `Name=targetType,Value="%TYPE%"`, and the command can access the target type from environment variable "ENVtargetType".

Credentials

- **Host Credentials** - The credential used to launch the OS Command.
- **Input Credentials** - Additional credentials passed to the OS Command's standard input stream.

Example 1

Read the contents of a log file, and dump out all lines containing references to the target.

- **Approach 1** - Use the grep command, and specify the target name using %NAME% parameter.

  ```
  Command = /bin/grep %NAME% mytrace.log
  ```

- **Approach 2** - Run a perl script

  ```
  Command = %perlBin%/perl
  ```

  ```
  Script = %scriptsDir%/filterLog.pl
  ```

  Input Properties:

  ```
  targetName = %NAME%
  ```

  ```
  targetType = %TYPE%
  ```

filterLog.pl:

```
require "emd_common.pl";

my %stdinVars = get_stdinvars();
my $targetName = $stdinVars{"targetName"};
my $targetType = $stdinVars{"targetType"};
open (MYTRACE, mytrace.log);
foreach $line (<MYTRACE >)
{
    # Do line-by-line processing
}

close (MYTRACE);
```

Example 2

Connect to a database instance from a PERL script and query the HR.JOBS sample schema table.

- Approach 1 - Pass credentials from target type properties into using Input Properties:

  ```
  Command = %perlBin%/perl
  ```

  ```
  Script = %scriptsDir%/connectDB.pl
  ```

Input Properties:

```
EM_DB_USERNAME = %Username%
EM_DB_PASSWORD = %Password%
EM_DB_MACHINE = %MachineName%
EM_DB_PORT = %Port%
EM_DB_SID = %SID%
```

connectDB.pl

```
use DBI;
require "emd_common.pl";

my %stdinVars = get_stdinvars();
my $dbUsername = $stdinVars{"EM_DB_USERNAME"};
my $dbPassword = $stdinVars{"EM_DB_PASSWORD"};
my $dbMachine = $stdinVars{"EM_DB_MACHINE"};
my $dbPort = $stdinVars{"EM_DB_PORT"};
my $dbSID = $stdinVars{"EM_DB_SID"};

my $dbAddress =
"(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=$dbMachine)(Port=$dbPort))(CONNECT_
DATA=(SID=$dbSID)))";

# Establish Target DB Connection
my $db = DBI->connect('dbi:Oracle:', "$dbUsername@".$dbAddress, "$dbPassword",
    {PrintError => 0, RaiseError => 0, AutoCommit => 0})
    or die (filterOraError("em_error=Could not connect to
$dbUsername/$dbAddress: $DBI::errstr\n", $DBI::err));

my $query = "SELECT JOB_TITLE, MIN_SALARY FROM HR.JOBS";
my $st = $db->prepare($query);
$st->execute();

while ( my ($job_title, $min_sal) = $st->fetchrow_array() )
{
    print "$job_title|$min_sal\n";
}

$db->disconnect
    or warn "disconnect $DBI::errstr\n";

exit 0;
```

- Approach 2 - Pass monitoring credential set using Input Credentials

```
Command = %perlBin%/perl
Script = %scriptsDir%/connectDB.pl
```

Input Credentials:

```
dbCreds = MyCustomDBCreds
```

connectDB.pl

```
use DBI;

require "emd_common.pl";
```

```
my %stdinVars = get_stdinvars();
my $credType = getCredType("dbCred", \%stdinVars);
my %credProps = getCredProps("dbCreds", \%stdinVars);
my $dbUsername = $credProps{"DBUserName"};
my $dbPassword = $credProps{"DBPassword"};
```

Example 3

Overriding default monitoring credentials by creating and using a custom monitoring credential set for host target.

We will only show a simple example here. Refer to the Credentials section of the Administrator's Guide for more details on creating and configuring custom monitoring credential sets.

Creating host credentials for the host target type:

```
> emcli create_credential_set -set_name=myCustomCreds -target_type=host -auth_
target_type=host -supported_cred_types=HostCreds -monitoring -description='My
Custom Credentials'
```

When you go to the Credentials page of the Metric Extension wizard, and choose to "Specify Credential Set" for Host Credentials, you will see "My Custom Credentials" show up as an option in the drop down list.

Note that this step only creates the Monitoring Credential Set for the host target type, and you need to set the credentials on each target you plan on deploying this metric extension to. You can set credentials from Enterprise Manager by going to Setup, then Security, then Monitoring Credentials. Alternatively, this can be done from the command line.

```
> emcli set_monitoring_credential -target_name=target1 -target_type=host -set_
name=myCustomCreds -cred_type=HostCreds -auth_target_type=host
-attributes='HostUserName:myusername;HostPassword:mypassword'
```

7.4.2 OS Command Adapter- Multiple Values

Executes the specified OS command and returns each command output line as a separate value. The metric result is a multi-row, 1 column table.

For example, if the command output is:

```
em_result=foo
em_result=bar
```

then three columns are populated with values 1,2,3 respectively.

Basic Properties

- **Command** - The command to execute. For example, %perlBin%/perl.
- **Script** - A script to pass to the command. For example, `%scriptsDir%/myscript.pl`. You can upload custom files to the agent, which will be accessible under the `%scriptsDir%` directory.
- **Arguments** - Additional arguments to be appended to the Command.
- **Starts With** - The starting string of metric result lines.

  Example: If the command output is:

  ```
  em_result=4354
  update
  ```

```
test
```

setting *Starts With* = *em_result* specifies that only lines starting with *em_result* will be parsed.

Advanced Properties

- **Input Properties** - Additional properties to be passed to the command through its standard input stream. For example, you can add Input Property: Name=targetName, Value=%NAME%, which the command can read through its standard input stream as "STDINtargetName=<target name>". See usage examples in OS Command Adapter - Single Columns.
- **Environment Variables** - Additional properties can be accessible to the command from environment variables. For example, you can add Environment Variable: Name=targetType, Value="%TYPE%", and the command can access the target type from environment variable "ENVtargetType". See usage examples in OS Command Adapter - Single Columns.

Credentials

- **Host Credentials** - The credential used to launch the OS Command. See usage examples in OS Command Adapter - Single Columns.
- **Input Credentials** - Additional credentials passed to the OS Command's standard input stream. See usage examples in OS Command Adapter - Single Columns.

7.4.3 OS Command Adapter - Multiple Columns

Executes the specified OS command and parses each command output line (delimited by a user-specified string) into multiple values. The metric result is a mult-row, multi-column table.

Example: If the command output is

```
em_result=1|2|3
em_result=4|5|6
```

and the Delimiter is set as "|", then there are two rows of three columns each:

Table 7–1 Multi-Column Output

1	2	3
4	5	6

Basic Properties

The complete command line will be constructed as: Command + Script + Arguments

- **Command** - The command to execute. For example, %perlBin%/perl.
- **Script** - A script to pass to the command. For example, %scriptsDir%/myscript.pl. You can upload custom files to the agent, which will be accessible under the %scriptsDir% directory.
- **Arguments** - Additional arguments.
- **Delimiter** - The string used to delimit the command output.
- **Starts With** - The starting string of metric result lines.

Example: If the command output is

```
em_result=4354
foo
bar
```

setting *Starts With* = *em_result* specifies that only lines starting with *em_result* will be parsed.

- **Input Properties** - Additional properties can be passed to the command through its standard input stream. For example, you can add Input Property: *Name=targetName*, *Value=%NAME%*, which the command can read through it's standard input stream as *STDINtargetName=<target name>*. To specify multiple Input Properties, enter each property on its own line.
- **Environment Variables** - Additional properties can be accessible to the command from environment variables. For example, you can add Environment Variable: *Name=targetType*, *Value="%TYPE%*, and the command can access the target type from environment variable "ENVtargetType".

Advanced Properties

- **Input Properties** - Additional properties can be passed to the command through its standard input stream. For example, you can add Input Property: Name=targetName, Value=%NAME%, which the command can read through its standard input stream as STDINtargetName=<target name>. See usage examples in OS Command Adapter - Single Columns.
- **Environment Variables** - Additional properties can be accessible to the command from environment variables. For example, you can add Environment Variable: *Name=targetType*, *Value="%TYPE%*, and the command can access the target type from environment variable "ENVtargetType". See usage examples in OS Command Adapter - Single Columns.

Credentials

- **Host Credentials** - The credential used to launch the OS Command. See usage examples in OS Command Adapter - Single Columns
- **Input Credentials** - Additional credentials passed to the OS Command's standard input stream. See usage examples in OS Command Adapter - Single Columns.

7.4.4 SQL Adapter

Executes custom SQL queries or function calls supported against single instance databases and instances on Real Application Clusters (RAC).

Properties

- **SQL Query** - The SQL query to execute. Normal SQL statements should not be semi-colon terminated. For example, SQL Query = "select a.ename, (select count(*) from emp p where p.mgr=a.empno) directs from emp a". PL/SQL statements are also supported, and if used, the "Out Parameter Position" and "Out Parameter Type" properties should be populated.
- **SQL Query File** - A SQL query file. Note that only one of "SQL Query" or "SQL Query File" should be used. For example, %scriptsDir%/myquery.sql. You can upload custom files to the agent, which will be accessible under the %scriptsDir% directory.

- **Transpose Result** - Transpose the SQL query result.
- **Bind Variables** - Declare bind variables used in normal SQL statements here. For example, if the SQL Query = "select a.ename from emp a where a.mgr = :1", then you can declare the bind variable as Name=1, Value=Bob.
- **Out Parameter Position** - The bind variable used for PL/SQL output. Only integers can be spcified.

  Example: If the SQL Query is

```
DECLARE
    l_output1 NUMBER;
    l_output2 NUMBER;
BEGIN
    .....
    OPEN :1 FOR
        SELECT l_output1, l_output2 FROM dual;
END;
```

  you can set Out Parameter Position = 1, and Out Parameter Type = SQL_CURSOR
- **Out Parameter Type** - The SQL type of the PL/SQL output parameter. See comment for Out Parameter Position

Credentials

- **Database Credentials** - The credential used to connect to the database.

Example

Overiding default monitoring credentials by creating and using a custom monitoring credential set for database target.

We will only show a simple example here. Refer to the Credentials guide of the Admin Guide for more details on creating and configuring custom monitoring credential sets.

Creating host credentials for the database target type:

```
> emcli create_credential_set -set_name=myCustomDBCreds -target_type=oracle_
database -auth_target_type=oracle_database -supported_cred_types=DBCreds
-monitoring -description='My Custom DB Credentials'
```

When you go to the Credentials page of the Metric Extension wizard, and choose to "Specify Credential Set" for Database Credentials, you will see "My Custom DB Credentials" show up as an option in the drop down list.

Note that this step only creates the Monitoring Credential Set for the host target type, and you need to set the credentials on each target you plan on deploying this metric extension to. You can set credentials from Enterprise Manager by going to **Setup**, then selecting **Security**, then selecting **Monitoring Credentials**. Alternatively, this can be performed using the Enterprise Manager Command Line Interface.

```
> emcli set_monitoring_credential -target_name=db1 -target_type=oracle_database
-set_name=myCustomDBCreds -cred_type=DBCreds -auth_target_type=oracle_database
-attributes='DBUserName:myusername;DBPassword:mypassword'
```

7.4.5 SNMP (Simple Network Management Protocol) Adapter

Allow Enterprise Manager Management Agents to query SNMP agents for Management Information Base (MIB) variable information to be used as metric data.

Basic Properties

- **Object Identifiers (OIDs)**: Object Identifiers uniquely identify managed objects in a MIB hierarchy. One or more OIDs can be specified. The SNMP adapter will collect data for the specified OIDs. For example, 1.3.6.1.4.1.111.4.1.7.1.1

Advanced Properties

- **Delimiter** - The delimiter value used when specifying multiple OID values for an OID's attribute. The default value is space or \n or \t
- **Tabular Data** - Indicates whether the expected result for a metric will have multiple rows or not. Possible values are TRUE or FALSE. The default value is FALSE
- **Contains V2 Types** - Indicates whether any of the OIDs specified is of SNMPV2 data type. Possible values are TRUE or FALSE. The default value is FALSE. For example, if an OID value specified is of counter64 type, then this attribute will be set to TRUE.

7.4.6 JMX Adapter

Retrieves JMX attributes from JMX-enabled servers and returns these attributes as a metric table.

Properties

- **Metric** -- The MBean ObjectName or ObjectName pattern whose attributes are to be queried. Since this is specified as metric metadata, it needs to be instance-agnostic. Instance-specific key properties (such as *servername*) on the MBean ObjectName may need to be replaced with wildcards.
- **ColumnOrder** -- A semi-colon separated list of JMX attributes in the order they need to be presented in the metric.

Advanced Properties

- **IdentityCol** -- The MBean key property that needs to be surfaced as a column when it is not available as a JMX attribute. For example:

  ```
  com.myCompany:Name=myName,Dept=deptName, prop1=prop1Val, prop2=prop2Val
  ```

  In this example, setting *identityCol* as *Name;Dept* will result in two additional key columns representing Name and Dept besides the columns representing the JMX attributes specified in the *columnOrder* property.

- **AutoRowPrefix** -- Prefix used for an automatically generated row. Rows are automatically generated in situations where the MBean *ObjectName* pattern specified in metric property matches multiple MBeans and none of the JMX attributes specified in the *columnOrder* are unique for each. The *autoRowId* value specified here will be used as a prefix for the additional key column created. For example, if the metric is defined as:

  ```
  com.myCompany:Type=CustomerOrder,* columnOrder
  ```

  is

  ```
  CustomerName;OrderNumber;DateShipped
  ```

  and assuming *CustomerName;OrderNumber;Amount* may not be unique if an order is shipped in two parts, setting *autoRowId* as "ShipItem-" will populate an

additional key column for the metric for each row with ShipItem-0, ShipItem-1, ShipItem-2...ShipItem-n.

- **Metric Service** -- True/False. Indicate whether *MetricService* is enabled on a target Weblogic domain. This property would be false (unchecked) in most cases for Metric Extensions except when metrics that are exposed via the Oracle DMS MBean needs to be collected. If *MetricService* is set to true, then the basic property *metric* becomes the *MetricService* table name and the basic property *columnOrder* becomes a semicolon-separated list of column names in the *MetricService* table.

> **Note:** Refer to the Monitoring Using Web Services and JMX chapter in the *Oracle® Enterprise Manager Extensibility Programmer's Reference* for an in-depth example of creating a JMX based Metric Extension.

7.5 Converting User-defined Metrics to Metric Extensions

For targets monitored by Enterprise Manager 12c Agents, both older user-defined metrics and metric extensions will be supported. After release 12c, only metric extensions will be supported. If you have existing user-defined metrics, it is recommended that you migrate them to metric extensions as soon as possible to prevent potential monitoring disruptions in your managed environment.

Migration of user-defined metric definitions to metric extensions is not automatic and must be initiated by an administrator. The migration process involves migrating user-defined metric metadata to metric extension metadata.

> **Note:** Migration of collected user-defined metric historic data is not supported.

After the user-defined metric is migrated to the metric extension and the metric extension has been deployed successfully on the target, the user-defined metric should be either disabled or deleted. Disabling the collection of the user-defined metric will retain the metadata definition of the user-defined metric) but will clear all the open alerts, remove the metric errors and prevent further collections of the user-defined metric. Deleting the user-defined metric will delete the metadata, historic data, clear open alerts and remove metric errors.

7.5.1 Overview

The User Defined Metric (UDM) to Metric Extension (ME) migration replaces an existing UDM with a new or existing ME. The idea behind the migration process is to consolidate UDMs with the same definition that have been created on different targets into a single ME. In addition, MEs support multiple metric columns, allowing the user to combine multiple related UDMs into a single ME.

This migration process is comprised of the following steps:

1. Identify the UDMs that need to be migrated.
2. Use the provided EM CLI commands to create or select a compatible metric extension.
3. Test and publish the metric extension.
4. Deploy the metric extension to all targets and templates where the original UDMs are located. Also update the existing notification rules to refer to the ME.

5. Delete the original UDMs. Note that the historical data and alerts from the old UDM is still accessible from the UI, but the new ME will not inherit them.

Note that the credentials being used by the UDM are NOT migrated to the newly created ME. The user interface allows a user to specify the credential set required by the metric extension. If the ME does not use the default monitoring credentials, the user will need to create a new credential set to accommodate the necessary credentials through the relevant EM CLI commands. This set will then be available in the credentials page of the metric extension wizard.

The migration process is categorized by migration sessions. Each session is responsible for migrating one or more UDMs. The process of migrating an individual session is referred to as a task. Therefore, a session is comprised of one or more tasks. In general terms, the migration involves creating a session and providing the necessary input to complete each tasks within that session. The status of the session and tasks is viewable throughout the workflow.

7.5.2 Commands

A number of EM CLI commands are responsible for completing the various steps of this process. For a more detailed explanation of the command definition, please use the 'EM CLI help <command>' option.

- **list_unconverted_udms** - Lists the UDMs that have yet to be migrated and not in a session
- **create_udmmig_session** - Creates a session to migrate one or more UDMs
- **udmmig_summary** - Lists the migration sessions in progress
- **udmmig_session_details** - Provides the details of a specific session
- **udmmig_submit_metricpics** - Provides a mapping between the UDM and the ME in order to create a new ME or use an existing one
- **udmmig_retry_deploys** - Deploys the ME to the targets where the UDM is present. Note that the ME has to be in a deployable draft or published state for this command to succeed
- **udmmig_request_udmdelete** - Deletes the UDM and completing the migration process

Usage Examples

The following exercise outlines a simple use case to showcase the migration

Consider a system with one host (host1) that has one host UDM (hostudm1) on it. The goal is to create a new ME (me1) that represents the UDM. The sequence of commands would be as follows

```
$ emcli list_unconverted_udms

-------------+----------------------+-----------+--------------------
Type         | Name                 | Metric    | UDM
-------------+----------------------+-----------+--------------------
host         | host1                |UDM        | hostudm1
```

The command indicates that there is only one UDM that has not been migrated or in the process of migration at this stage. Now proceed with the creation of a session.

```
$ emcli create_udmmig_session -name=migration1 -desc="Convert UDMs for host
target" -udm_choice=hostudm1 -target=host:host1
```

```
Migration session created - session id is 1
```

The command creates a migration session with name migration1 and the description "convert UDMs for host target". The udm_choice flag indicates the UDM chosen and the target flag describes the target type and the target on which the UDM resides. Migration sessions are identified by session IDs. The current session has an ID of 1.

```
$ emcli udmmig_summary

------+--------------+------------------+------+------+--------+------+--------
ID    | Name         | Description      |#Tgts |Todo  |#Tmpls  |Todo  |IncRules
------+--------------+------------------+------+------+--------+------+--------
1     |migration1    |Convert UDMS      |      | 1/1  |  0     | -/0  | -/0
------+--------------+------------------+------+------+--------+------+--------
```

The command summarizes all the migrations sessions currently in progress. The name and description fields identify the session. The remaining columns outline the number of targets, templates and incident rules that contain references to the UDM that is being converted to a metric extension. The 'Todo' columns indicate the number of targets, templates and incident rules whose references to the UDM are yet to be updated. Since a migration session can be completed over a protracted period of time, the command provides an overview of the portion of the session that was been completed.

```
$ emcli list_unconverted_udms

There are no unconverted udms
```

Since the UDM is part of a migration session, it no longer shows up in the list of unconverted UDMs.

```
$ emcli udmmig_session_details -session_id=1

Name: migration1
Desc: Convert UDMs for host target
Created: <date> <time>
UDM Pick: [hostudm1]
UDMs being converted:
----------+----------+---------+------+------------+---------+---------+-----
Type      |Name      |UDM      |#MC   |Metric      |Column   |DepS     |DelS
----------+----------+---------+------+------------+---------+---------+-----
host      |host1     |hostudm1 |  0   |            |         |WAIT     |WAIT
----------+----------+---------+------+------------+---------+---------+-----
```

The command provides the status of a single migration session. It lists the name of the UDM and the target type and name of the target on which the UDM resides. In addition, it also outlines the metric extensions currently in the EM instance that match the UDM. The user can elect to use one of the existing choices or create an entirely new metric extension.

The system attempts to find compatible metric extensions by matching the properties of the UDM. For example, in the case of a host UDM, the system tries to find a metric extension that has the same command, script and argument fields. In the case of a database UDM, the system attempts to match the SQL query.

Finally, the DepS column indicates whether the metric extension that was matched to the UDM has been deployed to the target on which the UDM is defined. The DelS column tells the user whether the UDM has been deleted after the metric extension has been deployed. As the user proceeds with the migration, the above table is updated

from left to right. When the delete status column is set to complete, the migration session has ended.

```
$ emcli udmmig_submit_metricpicks -session_id=1 -input_file=metric_picks:filename

Successfully submitted metric picks for migration session
```

The command instructs the Enterprise Manager instance to use an existing metric extension or create a new one to replace the UDM. The various options are presented through a file, which is filename in the above command. The contents of the file are shown below

"host,host1,hostudm1,N,ME$me1,Usage"

Each line in the file represents a mapping from n UDM to an ME. The line provides the target type, the name of the target, the name of the UDM, a flag to indicate whether the metric extension is new (N) or existing (E), the name of the metric extension (note that ME$ must be prefixed) and the column name.

The types of UDMs supported are:

- Host (host)
- Database (oracle_database)
- RAC (rac_database)

A user can only specify the names of the data columns via the collection item portion of the file. A metric extension created through migration will always have two columns to represent the structure of the UDM. The first column is an index column for single column UDMs while the second column uses the column name mentioned in the file. In the case of two column UDMs, the first column of the ME is termed as the 'KEY' column and the collection name is used for the second column.

At this stage, the metric extension has been created and is visible in the metric extensions library.

```
$ emcli udmmig_session_details -session_id=1

Name: migration1
Desc: Convert UDMs for host target
Created: <date> <time>
UDM Pick: [hostudm1]
Udms being converted:
----------+--------+---------+------+----------+------------+---------+-----
Type      |Name    |UDM      |#MC   |Metric    |Column      |DepS     |DelS
----------+--------+---------+------+----------+------------+---------+-----
host      |host1   |hostudm1 |   1  | ME$me1   | Usage      |WAIT     |WAIT
----------+--------+---------+------+----------+------------+---------+-----

        #MC : There are 1 matches for udms in this session.
        Use emcli udmmig_list_matches to list available matches
```

The session details command indicates that there is one matching metric extension for this UDM (the value of the MC column is 1) and that metric extension is named as ME$me1. At this stage, we are ready to test the metric extension through the library page. Once the testing is complete and the user is satisfied with the metric extension that has been created, it is ready to be deployed. In order to deploy, the metric extension has to be minimally saved as a deployable draft.

```
$ emcli udmmig_retry_deploys -session_id=1 -input_file=metric_tasks:filename2

Metric Deployments successfully submitted
```

Note that the system will trigger a job to automatically deploy the metric extension to all targets where the UDM was present once the metric extension is published. If the user is interested in manually controlling the operation, the above command will perform the necessary steps. The command is similar to the submit_metricpicks option in that a file with the UDM to target mapping is provided. It is referred to by filename2 above. The contents of the file are as follows

"host,host1,hostudm1"

Each line in the file is a mapping from the UDM to the targets type and target on which it resides. Once the command is executed, jobs to deploy the metric extensions to various targets have been launched and can be tracked through the user interface.

```
$ emcli udmmig_request_udmdelete -session_id=1 -input_file=metric_tasks:demo_tasks

Udm deletes successfully submitted
```

The final command deletes the UDMs that were migrated to metric extensions. Note that this command might partially finish based on how many of the deployments were completed when the command was run.

```
$ emcli udmmig_session_details -session_id=1

Name: migration1
Desc: Convert UDMs for host target
Created: <date > <time>
Completed: <date > <time>
UDM Pick: [hostudm1]
Udms being converted:
--------+----------+---------+------+------------+------------+---------+-----
Type    |Name      |UDM      |#MC   |Metric      |Column      |DepS     |DelS
--------+----------+---------+------+------  ----+------------+---------+-----
host    |host1     |hostudm1 |  1   | ME$me1     | Usage      |COMP   |COMP
--------+----------+---------+------+------------+------------+---------+-----

        #MC : There are 1 matches for udms in this session.
          Use emcli udmmig_list_matches to list available matches
```

The session details command shows that the migration process is indeed complete.

7.6 Metric Extension Command Line Verbs

Metric extensions can be manipulated outside the UI via the Enterprise Manager Command Line Interface (EM CLI). Two categories of verbs are available:

- Metric Extension Verbs
  - *export_metric_extension*: Export a metric extension to an archive file
  - *get_unused_metric_extensions*: Get a list of unused metric extensions.
  - *import_metric_extension*: Import a metric extension archive file.
  - *publish_metric_extension*: Publish a metric extension for use by all administrators.
  - *save_metric_extension_draft*: Save a deployable draft of a metric extension.
- User-defined Metric Migration Verbs
  - *abort_udmmig_session*: Abort (partially) user-defined metric migration session.
  - *analyze_unconverted_udms*: Analyze the unconverted user-defined metrics.

- *create_udmmig_session*: Create a user-defined metric migration session.
- *list_unconverted_udms*: List the user-defined metrics that are not yet in a migration session.
- *udmmig_list_matches*: List the matching metrics per user-defined metric in a specific user-defined metric migration session.
- *udmmig_request_udmdelete*: Request deletion of user-defined metrics from targets.
- *udmmig_retry_deploys*: Retry deployment of metric extensions to targets.
- *udmmig_session_details*: Retrieve the details of a specific user-defined metric migration session.
- *udmmig_submit_metricpicks*: Select the metrics to replace user-defined metrics in a session.
- *udmmig_summary*: Summarize the status of all user-defined metric migration sessions.
- *udmmig_update_incrules*: Update user-defined metric incident rules to include replacement metric references.

Metric Extension Verbs

```
emcli export_metric_extension
     -file_name=<name of the metric extension archive>
     -target_type=<target type of the metric extension>
     -name=<name of the metric extension
     -version=<version of the metric extension>

Description:
  Export a metric extension archive file.

Options:
  -file_name=<file name>
    The name of the metric extension archive file to export into.
  -target_type=<target type>
    Target type of the metric extension.
  -name=<name>
    Name of the metric extension.
  -version=<version>
    Version of the metric extension to be exported.


emcli get_unused_metric_extensions

Description:
   Get a  list of metric extensions that are deployed to agents but not attached
to any targets.


emcli import_metric_extension
      -file_name=<name of the metric extension archive>
      -rename_as=<name of the metric extension to import as>

Description:
  Import a metric extension archive file.

Options:
  -file_name=<file name>
```

```
      The name of the metric extension archive file to be imported.
    -rename_as=<metric extension name>
      Import the metric extension using the specified name, replacing the name
given in the archive.


  emcli publish_metric_extension
       -target_type=<target type of the metric extension>
       -name=<name of the metric extension
       -version=<version of the metric extension>

  Description:
    Publish a metric extension for use by all administrators.
  The metric extension must currently be a deployable draft.

  Options:
    -target_type=<target type>
      Target type of the metric extension.
    -name=<name>
      Name of the metric extension.
    -version=<version>
      Version of the metric extension to be published.


  emcli save_metric_extension_draft
       -target_type=<target type of the metric extension>
       -name=<name of the metric extension
       -version=<version of the metric extension>

  Description:
    Save a deployable draft of a metric extension. The metric
  extension must currently be in editable state. Once saved as
  draft, the metric extension will no longer be editable.

  Options:
    -target_type=<target type>
      Target type of the metric extension.
    -name=<name>
      Name of the metric extension.
    -version=<version>
      Version of the metric extension to be saved to draft.
```

User-Defined Metric Verbs

```
  emcli abort_udmmig_session
        -session_id=<sessionId>
        [-input_file=specific_tasks:<complete path to file>]

  Description:
    Abort the migration of user-defined metrics to MEs in a session

  Options:
    -session_id=<id of the session>
      Specify the id that was returned at time of session created,
      or from the output of udmmig_summary
    [-input_file=specific_tasks:<complete file path>]
      This optional parameter points at a file name that contains a
         target, user-defined metric,
      one per line in the following format:
```

```
      <targetType>,<targetName>,<collection name>
      Use targetType=Template to indicate a template
      Use * for collection name to abort all user-defined metrics for a target


  emcli analyze_unconverted_udms [-session_id=<sessionId>]

  Description:
    Analyze user-defined metrics and list unique user-defined metrics, any
possible matches, and
    templates that can apply these matching metric extensions
  Options:
    -session_id=<id of a session to be reanalyzed>
      Not specifying a session id causes the creation of a analysis
      session that contains all unconverted user-defined metrics. You can specify
      this session id in future invocations to get fresh analysis.


  emcli create_udmmig_session
      -name=<name of the session>
      -desc=<description of the session>
      [-udm_choice=<specific udm to convert>]*
      {-target=<type:name of the target to migrate> }*
      | {-input_file=targetList:<complete path to file>};      {-template=<name of
the template to update> }*
      | {-input_file=templateList:<complete path to file>}
      [-allUdms]

  Description:
    Creates a session to migrate user-defined metrics to metric extensions for
targets.

  Options:
    -name=<session name>
      The name of the migration session to be created.
    -desc=<session session description>
      A description of the migration session to be created.
    -udm_choice=<udm name>
      If the session should migrate specific user-defined metrics, specify them
      Otherwise, all user-defined metrics will be migrated
    -target=<type:name of target to migrate>
      The type:name of the target to be updated.
      Multiple values may be specified.
    -input_file=targetList:<complete file path>
      This takes a file name that contains a list of targets,
      one per line in the following format:
      <targetType>:<targetName>
    -template=<name of template to migrate>
      The name of the template to update.Multiple values may be specified
    -input_file=templateList:<complete file path>
      This takes a file name that contains a list of templates,
      one name per line
    -allUdms
      This forces the session to contain all user-defined metrics from targets and
      templates (default behavior just picks those not in a session)

  emcli list_unconverted_udms [-templates_only]

  Description:
    Get the list of all user-defined metrics that are not yet in a migration
```

```
session

  Options:
    -templates_only
      Only lists unconverted user-defined metrics in templates.


  emcli udmmig_list_matches
        -session_id=<sessionId>

  Description:
    Lists the matching metrics per user-defined metric in a migration session
  Options:
    -session_id=<id of the session>
      Specify the id that was returned at time of session created,
      or from the output of udmmig_summary


  emcli udmmig_request_udmdelete
        -session_id=<sessionId>
        -input_file=metric_tasks:<complete path to file>

  Description:
    Delete the user-defined metrics that have been replaced by Metric Extenions

  Options:
    -session_id=<id of the session>
      Specify the id that was returned at time of session created,
      or from the output of udmmig_summary
    -input_file=metric_tasks:<complete file path>
      This takes a file name that contains a target, user-defined metric,
      one per line in the following format:
      <targetType>,<targetName>,<collection name>


  emcli udmmig_retry_deploys
        -session_id=<sessionId>
        -input_file=metric_tasks:<complete path to file>

  Description:
    Retry the deployment of metric extensions to a target

  Options:
    -session_id=<id of the session>
      Specify the id that was returned at time of session created,
      or from the output of udmmig_summary
    -input_file=metric_tasks:<complete file path>
      This takes a file name that contains a target, user-defined metric,
      one per line in the following format:
      <targetType>,<targetName>,<collection name>


  emcli udmmig_submit_metricpicks
        -session_id=<sessionId>
        -input_file=metric_picks:<complete path to file>

  Description:
    Supply the metric picks to use to replace user-defined metrics per target in a
session
```

```
  Options:
    -session_id=<id of the session>
      Specify the id that was returned at time of session created,
      or from the output of udmmig_summary
    -input_file=metric_picks:<complete file path>
      This takes a file name that contains a target, user-defined metric, metric
pick,
      one per line in the following format:
      <targetType>,<targetName>,<collection name>,[N/E],<metric>,<column>
       using N if a new metric should be created or E if an existing
       metric is referenced.


  emcli udmmig_summary
      [-showAll]

  Description:
    Gets the summary details of all migration sessions in progress

  Options:
    -showAll
      This prints out all sessions including those that are complete.
      By default, only in-progress sessions are listed.


  emcli udmmig_update_incrules
        -session_id=<sessionId>
        -input_file=udm_inc_rules:<complete path to file>

  Description:
    Update Incident Rules that reference user-defined metrics with a reference to
    replacing metric extension.

  Options:
    -session_id=<id of the session>
      Specify the id that was returned at time of session created,
      or from the output of udmmig_summary
    -input_file=udm_inc_rules:<complete file path>
      This takes a file name that contains rule, user-defined metric, metric,
      one per line in the following format:
      <ruleset id>,<rule id>,<udm name>,<metric name>
```

8

Utilizing the Job System and Corrective Actions

Today's IT environments have many sets of components, so it is beneficial to minimize the time needed to support these IT components and eliminate the human error associated with component maintenance. The Enterprise Manager Cloud Control Job System can automate routine administrative tasks and synchronize components in your environment so you can manage them more efficiently.

This chapter facilitates your usage of the Job System by presenting instructional information in the following sections:

- Job System Purpose and Overview
- Preliminary Considerations
- Creating Jobs
- Analyzing Job Activity
- Generating Job Event Criteria
- Creating Event Rules For Job Status Change
- Creating Corrective Actions

8.1 Job System Purpose and Overview

The Enterprise Manager Job System serves these purposes:

- Automates many administrative tasks; for example: backup, cloning, and patching
- Enables you to create your own jobs using your own custom OS and SQL scripts
- Enables you to create your own multi-task jobs comprised of multiple tasks

A job is a unit of work that you define to automate commonly-run tasks. Scheduling flexibility is one of the advantages of jobs. You can schedule a job to start immediately or start at a later date and time. You can also run the job once or at a specific interval, such as three times every month.

The Job Activity page (Figure 8–1) is the hub of the Job System. From this page, you can:

- Search for existing job runs and job executions filtered by name, owner, status, scheduled start, job type, target type, and target name
- Create a job
- View or edit the job definition

- Create like, copy to library, suspend, resume, stop, and delete a job
- View results, edit, create like, suspend, resume, retry, stop, and delete a job run or execution

Figure 8–1 Job Activity Page

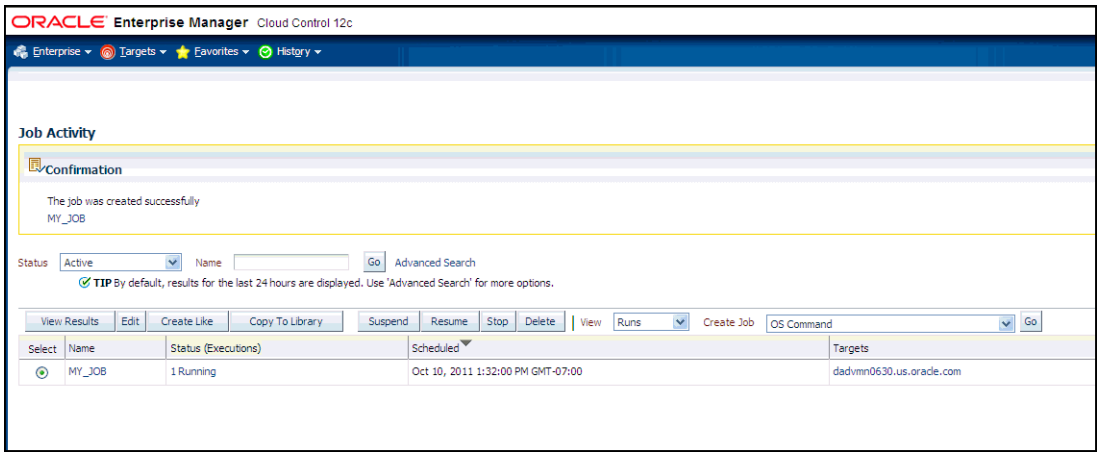

Besides accessing the Job Activity page from the Enterprise menu, you can also access this page from any Database or Cluster Database property page (Home, Performance, Availability, and so forth) by selecting Job Activity from the Oracle Database menu. When you access this page from these alternate locations, rather than showing the entire list of jobs, the Job Activity page shows a subset of the jobs associated with the particular target.

8.1.1 What Are Job Executions and Job Runs?

Job executions are usually associated with one target, such as a patch job on a particular database. When a job is run against multiple targets, each execution may execute on one target.

Job executions are not always a one-to-one mapping to a target. Some executions have multiple targets, such as comparing hosts. A few jobs have no target.

When you submit a job to many targets, it would be tedious to examine the status of each execution of the job against each target. For example, suppose you run a backup job against several databases. A typical question would be: Were all the backup jobs successful, and if not, which jobs failed? If this backup job runs every week, you would want to know which backups were successful and those that failed each week.

With the Job System, you can easily get these answers by viewing the *job run*. A job run is the summary of all job executions of a job that ran on a particular scheduled date. For example, if you have a job scheduled for March 5th, you will have a March 5 job run. The job table that shows the job run provides a roll-up of the status of the executions, such as Succeeded, Failed, or Error.

8.1.2 Operations on Job Executions and Job Runs

Besides supporting the standard job operations of create, edit, create like, and delete, the Job System enables you to:

- **Suspend jobs —**

You can suspend individual executions or entire jobs. For example, you may need to suspend a job if a needed resource was unavailable, or the job needs to be postponed.

If a job is scheduled to repeat but is suspended past the scheduled repeat time, the execution of this job would be marked "Skipped."

- **Resume jobs —**

  After you suspend a job, any scheduled executions do not occur until you decide to resume the job.

- **Retry failed executions —**

  When analyzing individual executions or entire jobs, it is useful to retry a failed execution after you determine the cause of the problem. This alleviates the need to create a new job for that failed execution. When you use the Retry operation in the Job System, Enterprise Manager provides links from the failed execution to the retried execution and vice versa, should it become useful to retroactively examine the causes of the failed executions. Only the most recent retry is shown in the Job Run page.

With regard to job runs, the Job System enables you to:

- **Delete old job runs**
- **Stop job runs**
- **Retry job runs**

> **See Also:** For more information on job executions and runs, refer to Enterprise Manager Cloud Control online help.

8.2 Preliminary Considerations

Before proceeding to the procedural information presented in Section 8.3, "Creating Jobs" on page 8-4, it is advisable to read the topics presented in the sections below:

- Creating Scripts
- Sharing Job Responsibilities
- Submitting Jobs for Groups

8.2.1 Creating Scripts

Besides predefined job tasks, you can define your own job tasks by writing code to be included in OS and SQL scripts. The advantages of using these scripts include:

- When defining these jobs, you can use target properties.
- When defining these jobs, you can use the job library, which enables you to share the job and make updates as issues arise. However, you need to resubmit modified library jobs for them to take effect.
- You can submit the jobs against multiple targets.
- You can submit the jobs against a group. The job automatically keeps up with changes to group membership.
- For host command jobs, you can submit to a cluster.
- For SQL jobs, you can submit to a Real Application Cluster.

8.2.2 Sharing Job Responsibilities

To allow you to share job responsibilities, the Job System provides job privileges. These job privileges allow you to share the job with other administrators. Using privileges, you can:

- Grant access to the administrators who need to see the results of the job.
- Grant Full access to the administrators who may need to edit the job definition or control the job execution (suspend, resume, stop).

You can grant these privileges on an as-needed basis.

8.2.3 Submitting Jobs for Groups

Besides submitting jobs to individual targets, you can submit jobs against a group of targets. Any job that you submit to a group is automatically extended to all of its member targets that match the target type of the job, and accounts for the membership of the group as it changes.

For example, if a Human Resources job is submitted to the Payroll group, then a new host is added to this group, the host automatically becomes part of future Human Resources job runs. For instance, for a daily repeating job scheduled for 10:00 a.m. today, if you add a target before that time, the new target would be part of the job run. However, if you add a target after that time today, the target would not be part of today's run, but would be part of the next run. Additionally, if the Payroll group is comprised of diverse targets (for example: databases, hosts, and application servers), the job only runs against applicable targets in the group.

By accessing the Group Home page, you can analyze the job activity for that group.

See Also: Chapter 5, "Managing with Groups"

8.3 Creating Jobs

Your first task in creating a job from the Job Activity page is to choose a job type, which the next section, Selecting a Job Type, explains. The most typical job types are OS command jobs, script jobs, and multi-task jobs, which are explained in these subsequent sections:

- Creating an OS Command Job
- Creating a SQL Script Job
- Creating a Multi-task Job

8.3.1 Selecting a Job Type

Using the Job System, you can create a job by selecting one of the job types from the Create Job drop-down in the Job Activity page. The most commonly used types are as follows:

- **OS Command** — Runs an operating system command or script.
- **SQL Script** — Runs a user-defined SQL or PL/SQL script.
- **Multi-Task** — Use to specify primary characteristics for multi-task jobs or corrective actions. Multi-task jobs enable you to create composite jobs by defining tasks, with each task functioning as an independent job. You edit and define tasks similarly to a regular job.

Blocked Agents Job Type

A blocked Agent is a condition where the Oracle Management Server (OMS) rejects all heartbeat or upload requests from the blocked Agent. Therefore, a blocked Agent cannot upload any alerts or metric data to the OMS. However, blocked Agents continue to collect monitoring data.

On a blocked Agent, the OMS “ignores” requests from the blocked Agent, thereby reducing the workload on the OMS. For example, by using this feature for an Agent that fails to upload properly, you can block the Agent until you can resolve the upload issue.

An Agent can become blocked under the following circumstances:

- The system detects that the Agent is no longer sending the correct state. This can occur after a failed recovery, or when users have corrupted state files. The OMS can detect some of the corruptions, and when it finds one, it blocks the Agent until the problem has been resolved.
- A superuser has blocked an Agent to prevent a "rogue" Agent from flooding the system with errors and bad data.

When an Agent is blocked for a long period of time and the Agent is kept running, it eventually must stop monitoring, because it will run out of local disc space to store all of the results. However, this is not an issue, because the "state" of the Agent was corrupt anyway. Therefore, unless corrective actions were taken, the Agent should remain blocked so that no data then penetrates the system.

8.3.2 Creating an OS Command Job

Use this type of job to run an operating system command or script. Tasks and their dependent steps for creating an OS command are discussed below.

Task 1 Initiate Job Creation

1. From the Enterprise menu, select **Jobs**, then Job Activity.
2. Select **OS Command** from the Create Job drop-down, then click **Go**. The **General property page** of the Create OS Command Job page appears.

Task 2 Specify General Job Information

Perform these steps on the General property page:

1. Provide a required Name for the job, then select a Target Type from the drop-down.

   After you have selected a target of a particular type for the job, only targets of that same type can be added to the job. If you change target types, the targets you have populated in the Targets table disappear, as well as parameters and credentials for the job.

   If you specify a composite as the target for this job, the job executes only against targets in the composite that are of the selected target type. For example, if you specify a target type of host and a group as the target, the job only executes against the hosts in the group, even if there are other non-host targets in the group.
2. Click **Add**, then select one or more targets from the Search and Select: Targets pop-up window. The targets now appear in the Targets table.
3. Click the **Parameters** property page link.

Task 3 Specify Parameters

Perform these steps on the Parameters property page:

1. Select either **Single Operation** or **Script** from the Command Type drop-down.

   The command or script you specify executes against each target specified in the target list for the job. The Management Agent executes it for each of these targets.

   Depending on your objectives, you can choose one of the following options:

   - Single Operation to run a specific command
   - Script to run an OS script and optionally provide an interpreter, which processes the script; for example, `%perlbin%/perl or /bin/sh`.

   Sometimes, a single command line is insufficient to specify the commands to run, and you may not want to install and update a script on all hosts. In this case, you can use the Script option to specify the script text as part of the job.

2. Based on your objectives, follow the instructions in Section 8.3.2.1, "Specifying a Single Operation" or Section 8.3.2.2, "Specifying a Script".

3. Click the **Credentials** property page link.

Task 4 Specify Credentials - (optional)

You do not need to provide input on this page if you want to use the system default of using preferred credentials.

On the Credentials property page, you can specify the credentials that you want the Oracle Management Service to use when it runs the OS Command job against target hosts. The job can use either the job submitter's preferred credentials for hosts, or you can specify other credentials to override the preferred credentials.

You do not need to provide input on this page if you have already set preferred credentials.

- **To use preferred credentials:**

  1. Select the **Preferred Credential** radio button, which is the default selection.

     If the target for the OS Command job is a host or host group, the preferred host credentials are used. You specify these on the Database Preferred Credentials page, and they are different from the host credentials for the host on which the database resides.

  2. Select either **Normal Host Credentials** or **Privileged Host Credentials** from the Host Credentials drop-down.

     You specify these separately on the Preferred Credentials page, which you can access by selecting **Security** from the **Setup** menu, then **Preferred Credentials**. The Preferred Credentials page appears, where you can click the Manage Preferred Credentials button to set credentials.

- **To use named credentials:**

  1. Select the **Named Credential** radio button to override database or host preferred credentials.

     The drop-down list is a prepopulated credential set with values saved with names. These are not linked to targets, and you can use them to provide credential and authentication information to tasks.

- **To use other credentials:**

1. Select the **New Credential** radio button to override previously defined preferred credentials.

   Note that override credentials apply to all targets.

2. Optionally select Sudo or PowerBroker as the run privilege.

   Sudo enables you to authorize certain users (or groups of users) to run some (or all) commands as root while logging all commands and arguments. PowerBroker provides access control, manageability, and auditing of all types of privileged accounts.

   If you provide Sudo or PowerBroker details, they must be applicable to all targets. It is assumed that Sudo or PowerBroker settings are already applied on all the hosts on which this job is to run.

   See your Super Administrator about setting up these features if they are not currently enabled.

   Tip: For information on using Sudo, see the Sudo Manual at:

   `http://www.sudo.ws/sudo/man/1.7.4p6/sudo.man.html`

   For information on using PowerBroker, see the PowerBroker Desktops User Guide at:

   `http://www.ubm-global.com/docs/powerbroker/PBWD_User_Guide_V5%200.pdf`

Task 5 Schedule the Job - (optional)

You do not need to provide input on this page if you want to proceed with the system default of running the job immediately after you submit it.

1. Select the type of schedule:

   - **One Time (Immediately)**

     If you do not set a schedule before submitting a job, Enterprise Manager executes the job immediately with an indefinite grace period. You may want to run the job immediately, but specify a definite grace period in case the job is unable to start for various reasons, such as a blackout, for instance.

     A grace period is a period of time that defines the maximum permissible delay when attempting to start a scheduled job. If the job system cannot start the execution within a time period equal to the scheduled time plus grace period, it sets the job status to Skipped.

   - **One Time (Later)**

     – Setting up a custom schedule:

       You can set up a custom schedule to execute the job at a designated time in the future. When you set the Time Zone for your schedule, the job runs simultaneously on all targets when this time zone reaches the start time you specify. If you select each target's time zone, the job runs at the scheduled time using the time zone of the managed targets. The time zone you select is used consistently when displaying date and time information about the job, such as on the Job Activity page, Job Run page, and Job Execution page.

       For example, if you have targets in the Western United States (US Pacific Time) and Eastern United States (US Eastern Time), and you specify a schedule where Time Zone = US Pacific Time and Start Time = 5:00 p.m., the job runs simultaneously at 5:00 p.m. against the targets in the Western

United States and at 8:00 p.m. against the targets in the Eastern United States. If you specify 5:00 p.m. in the Agent time zone, the executions do not run concurrently. The EST target would run 3 hours earlier.

– Specifying the Grace Period:

The grace period controls the latest start time for the job in case the job is delayed. A job might not start for many reasons, but the most common reasons are that the Agent was down or there was a blackout. By default, jobs are scheduled with indefinite grace periods.

A job can start any time before the grace period expires. For example, a job scheduled for 1 p.m. with a grace period of 1 hour can start any time before 2 p.m., but if it has not started by 2 p.m., it is designated as skipped.

- **Repeating**
  - Defining the repeat interval:

    Specify the Frequency Type (time unit) and Repeat Every (repeat interval) parameters to define your job's repeat interval. The Repeat Until options are as follows:

    Note that both the end date and time determine the last execution. For example, for a job that runs daily at 6 p.m., where...

    Start Time is June 1, 2010 at 6 p.m.
    End Time is June 30, 2010 at 4 a.m.

    ... the last execution runs on June 29, not June 30, since the June 30 end time occurs before the daily time of the job.

  – Specifying the Grace Period:

    The grace period controls the latest start time for the job in case the job is delayed. A job might not start for many reasons, but the most common reasons are that the Agent was down or there was a blackout. By default, jobs are scheduled with indefinite grace periods.

    If the job starts on time, the grace period is ignored. For example, a job scheduled for 1 p.m. with a grace period of 1 hour can start any time before 2 p.m., but if it has not started by 2 p.m., it is designated as skipped.

2. Click the **Access** property page link.

Task 6 Specify Who Can Access the Job - (optional)

You do not need to provide input on this page if you want to proceed with the system default of not sharing the job. The table shows the access that administrators and roles have to the job. Only the job owner (or Super Administrator) can make changes on the Job Access page.

1. Change access levels for administrators and roles, or remove administrators and roles. Your ability to make changes depends on your function.

   If you are a job owner, you can:

   - Change the access of an administrator or role by choosing the Full or View access privilege in the Access Level column in the table.
   - Remove all access to the job for an administrator or role by clicking the icon in the Remove column for the administrator or role. All administrators with Super Administrator privileges have the View access privilege to a job. If you

choose to provide access privileges to a role, you can only provide the View access privilege to the role, not the Full access privilege.

If you are a Super Administrator, you can:

- Grant View access to other Enterprise Manager administrators or roles.
- Revoke all administrator access privileges.

Note: Neither the owner nor a super user can revoke View access from a super user. All super users have View access.

For more information on access levels, see Section 8.3.2.3, "Access Level Rules".

2. Click **Add** to add administrators and roles. The Create Job Add Administrators and Roles page appears.

   a. Specify a **Name** and **Type** in the Search section and click **Go**. If you just click Go without specifying a Name or Type, all administrators and roles in the Management Repository appear in the table.

      The value you specify in the Name field is not case-sensitive. You can specify either * or % as a wildcard character at any location in a string (the wildcard character is implicitly added to the end of any string). For example, if you specify `%na` in the Name field, names such as ANA, ANA2, and CHRISTINA may be returned as search results in the Results section.

   b. Select one or more administrators or roles in the Results section, then click **Select** to grant them access to the job. Enterprise Manager returns to the Create Job Access page or the Edit Job Access page, where you can modify the access of administrators and roles.

3. Define a notification rule.

   You can use the Notification system (rule creation) to easily associate specific jobs with a notification rule. The Cloud Control Notification system enables you to define a notification rule that sends e-mail to the job owner when a job enters one of these chosen states:

   - Scheduled
   - Running
   - Suspended
   - Succeeded
   - Problems
   - Action Required

   Note: Before you can specify notifications, you need to set up your email account and notification preferences. See Chapter 4, "Notifications" for this information.

Task 7 Conclude Job Creation

At this point, you can either submit the job for execution or save it to the job library.

- **Submitting the job —**

Click **Submit** to send the active job to the job system for execution, and then view the job's execution status on the main Job Activity page. If you are creating a library job, Submit saves the job to the library and returns you to the main Job Library page where you can edit or create other library jobs.

If you submit a job that has problems, such as missing parameters or credentials, an error appears and you will need to correct these issues before submitting an active job. For library jobs, incomplete specifications are allowed, so no error occurs.

Note: If you click Submit without changing the access, only Super Administrators can view your job.

- **Saving the job to the library** —

  Click **Save to Library** to the job to the Job Library as a repository for frequently used jobs. Other administrators can then share and reuse your library job if you provide them with access privileges. Analogous to active jobs, you can grant View or Full access to specific administrators. Additionally, you can use the job library to store:

  – Basic definitions of jobs, then add targets and other custom settings before submitting the job.

  – Jobs for your own reuse or to share with others. You can share jobs using views or giving Full access to the jobs.

  – Critical jobs for resubmitting later, or revised versions of these jobs as issues arise.

8.3.2.1 Specifying a Single Operation

Note: The following information applies to step 2 in Task 3, "Specify Parameters" on page 8-6.

Enter the full command in the **Command** field. For example:

```
/bin/df -k /private
```

Note the following points about specifying a single operation:

- You can use shell commands as part of your command. The default shell for the platform is used, which is `/bin/sh` for Linux and `cmd/c` for Windows.

  ```
  ls -la /tmp > /tmp/foobar.out
  ```

- If you need to execute two consecutive shell commands, you must invoke the shell in the Command field and the commands themselves in the OS Script field. You would specify this as follows in the Command field:

  ```
  sleep 3; ls
  ```

8.3.2.2 Specifying a Script

Note: The following information applies to step 2 in Task 3, "Specify Parameters" on page 8-6.

The value you specify in the OS Script field is used as stdin for the command interpreter, which defaults to `/bin/sh` on Linux and `cmd/c` on Windows. You can override this with another interpreter; for example: `%perlbin%/perl`. The shell scripts size is limited to 2 GB.

To control the maximum output size, set the `mgmt_job_output_size_limit` parameter in `MGMT_PARAMETERS` to the required limit. Values less than 10 KB and greater than 2 GB are ignored. The default output size is 10 MB.

You can run a script in several ways:

- **OS Scripts** — Specify the path name to the script in the OS Script field. For example:

  OS Script field: `/path/to/mycommand`
  Interpreter field:

- **List of OS Commands** — You do not need to enter anything in the Interpreter field for the following example of standard shell commands for Linux or Unix systems. The OS's default shell of `/bin/sh` or `cmd/c` will be used.

  ```
  /usr/local/bin/myProg arg1 arg2
  mkdir /home/$USER/mydir
  cp /dir/to/cp/from/file.txt /home/$USER/mydir
  /usr/local/bin/myProg2 /home/$USER/mydir/file.txt
  ```

  When submitting shell-based jobs, be aware of the syntax you use and the targets you choose. This script does not succeed on NT hosts, for example.

- **Scripts Requiring an Interpreter** — Although the OS shell is invoked by default, you can bypass the shell by specifying an alternate interpreter. For example, you can run a Perl script by specifying the Perl script in the OS Script field and the location of the Perl executable in the Interpreter field:

  OS Script field: `<Enter-Perl-script-commands-here>`
  Interpreter field: `%perlbin%/perl`

  The following example shows how to run a list of commands that rely on a certain shell syntax:

  ```
  setenv VAR1 value1
  setenv VAR2 value2
  /user/local/bin/myProg $VAR1 $VAR2
  ```

  You would need to specify csh as the interpreter. Depending on your system configuration, you may need to specify the following string in the Interpreter field:

  ```
  /bin/csh
  ```

  When submitting shell-based jobs, be aware of the syntax you use and the targets you choose. This script would not succeed on NT hosts, for example. However, you do have the option of running a script for a list of Windows shell commands, as shown in the following example. The default shell of `cmd/c` is used for Windows systems.

  ```
  C:\programs\MyApp arg1 arg2
  md C:\MyDir
  copy C:\dir1x\copy\from\file.txt \home\$USER\mydir
  ```

8.3.2.3 Access Level Rules

> **Note:** The following rules apply to Task 6, "Specify Who Can Access the Job - (optional)" on page 8-8.

- Super Administrators always have View access on any job.
- The Enterprise Manager administrator who owns the job can make any access changes to the job, except revoking View from Super Administrators.
- Super Administrators with a View or Full access level on a job can grant View (but not Full) to any new user. Super Administrators can also revoke Full and View from normal users, and Full from Super Administrators.
- Normal Enterprise Manager administrators with Full access levels cannot make any access changes on the job.
- If the job owner performs a Create Like operation on a job, all access privileges for the new job are identical to the original job. If the job owner grants other administrators View or Full job access to other administrators, and any of these administrators perform a Create Like operation on the job, ALL administrators will, by default, have View access on the newly created job.

8.3.3 Creating a SQL Script Job

The basic process for creating a SQL script job is the same as described in Section 8.3.2, "Creating an OS Command Job." The following sections provide supplemental information specific to script jobs:

- Specifying Targets
- Specifying Options for the Parameters Page
- Specifying Host and Database Credentials
- Returning Error Codes from SQL Script Jobs

8.3.3.1 Specifying Targets

You can run a SQL Script job against database and cluster database target types. You select the targets to run the job against by doing the following:

1. Click **Add** in the Targets section.
2. Select the database target(s) from the pop-up.

Your selection(s) now appears in the Target table.

> **Note:** For a cluster host or RAC database, a job runs only once for the target, regardless of the number of database instances. Consequently, a job cannot run on all nodes of a RAC cluster.

8.3.3.2 Specifying Options for the Parameters Page

In a SQL Script job, you can specify any of the following in the SQL Script field of the Parameters property page:

- Any directives supported by SQL*Plus
- Contents of the SQL script itself

- Fully-qualified SQL script file; for example:

```
@/private/oracle/scripts/myscript.sql
```

  Make sure that the script file is installed in the appropriate location on all targets.

- PL/SQL script using syntax supported by SQL*Plus; for example, one of the following:

```
EXEC plsql_block;
```

  or

```
DECLARE
   local_date DATE;
BEGIN
   SELECT SYSDATE INTO local_date FROM dual;
END;
/
```

  You can use target properties in the SQL Script field, a list of which appears in the Target Properties table. Target properties are case-sensitive. You can enter optional parameters to SQL*Plus in the Parameters field.

8.3.3.3 Specifying Host and Database Credentials

In the Credentials property page, you specify the host credentials and database credentials. The Management Agent uses the host credentials to launch the SQL*Plus executable, and uses database credentials to connect to the target database and run the SQL script. The job can use either the preferred credentials for hosts and databases, or you can specify other credentials that override the preferred credentials.

- **Use Preferred Credentials —**

  Select this choice if you want to use the preferred credentials for the targets for your SQL Script job. The credentials used for both host and database are those you specify in the drop-down. If you choose Normal Database Credentials, your normal database preferred credentials are used. If you choose SYSDBA Database Credentials, the SYSDBA preferred credentials are used. For both cases, the host credentials associated with the database target are used. Each time the job executes, it picks up the current values of your preferred credentials.

- **Named Credentials —**

  Select this choice if you want to override the preferred credentials for all targets, then enter the named credentials you want the job to use on all targets.

  Many IT organizations require that passwords be changed on regular intervals. You can change the password of any preferred credentials using this option. Jobs and corrective actions that use preferred credentials automatically pick up these new changes, because during execution, Enterprise Manager uses the current value of the credentials (both user name and password). Named credentials are also centrally managed. A change to a named credential is propagated to all jobs or corrective actions that use it.

  For corrective actions, if you specify preferred credentials, Enterprise Manager uses the preferred credentials of the last Enterprise Manager user who edited the corrective action. For this reason, if a user attempts to edit the corrective action that a first user initially specified, Enterprise Manager requires this second user to specify the credentials to be used for that corrective action.

8.3.3.4 Returning Error Codes from SQL Script Jobs

The SQL Script job internally uses SQL*Plus to run a user's SQL or PL/SQL script. If SQL*Plus returns 0, the job returns a status of Succeeded. If it returns any other value, it returns a job status of Failed. By default, if a SQL script runs and encounters an error, it may still result in a job status of Succeeded, because SQL*Plus still returned a value of 0. To make such jobs return a Failed status, you can use SQL*Plus EXIT to return a non-zero value.

The following examples show how you can return values from your PL/SQL or SQL scripts. These, in turn, will be used as the return value of SQL*Plus, thereby providing a way to return the appropriate job status (Succeeded or Failed). Refer to the *SQL*Plus User's Guide and Reference* for more information about returning EXIT codes.

Example 1

```
WHENEVER SQLERROR EXIT SQL.SQLCODE
select column_does_not_exist from dual;
```

Example 2

```
-- SQL*Plus will NOT return an error for the next SELECT statement
SELECT COLUMN_DOES_NOT_EXIST FROM DUAL;

WHENEVER SQLERROR EXIT SQL.SQLCODE;
BEGIN
  -- SQL*Plus will return an error at this point
  SELECT COLUMN_DOES_NOT_EXIST FROM DUAL;
END;
/
WHENEVER SQLERROR CONTINUE;
```

Example 3

```
variable exit_code number;

BEGIN
 DECLARE
 local_empno number(5);
 BEGIN
  -- do some work which will raise exception: no_data_found
  SELECT 123 INTO local_empno FROM sys.dual WHERE 1=2;
 EXCEPTION
  WHEN no_data_found THEN
    :exit_code := 10;
  WHEN others THEN
    :exit_code := 2;
  END;
 END;
/
exit :exit_code;
```

8.3.4 Creating a Multi-task Job

The basic process for creating a multi-task job is the same as described in Section 8.3.2, "Creating an OS Command Job." The following sections provide supplemental information specific to multi-task jobs:

- Job Capabilities
- Specifying Targets for a Multi-task Job

- Adding Tasks to the Job

8.3.4.1 Job Capabilities

Multi-task jobs enable you to create complex jobs consisting of one or more distinct tasks. Because multi-task jobs can run against targets of the same or different type, they can perform ad hoc operations on one or more targets of the same or different type.

The Job System's multi-task functionality makes it easy to create extremely complex operations. You can create multi-task jobs in which all tasks run on a single target. You can also create a multi-task job consisting of several tasks, each of which has a different job type, and with each task operating on separate (and different) target types. For example:

- Task 1 (OS Command job type) performs an operation on Host 1.
- If Task 1 is successful, run Task2 (SQL Script job type) against Database 1 and Database 2.

8.3.4.2 Specifying Targets for a Multi-task Job

You can run a multi-task job against any targets for which jobs are defined that can be used as tasks. Not all job types can be used as tasks.

The Target drop-down in the General page enables you to choose between running the job against the same targets for all tasks, or different targets for different tasks. Because each task of a multi-task job can be considered a complete job, when choosing the **Same targets for all tasks** option, you add all targets against which the job is to run from the General page. If you choose the **Different targets for different tasks** option, you specify the targets (and required credentials) the tasks will run against as you define each task.

After making your choice from the Target drop-down, you then select the targets to run the job against by clicking Add in the Targets section.

8.3.4.3 Adding Tasks to the Job

You can use the Tasks page to:

- Add, delete, or edit tasks of various job types
- Set task condition and dependency logic
- Add task error handling

You must define at least two tasks in order to set Condition and Depends On options. Task conditions define states in which the task will be executed. Condition options include:

- **Always** — Task is executed each time the job is run.
- **On Success** — Task execution **Depends On** the successful execution of another task.
- **On Failure** — Task execution **Depends On** the execution failure of another task.

The Error Handler Task is often a "clean-up" step that can undo the partial state of the job. The Error Handler Task executes if any task of the multi-task job has an error. Errors are a more severe form of failure, usually meaning that the job system could not run the task. Failures normally indicate that the task ran, but failed. The Error Handler Task does not affect the job execution status. Use the Select Task Type page to specify the job type of the task to be used for error handling.

8.4 Analyzing Job Activity

After you submit jobs, the status of all job executions across all targets is automatically rolled up and available for review on the Enterprise Summary page. Figure 8–2 shows the Jobs section at the bottom of the Enterprise Summary page.

Figure 8–2 Summary of Target Jobs on the Enterprise Summary Page

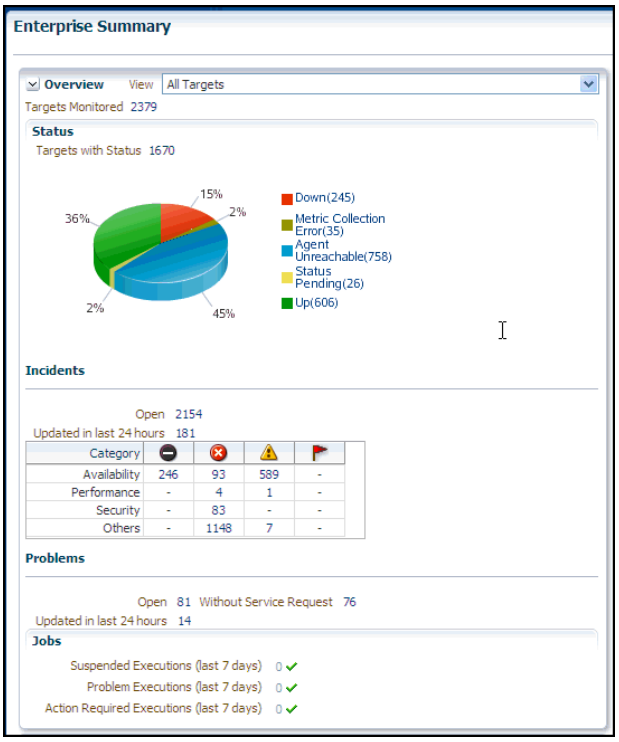

This information is particularly important when you are examining jobs that execute against hundreds or thousands of systems. You can determine the job executions that have failed. By clicking the number associated with a particular execution, you can drill down to study the details of the failed jobs.

8.5 Generating Job Event Criteria

The job system publishes status change events when a job changes its execution status, and these events have different severities based on the execution status.

Use the Job Event Generation Criteria page (Figure 8–3) to set up targets for job event notifications. This page enables you to decide about the jobs or targets or statuses for which you want to raise events or notifications. This ensures that users raise only useful events. Any settings you make on this page do not change the job behavior whatsoever. You can set up notifications on job events through incident rule sets.

Figure 8–3 Job Event Generation Criteria Page

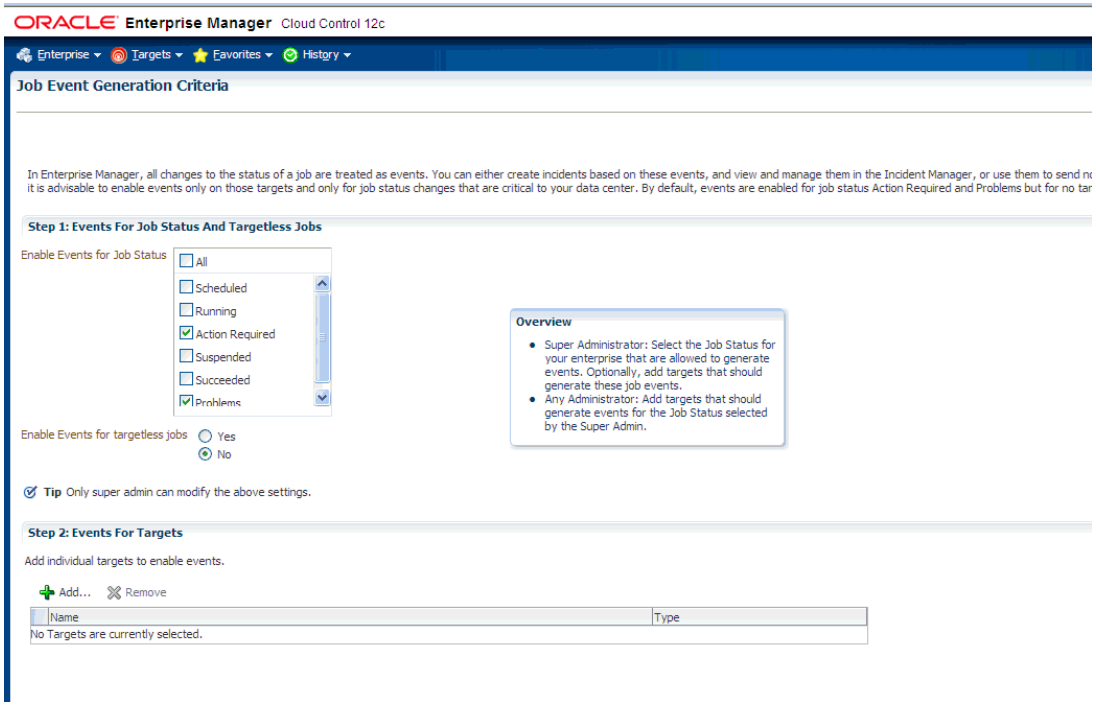

You need to add targets in the target filter if you want to set automatic job event generation for these targets. The job event generation settings you make on this page apply to all users. If you do not add any targets in the target filter, no targets will be set up for automatic job event generation.

On this page, you can do the following:

- Enable events for job status and targetless jobs
- Add targets to generate events for job status

8.5.1 Enabling Events For Job Status and Targetless Jobs

To enable events for job status and targetless jobs, do the following:

1. Ensure that you have Super Administrator privileges to select the job status for which you want to generate events.
2. Ensure that you are an administrator with View Target privileges to add targets for which you want to generate events for the job status set by the Super Administrator.
3. Log into Cloud Control as a Super Administrator.
4. From the **Setup** menu, select **Incidents** and then select **Job Events**. The Job Event Generation Criteria Page is displayed.
5. In the Job Event Generation Criteria page, do the following:
   a. In the Events For Job Status And Targetless Jobs section, from the **Enable Events for Job Status** check boxes, select the status for which you want to publish events. In **Enable Events for targetless jobs**, select **Yes** to create events for jobs that are not associated with any target.

b. In the Events For Targets section, click **Add** to add targets for which you want the job events to be enabled.

6. Click **Apply**.

8.5.2 Adding Targets To Generate Events For Job Status

After a Super Administrator selects events for which job status will be published, administrators can add targets to generate events. To add targets to generate events for job status, do the following:

1. Ensure that you are an administrator with View Target privileges to add targets for which you want to generate events for the job status set by a Super Administrator.
2. Log into Cloud Control as an administrator.
3. From the **Setup** menu, select **Incidents** and then select **Job Events**. The Job Event Generation Criteria Page is displayed.
4. In the Job Event Generation Criteria page, do the following:
   a. In the Events For Job Status And Targetless Jobs section, you can view the status for which events can be published. You can also see if events have been enabled for targetless job filters.
   b. In the Events For Targets section, click **Add** to add targets for which you want the job events to be enabled. You can also remove targets for which you do not want the job events to be enabled by clicking **Remove**.

   Note: Your selected settings in the Events for Targets section are global. Adding or removing targets for events also affect other Enterprise Manager users.

5. Click **Apply**.

8.6 Creating Event Rules For Job Status Change

Enterprise Manager enables you to create and apply rules to events, incidents, and problems. A rule is applied when a newly created or updated event, incident, or problem matches the conditions defined in the rule. The following sections explain how to create event rules for job status change events:

- Creating Job Status Change Event Rules For Jobs
- Creating Job Status Change Event Rules For Targets

8.6.1 Creating Job Status Change Event Rules For Jobs

To create job status change event rules for jobs, do the following:

1. Ensure that the relevant job status is enabled and required targets have been added to job event generation criteria.
2. Ensure that you have administrator privileges to create event rules for job status change events.
3. Log into Cloud Control as an administrator.
4. From the **Setup** menu, select **Incidents** and then **Incident Rules**. The Incident Rules Page is displayed.

5. In the Incident Rules page, click **Create Rule Set** to create rule sets for incidents.
6. Specify the **Name**, **Description**, and select **Enabled** to enable the rule set. Select Type as **Enterprise** if you want to set the rule for all Enterprise Manager users or Private if you want to set the rule for a specific user only. Select **Applies to Job**.

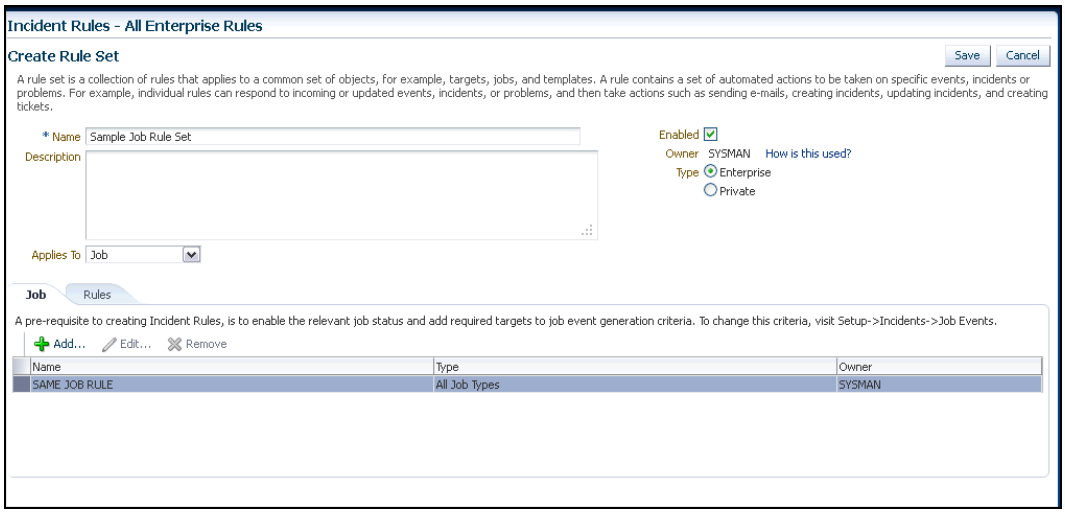

In the Job section, click **Add** to add jobs for which you want to create event rules.

7. In the Add Jobs dialog box, if you select Job **By Pattern**, provide **Job name like** and select the **Job Type**. Specify **Job owner like**. For **Specific jobs**, select the job. Click **OK**.
8. In the **Rules** tab, click **Create**.
9. In the Select Type of Rule to Create dialog box, select from the following choices according to the rule set you want to create:
   - **Incoming events and updates to events** to receive notification or create incidents for job rules. If you are operating on events (for example, if you want to create incidents for incoming events, such as job failed, or notify someone), choose this option.
   - **Newly created incidents or updates to incidents** receive notifications or create rules for incidents even though the events for which incidents are generated do not have associated rules. If you are operating on incidents already created or newly created (for example, you want to direct all incidents related to a group, say foo, to a particular user or escalate all incidents open for more than 3 days), choose this option.
   - **Newly created problems or updates to problems** to receive notifications or create rules for problems even though the incidents for which problems are generated do not have associated rules. This option does not apply for jobs.
10. Select **Incoming events and updates to events**, and in the Create New Rule: Select Events page, do the following:
    - **Select By Type** to **Job Status Change**. Select **All events of type Job Status Change** if you want to take an action for all job state change events for the selected jobs. Select **Specific events of type Job Status Change** if you only want to act on specific job states. If you have selected Specific events of type Job Status Change, select Job Status for events for which you want to create the rule.

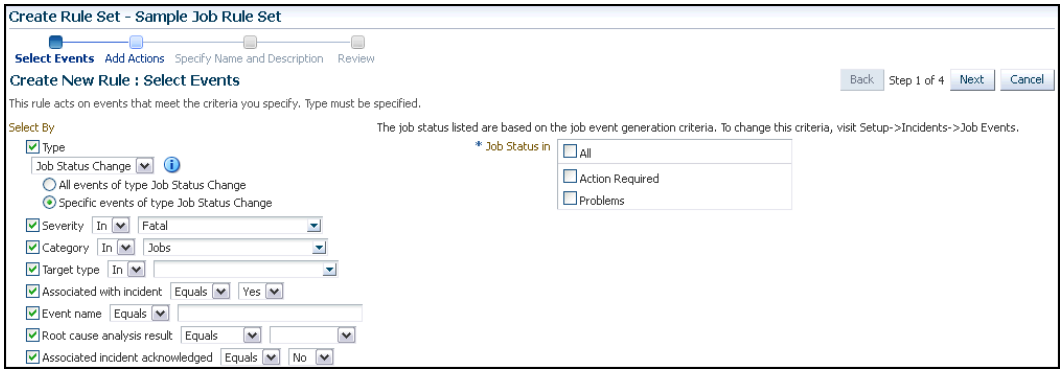

- Set the other criteria for which you want to set the rule as displayed in the above graphic.

11. Select **Newly created incidents or updates to incidents** if you want to create rules for an incident, though the event associated with the incident does not have notification rules. In the Create New Rule: Select Incidents page, select any of the following:

    - **All new incidents and updated incidents** to apply the rule to all new and updated incidents
    - **All new incidents** to apply the rule to all new incidents
    - **Specific incidents** and then select the criteria for the incidents

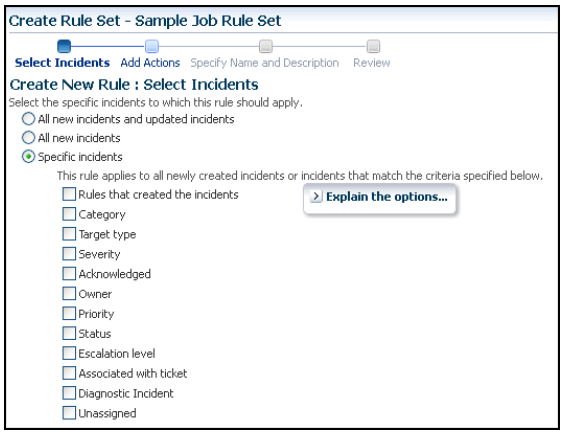

12. In the Create New Rule: Add Actions page, click **Add** to add actions to the rule.
13. In the Add Conditional Actions page, specify actions to be performed when the event matches the rule.

    In the Conditions for actions section, select:

    - **Always execute the actions** to execute actions regardless of event.
    - **Only execute the actions if specified conditions match** to execute actions to match specific criteria.

    When adding actions to events, specify the following:

    - Select **Create Incident** to create an incident for the event to manage and track its resolution.

- In the Notifications section, specify recipients for notifications in the **E-mail To**, **E-mail Cc**, and **Page** fields who will receive e-mail when the event for which a condition is set occurs. If Advanced and Repeat Notifications options have been set, specify them.
- In the Clear events section, select **Clear permanently** if you want to clear an event after the issue that generated the event is resolved.
- If you have configured event connections, in the Forward to Event Connectors section, you can send the events to third-party event management systems.

When adding actions to incidents, specify the following:

- In the Notifications section, specify recipients for notifications in the **E-mail To**, **E-mail Cc**, and **Page** fields who will receive e-mail when the event for which a condition is set occurs. If Advanced and Repeat Notifications options have been set, specify them.
- In the Update Incident section, specify the details to triage incidents when they occur. Specify **Assign to**, **Set priority to**, **Set status to**, and **Escalate to** details.
- In the Create Ticket section, if a ticket device has been configured, specify details to create the ticket.

Click **Continue**.

14. In the Specify Name and Description page, specify a **Name** and **Description** for the event rule. Click **Next**.
15. In the Review page, verify the details you have selected for the event rule and click **Continue** to add this rule in the rule set.
16. On the Create Rule Set page, click **Save** to save the rule set.

8.6.2 Creating Job Status Change Event Rules For Targets

To create job status change event rules for targets, do the following:

1. Ensure that the relevant job status is enabled and required targets have been added to job event generation criteria.
2. Ensure that you have administrator privileges to create event rules for job status change events.
3. Log into Cloud Control as an administrator.
4. From the **Setup** menu, select **Incidents**, then **Incident Rules**. The Incident Rules Page is displayed.
5. In the Incident Rules page, click **Create Rule Set** to create rule sets for incidents.
6. Specify the **Name**, **Description**, and select **Enabled** to enable the rule set. Select Type as **Enterprise** if you want to set the rule for all Enterprise Manager users, or Private if you want to set the rule for a only specific user. Select **Applies to Targets**.

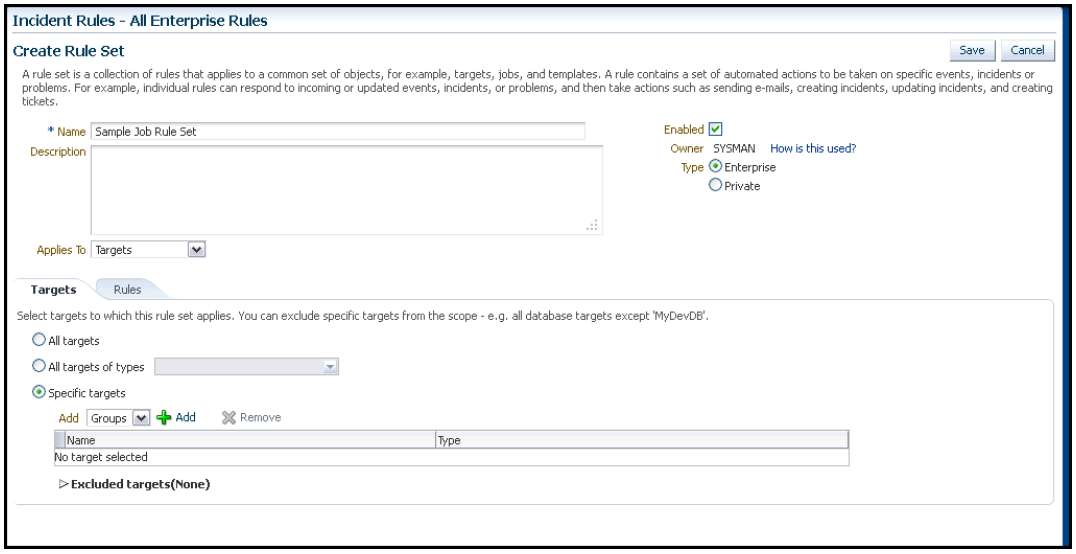

In the **Targets** tab, select one of the following:

- **All targets** to apply to all targets. In the Excluded Targets section, click **Add** to search and select the target that you want to exclude from the rule set. Click **Select**.
- **All targets of types** to select the types of targets to which you want to apply the rule set.
- **Specific targets** to individually specify the targets. Select to Add **Groups** or **Targets** to add groups or targets and click **Add** to search and select the targets to which you want to apply the rule set. Click **Select**. In the Excluded Targets section, click **Add** to search and select the target that you want to exclude from the rule set. Click **Select**.

7. In the **Rules** tab, click **Create**.
8. In the Select Type of Rule to Create dialog box, select from the following choices according to the rule set you want to create:
   - **Incoming events and updates to events** to receive notifications or create incidents for job rules. If you are operating on events (for example, if you want to create incidents for incoming events, such as job failed, or notify someone), choose this option.
   - **Newly created incidents or updates to incidents** receive notifications or create rules for incidents even though the events for which incidents are generated do not have associated rules. If you are operating on incidents already created or newly created (for example, you want to direct all incidents related to a group, say foo, to a particular user or escalate all incidents open for more than 3 days), choose this option.
   - **Newly created problems or updates to problems** to receive notifications or create rules for problems even though the incidents for which problems are generated do not have associated rules. This option does not apply for jobs.
9. Select **Incoming events and updates to events**, and in the Create New Rule: Select Events page, do the following:
   - **Select By Type** to **Job Status Change**. Select **All events of type Job Status Change** if you want to take an action for all job state change events for the

selected jobs. Select **Specific events of type Job Status Change** if you only want to act on specific job states. If you have selected Specific events of type Job Status Change, select Job Status for events for which you want to create the rule.

- Set the other criteria for which you want to set the rule as displayed in the above graphic.

10. Select **Newly created incidents or updates to incidents** if you want to create rules for an incident, though the event associated with the incident does not have notification rules. In the Create New Rule: Select Incidents page, select any of the following:
    - **All new incidents and updated incidents** to apply the rule to all new and updated incidents.
    - **All new incidents** to apply the rule to all new incidents.
    - **Specific incidents** and then select the criteria for the incidents.

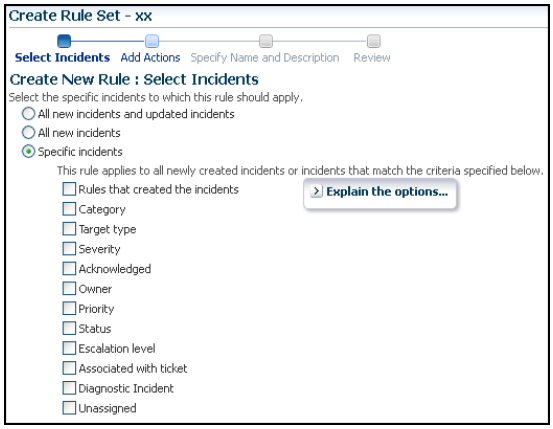

11. In the Create New Rule: Add Actions page, click **Add** to add actions to the rule.
12. In the Add Conditional Actions page, specify actions to be performed when the event matches the rule.

    In the Conditions for actions section, select:
    - **Always execute the actions** to execute actions regardless of event.
    - **Only execute the actions if specified conditions match** to execute actions to match specific criteria.

When adding actions to events, specify the following:

- Select **Create Incident** to create an incident for the event to manage and track its resolution.
- In the Notifications section, specify recipients for notifications in the **E-mail To**, **E-mail Cc**, and **Page** fields who will receive e-mail when the event for which a condition is set occurs. If Advanced and Repeat Notifications options have been set, specify them.
- In the Clear events section, select **Clear permanently** if you want to clear an event after the issue that generated the event is resolved.
- If you have configured event connections, in the Forward to Event Connectors section, you can send the events to third-party event management systems.

When adding actions to incidents, specify the following:

- In the Notifications section, specify recipients for notifications in the **E-mail To**, **E-mail Cc**, and **Page** fields who will receive e-mail when the event for which a condition is set occurs. If Advanced and Repeat Notifications options have been set, specify them.
- In the Update Incident section, specify the details to triage incidents when they occur. Specify **Assign to**, **Set priority to**, **Set status to**, and **Escalate to** details.
- In the Create Ticket section, if a ticket device has been configured, specify the details to create the ticket.

Click **Continue**.

13. In the Specify Name and Description page, specify a **Name** and **Description** for the event rule. Click **Next**.
14. In the Review page, verify the details you have selected for the event rule and click **Continue** to add this rule in the rule set.
15. On the Create Rule Set page, click **Save** to save the rule set.

8.7 Creating Corrective Actions

Corrective Actions enable you to specify automated responses to metric alerts. Corrective Actions ensure that routine responses to metric alerts are automatically executed, thereby saving you time and ensuring problems are dealt with before they noticeably impact end users.

Corrective actions share many features in common with the Job System. By default, a corrective action runs on the target on which the metric alert has triggered. Alternatively, you can specify a corrective action to contain multiple tasks, with each task running on a different target. You can also receive notifications for the success or failure of corrective actions.

You define corrective actions for individual metrics for monitored targets. The following sections provide instructions on setting up corrective actions:

- Providing Credentials
- Creating Corrective Actions for Metrics
- Creating a Library Corrective Action
- Specifying Access to Corrective Actions

- Setting Up Notifications for Corrective Actions
- Providing Agent-side Response Actions

8.7.1 Providing Credentials

Since corrective actions are associated with a target's metric thresholds, you can define corrective actions if you have been granted OPERATOR or greater privilege on the target. You can define separate corrective actions for both Warning and Critical thresholds. Corrective actions must run using the credentials of a specific user. For this reason, whenever a corrective action is created or modified, you must specify the credentials that the modified action runs with.

8.7.2 Creating Corrective Actions for Metrics

For any target, the Metric and Collection Settings page shows whether corrective actions have been set for various metrics. For each metric, the Corrective Actions column shows whether Critical and/or Warning severities of corrective actions have been set.

1. From any target's home page menu, select **Monitoring**, then **Metric and Collection Settings**. The Metric and Collection Settings page appears.

   Tip: For instance, on the home page for a host named `dadvmn0630.us.oracle.com`, you would select the Host menu, then Monitoring, then Metric and Collection Settings.

2. Click the pencil icon for a specific metric to access the Edit Advanced Settings page for the metric.
3. In the Corrective Actions section, click **Add** for the metric severity (Warning and/or Critical) for which you want to associate a corrective action.
4. Select the task type on the Add Corrective Actions page, then click **Continue**.
   - If you want to use a corrective action from the library, select **From Library** as the task type. Using a library corrective action copies the description, parameters, and credentials from the library corrective action. You must still define a name for the new corrective action. You can provide corrective action parameters if necessary.
   - If you want to create a corrective action to store in the library, see Section 8.7.3, "Creating a Library Corrective Action."
   - If you want to provide an Agent-side response action, select Agent Response Action as the task type. See Section 8.7.6, "Providing Agent-side Response Actions" for more information.
5. On the Corrective Action page, provide input for General, Parameters, and Credentials as you would similarly do when creating a job.
6. Click **Continue** to save the corrective action and return to the Edit Advanced Settings page, where your corrective action now appears.
7. *Optional*: To prevent multiple instances of a corrective action from operating simultaneously, enable the **Allow only one corrective action for this metric to run at any given time** checkbox.

   This option specifies that both Critical and Warning corrective actions will not run if a severity is reported to the Oracle Management Services when an execution of either corrective action is currently running. This can occur if a corrective action

runs longer than the collection interval of the metric it corrects; the value of the metric may be oscillating back and forth across one of the thresholds (leading to multiple executions of the same corrective action), or may be rising or falling quickly past both thresholds (in which case an execution of the Warning corrective action may overlap an execution of the Critical corrective action).

If you do not select this option, multiple corrective action executions are launched under the aforementioned circumstances. It is the administrator's responsibility to ensure that the simultaneous corrective action executions do not conflict.

8. Click **Continue** when you have finished adding corrective actions to return to the Metric and Collection Settings page.

   The page shows the corrective action value you have provided for the metric in the Corrective Actions column. Possible values are:

   - **None** — No corrective actions have been set for this metric.
   - **Warning** — A corrective action has been set for Warning, but not Critical, alerts for this metric.
   - **Critical** — A corrective action has been set for Critical, but not Warning, alerts for this metric.
   - **Warning and Critical** — Corrective actions have been set for both Warning and Critical alerts for this metric. If an Agent-side response action is associated with the metric, the value is also Warning and Critical, since Agent-side response actions are always triggered on either Critical or Warning alert severities.

9. Continue the process from step 2 forward, then click **OK** on the Metric and Collection Settings page to save your corrective actions and return to the target page you started from in step 1.

8.7.3 Creating a Library Corrective Action

For corrective actions that you use repeatedly, you can define a library corrective action. After a corrective action is in the library, you can reuse the corrective action definition whenever you define a corrective action for a target metric or policy rule.

1. From the Enterprise menu, select **Monitoring**, then **Corrective Actions**. The Corrective Action Library page appears.
2. Select a job type from the **Create Library Corrective Action** drop-down, then click **Go**.
3. Define the corrective action as you would for creating a job in Section 8.3, "Creating Jobs" for General, Parameters, and Credentials. For Access, go to the following optional step.
4. *Optional*: Select **Access** to define or modify the access you want other users to have for this corrective action.

   For more information, see Section 8.7.4, "Specifying Access to Corrective Actions."
5. Click **Save to Library** when you have finished. The Corrective Action Library page reappears, and your corrective action appears in the list.

   You can now create another corrective action based on this one (Create Like button), edit, or delete this corrective action.

You can access this library entry whenever you define a corrective action for a metric severity by selecting From Library as the task type in the Add Corrective Actions page.

See step 4 in Section 8.7.2, "Creating Corrective Actions for Metrics," for more information.

8.7.4 Specifying Access to Corrective Actions

As mentioned in the procedure above, you can determine the access to corrective actions by other users. You do not need to provide input for this page if you do not want to share the corrective action.

8.7.4.1 Defining or Modifying Access

The table on the Access page shows the access that administrators and roles have to the corrective action. Only the corrective action owner (or Super Administrator) can make changes on this page.

As the corrective action owner, you can do the following:

- Add other administrators and roles to the table by clicking **Add**, then selecting the appropriate type in the subsequent page that appears.
- Change the access of an administrator or role by choosing the **Full** or **View** access right in the Access Level column in the table.
- Remove all access to the corrective action for an administrator or role by clicking the icon in the **Remove** columns for this administrator or role. All administrators with Super Administrator privileges have the View access right to a corrective action.

If you choose to provide access rights to a role, you can only provide the View access right to the role, not the Full access right.

If you are a Super Administrator, you can:

- Grant View access to other Enterprise Manager administrators or roles.
- Revoke all administrator access privileges.

> **Note:** If a new user is being created, the user should have the CREATE_JOB privilege to create Corrective Actions.

8.7.4.2 Access Level Rules

Access level rules are as follows:

- Super Administrators always have View access for any corrective action.
- The Enterprise Manager administrator who owns the corrective action can make any access changes to the corrective action (except revoking View from Super Administrators).
- Super Administrators with a View or Full access level for a corrective action can grant View (but not Full) access to any new user. Super Administrators can also revoke Full and View access from normal users, and Full access from Super Administrators.
- Normal Enterprise Manager administrators with Full access levels cannot make any access changes on the corrective action.
- If the corrective action owner performs a Create Like operation on a corrective action, all access privileges for the new corrective action become identical to the original corrective action. If the corrective action owner grants other administrators View or Full access to other administrators, and any of these

administrators perform a Create Like operation on this corrective action, all administrators will, by default, have View access on the newly created corrective action.

8.7.5 Setting Up Notifications for Corrective Actions

Corrective actions are associated with metrics whose alerts trigger them. Any Enterprise Manager administrator with View or higher privileges on a target can receive notifications following the success or failure of a corrective action.

A single incident rule can contain any combination of alert and corrective action states. All metrics and targets selected by the incident rule are notified for the same alert and corrective action states. Therefore, if you want to be notified of corrective action success or failure for one metric, but only on failure for another, you need to use two incident rules. An incident rule can include corrective action states for metrics with which no corrective actions have been associated. In this case, no notifications are sent.

Note: Notifications cannot be sent for Agent-side response actions, regardless of the state of any incident rules applied to the target.

To create incident rules for notifications:

1. From the Setup menu, select **Incidents**, then **Incident Rules**.
2. Click **Create Rule Set**. The Create Rule Set wizard appears.
3. Provide the requisite information at the top of the Create Rule Set page, then select one of the target choices in the Targets sub-tab, supplying additional information as needed for the "All targets of types" and "Specific targets" choices.
4. Select the **Rules** sub-tab, then click **Create**.
5. In the pop-up that appears, select the default **Incoming events and update to events** choice, the click **Continue**.
6. On the Select Events page, enable the **Type** checkbox, then select **Metric Alert**.
7. Click the **Specific events of type Metric alert** radio button, then click **Add** in the table that appears.
8. In the pop-up that appears, select the Target Type, filter and select the metric, select a severity, then enable the desired corrective action status. Click **OK**.
9. From the Add Actions page, click **Add**.
10. Specify recipients in the Basic Notifications section of the Add Conditional Actions page.
11. Proceed through the final two pages of the wizard, then click **Continue**. Your new rule appears in the Create Rule Set page.
12. Click **Save** to save this rule.

After you have created one or more rule sets, you need to set up notification methods as follows:

1. From the Setup menu, select **Notifications**, then **Notification Methods**.
2. From the Notification Methods page, select **Help**, then **Enterprise Manager Help** for assistance on providing input for this page.

8.7.6 Providing Agent-side Response Actions

Agent-side response actions perform simple commands in response to an alert. When the metric triggers a warning or critical alert, the Management Agent automatically runs the specified command or script without requiring coordination with the Oracle Management Service (OMS). The Agent runs this command or script as the OS user who owns the Agent executable. Specific target properties can be used in the Agent response action script.

> **Note:** Use the Agent-side Response Action page to specify a single command-line action to be executed when a Warning or Critical severity is reached for a metric. For tasks that require alert context, contain more complex logic, or require that notifications be sent on success or failure, corrective actions should be used instead of an Agent-side response action.

To access this page, follow steps 1 through 4 in Section 8.7.2, "Creating Corrective Actions for Metrics."

8.7.6.1 Specifying Commands and Scripts

You can specify a single command or execute a script. You cannot specify special shell command characters (such as > and <) as part of the response action command. If you must include these types of special characters in your response action commands, you should use them in a script, then specify the script as the response action command.

If using a script, make sure the script is installed on the host machine that has the Agent. If using shell scripts, make sure the shell is specified either in the Response Action command line:

Script/Command: `/bin/csh myScript`

... or within the body of the script itself:

Script/Command: `myScript`

... where `myScript` contains the following:

```
!#/bin/csh<
<rest of script>
```

8.7.6.2 Using Target Properties in Commands

You can use target properties in a command. Click **Show Available Target Properties** to display target properties you can use in the Script/Command field. The list of available target properties changes according to the type of target the response action is to run against.

Use Target Properties as command-line arguments to the script or command, then have the script reference these command-line arguments. For example, to use the `%OracleHome%` and `%SID%` target properties, your command might appear as follows:

```
/bin/csh MyScript %OracleHome% %SID%
```

.... and your script, `MyScript`, can reference these properties as command-line arguments. For example:

```
IF $1 = 'u1/bin/OracleHome' THEN...
```

Target properties are case-sensitive. For example, if you want to access the Management Agent's Perl interpreter, you can specify `%perlBin%/perl <my_perl_script>` in the Script/Command field.

8.7.6.3 Using Advanced Capabilities

You can get other target properties from the target's XML file in the `OracleHome/sysman/admin/metadata` directory, where OracleHome is the Oracle home of the Management Agent that is monitoring the target. In the XML file, look for the `PROP_LIST` attribute of the `DynamicProperties` element to get a list of properties that are not listed in the `targets.xml` entry for the target.

The following example is an excerpt from the `hosts.xml` file:

```
<InstanceProperties>
 <DynamicProperties NAME="Config" FORMAT="ROW"
 PROP_LIST="OS;Version;OS_patchlevel;Platform;Boottime;IP_address">
   <ExecutionDescriptor>
   <GetTable NAME="_OSConfig"/>
   <GetView NAME="Config" FROM_TABLE="_OSConfig">
   <ComputeColumn NAME="osName" EXPR="Linux" IS_VALUE="TRUE"/>
   <Column NAME="osVersion"/>
   <Column NAME="osPatchLevel"/>
   <Column NAME="Platform"/>
   <Column NAME="Boottime"/>
   <Column NAME="IPAddress"/>
   </GetView>
   </ExecutionDescriptor>
   </DynamicProperties>
   <InstanceProperty NAME="Username" OPTIONAL="TRUE" CREDENTIAL="TRUE">
   <ValidIf>
   <CategoryProp NAME="OS" CHOICES="Linux"/>
   </ValidIf>
   <Display>
   <Label NLSID="host_username_iprop">Username</Label>
   </Display>
   </InstanceProperty>
   <InstanceProperty NAME="Password" OPTIONAL="TRUE" CREDENTIAL="TRUE">
   <ValidIf>
   <CategoryProp NAME="OS" CHOICES="Linux"/>
   </ValidIf>
   <Display>
   <Label NLSID="host_password_iprop">Password</Label>
   </Display>
   </InstanceProperty>
</InstanceProperties>
```

9

Configuring Software Library

This chapter describes how to configure a new Software Library, the various users and the privileges required to access the Software Library, and finally how to maintain an existing Software Library in the Enterprise Manager Cloud Control environment.

In particular, this chapter covers the following:

- Overview of Software Library
- Users, Roles, and Privileges
- Software Library Storage
- Prerequisites for Configuring Software Library
- Configuring Software Library Storage Location
- Using Software Library Entities
- Tasks Performed Using the Software Library Home Page
- Maintaining Software Library

9.1 Overview of Software Library

Oracle Software Library (Software Library) is one of the core features offered by Enterprise Manager Cloud Control. Technically, it is a repository that stores 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, from the **Enterprise** menu, select **Provisioning and Patching**, then click **Software Library.** On the Software Library home page, as shown in Figure 9–1, there are two types of folders: Oracle-owned folders (marked by a lock symbol) and User-owned folders.

Oracle-owned folders and their contents (including other subfolders and entities) ship with the product by default, and appear on the Software Library home page after Software Library is configured. User-owned folders are logical top level folders that the user creates to organize the entities that he/she intends to create.

Figure 9–1 Software Library Console

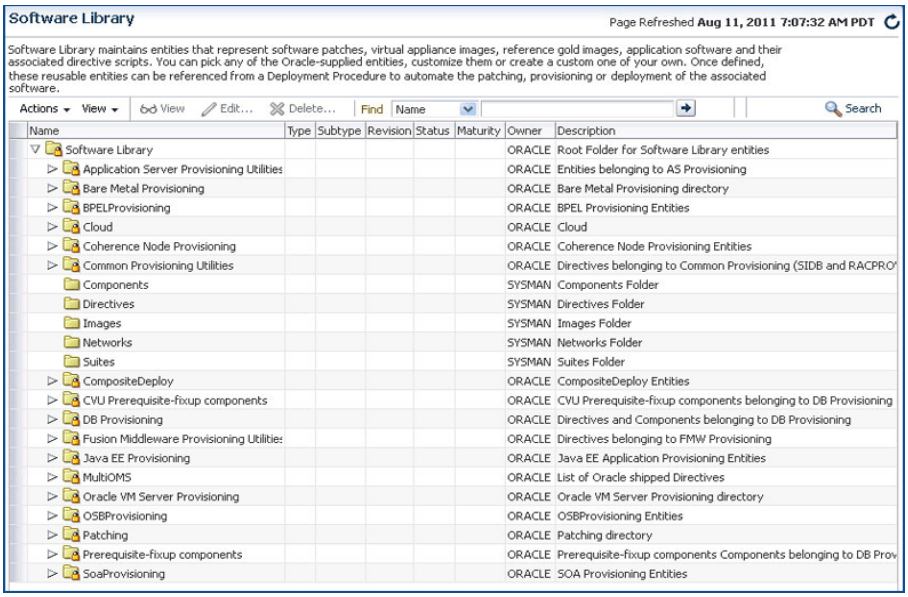

The Software Library Page is centralized media storage for all Enterprise Manager entities. For example,

- Self Update entities like plug-ins, connectors, DB workload, and so on.
- Provisioning and Patching entities like gold images, application archives, perl/shell scripts, and so on.

Advantages:

- Software Library facilitates patching and provisioning tasks in both the modes, which are Online mode and Offline mode. For example, if database patches cannot be downloaded directly from *My Oracle Support*, you can download then separately and stage them in Software Library for offline deployment.
- Starting with Enterprise Manager Cloud Control 12c, Referenced File Locations are supported, which means that the Software Library allows you to leverage your organizations existing IT infrastructure (like file servers, web servers, or storage systems) to stage the files to host targets as part of a provisioning or patching activity.
- Software Library allows you to organize the entities, which basically refer to the software binaries or directive scripts in your enterprise, into logical folders for efficient management.

From the Software Library Console page, you can perform the following tasks:

- Configure Software Library Storage, see Section 9.5, "Configuring Software Library Storage Location" for more information.
- Create Software Library Entities. For example, Creating a Generic Component, Creating Directives, and so on.
- Manage Software Library Entities. For example, Viewing Entities, Editing Entities, Deleting Entities, Searching Entities, and so on.

9.2 Users, Roles, and Privileges

Software Library folders and entities that ship with the product, by default are viewable by all the Enterprise Manager users. Fine grained privileges provide a way to control user access to the different entities in the Software Library. Administrator by default do not have any Software Library privileges, it is for the Super Administrator, to grant access, privileges to an Administrator.

Software Library users roles can be broadly classified as:

- **Designers** are administrators who perform design time tasks like setting up Software library, migrating entities, granting privileges to the Operators, deleting entities, and so on. They can perform both the design time activities, and run-time activities that the Operator can perform. Designers in Enterprise Manager Cloud Control can be granted Super Administrator role or the `EM_PROVISIONING_DESIGNER` role which allows him to create and maintain any Software Library entity.
- **Operators** are administrators who can perform run-time activities like creating components, creating directives, and so on. Operators are typically granted roles like `EM_PROVISIONING_OPERATOR` or `EM_PATCH_OPERATOR` and so on.

Any Enterprise Manager user requires, at the very least, a view privilege on an entity for the entity to be visible on the Software Library Home page. Users will not be able to see this entity till the Super Administrator grants them at least a view privilege on the entity.

> **Note:** All the folders and entities that ship with the product also know as the Oracle-owned entities, by default are viewable by all the Enterprise Manager users.

Administrator by default do not have any Software Library privileges, it is for the Super Administrator, to grant access, privileges to an Administrator. Table 9–1 describes all the available Software Library privileges that can be granted to a user or role.

Users and roles can be granted privileges on specific entities by the owner of the entity or the Super Administrator. For more details, see *Oracle Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching*.

Table 9–1 Software Library Privileges for Administrators

Resource Type	Description
View any Template Entity	Ability to view any Template Entity
Export Any Software Library Entity	Ability to export any Software entity
Edit any Software Library Entity	Ability to edit any Software Library entity
Manage Any Software Library Entity	Ability to create, view, edit, and delete any Software Library entity
Import Any Software Library Entity	Ability to import any Software Library entity
Create Any Software Library Entity	Ability to create any Software Library entity
View Any Software Library Entity	Ability to view any Software Library entity
View Any Assembly Entity	Ability to view any Assembly entity

Table 9–1 (Cont.) Software Library Privileges for Administrators

Resource Type	Description
Grant Any Entity Privilege	Ability to grant view, edit, and delete privileges on any Software Library entity.This privilege is required if the user granting the privilege on any entity is not a Super Administrator or owner of the entity.

Table 9–2 describes all the primary users of Software Library, and their associated privileges:

Table 9–2 Roles and Privileges

Role	Software Library Privileges
Super Administrator	All Software Library Privileges
`EM_PROVISIONING_DESIGNER` (Designer)	Create Any Software Library Entity
`EM_PROVISIONING_OPERATOR` (Operator)	View Any Software Library Entity
`EM_PATCH_OPERATOR`	Create Any Software Library Entity View Any Software Library Entity
`EM_USER` (Administrator)	Access Enterprise Manager

Super Administrators have complete privileges on all the entities present in Software Library, and can exercise access control on the entities by granting one or more privileges, and later revoking the previously granted privilege to another user or role.

Designers by default are given create privileges, which allow them to create entities and manage them.

Operators by default are given view privileges, which allow them to view all the entities in Enterprise Manager Cloud Control.

Any Enterprise Manager user requires, at the very least, a view privilege on an entity for the entity to be visible on the Software Library console. The Super Administrator can choose to grant additional privileges described in Table 9–1 to the user or role. Users will not be able to see this entity till the Super Administrator grants them at least a view privilege on the entity.

9.3 Software Library Storage

The Software Library Administration console allows you to configure and administer Software Library. To start using the Software Library, you must add at least one upload file storage location (OMS Shared location, or OMS Agent location) on the host where the OMS is running. A storage location in Software Library represents a repository of files, these files are either uploaded by Software Library, or generated and saved by some user-owned process.

To access the administration console, log into Enterprise Manager Cloud Control with Administration access, and follow these steps:

In Cloud Control, from **Setup** menu, select **Provisioning and Patching**, and then click **Software Library**.

OR

In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching** and then, click **Software Library.** On the Software Library home page, from **Actions** menu, select **Administration**.

Figure 9–2 Software Library Administration

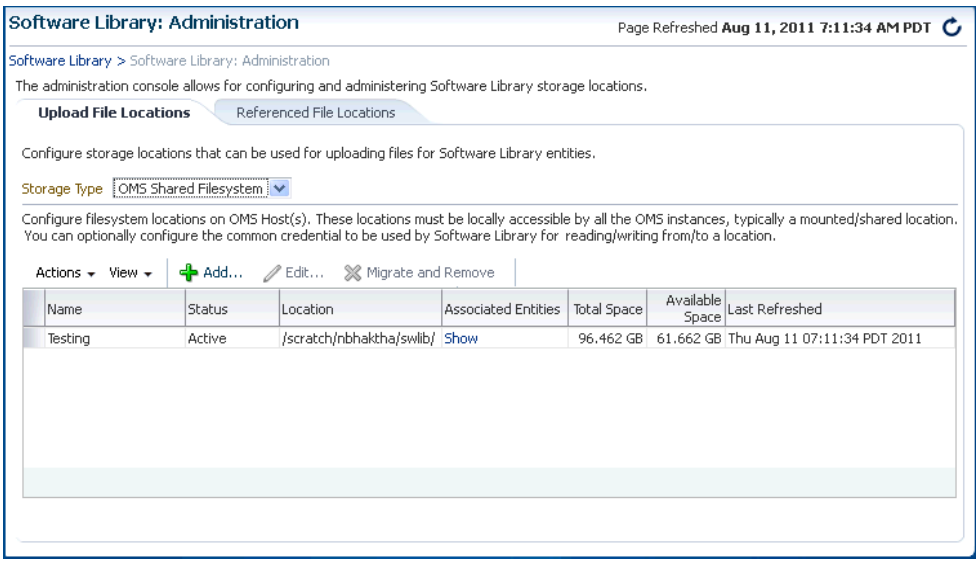

The Software Library Administration Page as shown in Figure 9–2 is a GUI based screen, that enables you to create one or more storage locations to store or refer to files that are associated with an entity. To view the entities present in the storage location, click **show** on the Administration page. You can create a storage location on the OMS or the agent running on the same host as the OMS. With Enterprise Manager 12*c*, a new feature called Referenced File Location has been introduced, wherein Software Library entities can refer to files that are stored on another host. These locations are however read-only for Software Library, will not be used for uploading files.

The space requirements for configuring Software Library depends on the amount of space required for storing the software binaries, and its associated scripts. Understandably, this space requirement increases over a period of time as you create more entities. Depending on the features or software required for provisioning and patching, you must decide and configure the Software Library storage space.

> **Note:** When a storage location starts running out of space, then it is important to deactivate the configured storage location so that no new uploads can happen to this location. For more information about removing a storage location, see Section 9.8

The following types of storage locations are available:

- Upload File Locations
- Referenced File Location

9.3.1 Upload File Locations

Upload File Locations are locations configured for storing files uploaded by Software Library as part of creating or updating an entity.

> **Note:** For Software Library to become usable, at least one upload file location must be configured. On adding the first upload file location, the default Software Library metadata for all installed plug-ins is imported from the OMS Oracle home, and a job is submitted. Ensure that you wait for this job to complete successfully, before performing other patching or provisioning operations.

As a prerequisite, before using Upload File Locations as storage option, you must set credentials for using an OMS Shared File System or OMS Agent File System:

- For multiple OMS environment, all the OMS hosts must have a preferred normal host credential set.

  > **Note:** When OMS instances are added, it is necessary to ensure that the configured locations are accessible from the designated host where the new OMS will be provisioned.
  >
  > For a OMS that will be provisioned using the Add OMS functionality, the shared location configured as upload location should be mounted on the designated host, and verified manually.

- For OMS Agent File System location configuration, a credential (preferred or named) has to be specified.

Upload File Locations support two storage options:

OMS Shared File System

An OMS Shared File System location is required to be shared (or mounted) across all the Oracle Management Server (OMS) hosts. This option is ideal for UNIX systems.

For single OMS environments, you can configure the Software Library either on the host where the OMS is running, or in a shared location. However, in future, if you plan to expand the single OMS setup to a multiple OMS setup, then local file system path is not recommended.

For multiple OMS environments, Oracle recommends you to configure the Software Library in a non-local, shared file system path that is accessible to all Oracle Management Servers in the environment. If you are implementing multiple management servers for high availability you should also make the Software Library file system highly available. Besides accessibility and availability, it is important to ensure that there is enough space available for the storage of software binaries, and associated scripts for the entities that you want to create and store

OMS Agent File System

An OMS Agent File System location should be accessible to the agent running on the host machine where the OMS is deployed, and is recommended for multiple OMS setup on Windows. To use this storage option, ensure that you have a preferred, or a named credential for the OMS host. Click **Change Credential** to change the associated credential to be used to access this location. Ensure that credential associated with the location is viewable by all Software Library designers, so that other designers can upload, and stage the files associated with any entity.

9.3.2 Referenced File Location

Referenced File Locations 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:

- **HTTP**: An HTTP storage location represents a base URL which acts as the source of files that can be referenced.

  For example, the base URL `http://my.files.com/scripts` could be configured as an HTTP location for sourcing files such as `http://my.files.com/scripts/perl/installMyDB.pl` or `http://my.files.com/scripts/linux/stopMyDB.sh.`

- **NFS**: An NFS storage location represents an exported file system directory on a server. The server need not be an Enterprise Manager host target.

  For example, the directory `/exported/scripts` is exported on server `my.file.server` could be configured as an NFS location for sourcing files such as `/exported/scripts/generic/installMyDB.pl` or `/exported/scripts/linux/stopMyDB.sh` once mounted on a target host file system.

- **Agent**: An Agent storage location is similar to the OMS Agent File System option, but can be any host monitored by an Enterprise Manager Agent. The Agent can be configured to serve the files located on that host.

  For example, the directory `/u01/binaries` on the Enterprise Manager Host `my.em.file.server` could be configured as an Agent location for sourcing files such as `/u01/binaries/rpms/myCustomDB.rpm` or `/u01/binaries/templates/myTemplate.tar.gz`.

  These locations require a named credential to be associated which will be used to access the files from the base location on the host through the Enterprise Manager Agent.

9.4 Prerequisites for Configuring Software Library

To administer the different storage types, and to configure software library, keep the following points in mind:

- Depending on the features or software required for provisioning and patching, you must decide and configure the Software Library storage space. The storage needs change based on the usage pattern.
- On UNIX systems, Oracle recommends that you configure at least one OMS Shared File System location that is accessible from all the OMS hosts.

  On Windows systems, Oracle recommends configuring OMS Agent File System storage.

- Each OMS host must have a preferred normal host credential set before configuring the location. For OMS Agent File System location configuration, a credential (preferred or named) has to be specified.
- You (the user configuring the Software Library) must have view privilege on all the OMS, and the agent targets running on the host machine. As per the accessibility verification, you must be able to view, and edit these newly configured locations using the credentials described in the proceeding point.

- All the credentials used while configuring the locations, must be viewable by all the Software Library designers, as this will enable the designers to upload the files to the Software Library storage. Additionally, designers can also stage the uploaded files to other hosts for provisioning or patching activities.
- To add an OMS Agent storage location, ensure that you have view privileges on the target OMS host, and the agents running on that target host.

9.5 Configuring Software Library Storage Location

System Administrators are responsible for configuring a storage location. Only after the storage location is configured, you can start uploading the entity files.

You can configure the Software Library in one of the following locations:

- Configuring an OMS Shared Filesystem Location
- Configuring an OMS Agent Filesystem Location
- Configuring a Referenced File Location

9.5.1 Configuring an OMS Shared Filesystem Location

To configure an OMS Shared File System storage location that can be used for uploading Software Library entity files, perform the following steps:

1. From the **Setup** menu, select **Provisioning and Patching**, then select **Software Library.**
2. On Software Library: Administration page, select **OMS Shared Filesystem**.
3. To add a new OMS Shared File System, click **+Add**.
4. In the Add OMS Shared File System location dialog box, provide a unique name, and location on the OMS host, where you want to set up the upload location.

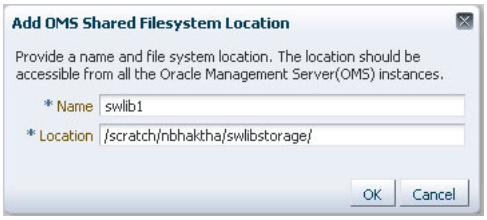

Ensure that the configured storage location is a shared location that is accessible by all the OMS instances. For a Multi OMS setup, set the Normal Preferred Credentials for all the OMS(s).

When you configure an upload location for the first time, a metadata registration job is submitted which imports all the metadata information of all the installed plug-ins from the Oracle home of the OMS.

To track the progress of the job, from **Enterprise** menu, select **Job**, and then click **Activity**. On the Job Activity Page, in the Advanced Search region, enter the name of the job, choose **Targetless** as the Target Type, and then click **Search**. Typically, the name of the job starts with `SWLIBREGISTERMETADATA_*`.

If the Import job fails, see Section 9.8 for information on Re-importing metadata for Oracle-owned files.

5. Click **OK** to create a new entry for the storage location in the table, with details like **Name**, **Location**, **Host**, **Status**, and **Host Credentials.**

In addition to this, you can click **Associated Entities** to view or search the entities present in the selected upload location.

9.5.2 Configuring an OMS Agent Filesystem Location

To configure an OMS Agent location, perform the following steps:

1. From the **Setup** menu, select **Provisioning and Patching,** then select **Software Library.**
2. On the Software Library: Administration page, select **OMS Agent Filesystem**.
3. Click **+Add,** in the Add OMS Agent File System Location dialog box, enter the following details:

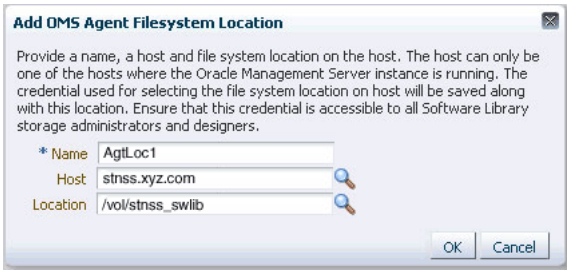

   a. In the **Name** field, enter a unique name for the storage.

   b. In the **Host** field, click the magnifier icon. From the Search and Select: Hosts dialog box, select a host where the OMS is running, and click **Select**.

      For example, `xyz.acme.com`

   c. In the **Location** field, click the magnifier icon. In the Remote File Browser dialog box, click **Login As** to log into the host machine with either Preferred, Named or New credentials.

      Navigate to the location on the host where you want to create the Agent File System, and click **OK**.

      The selected credential is saved along with the host and selected file system path. The saved credential is used to upload files and stage the uploaded files to a host target as part of some provisioning or patching activity.

      Note: The administrator configuring the Software Library must grant view privilege (at least) on the credential chosen to all designers uploading or staging the files to or from this location.

4. Click **OK** to create a new entry for the storage location in the table, with details like **Name**, **Location**, **Host**, **Status**, and **Host Credentials.**

   In addition to this, you can click **Associated Entities** to view or search the entities present in the selected upload location.

These newly configured OMS Agent locations are now available for storing entity files.

9.5.3 Configuring a Referenced File Location

To configure storage location that can be used for referring to files from the Software Library entities, perform the following steps:

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

2. On the Software Library: Administration page, click **Referenced File Locations** tab.
3. To add an HTTP location that can be accessed through a HTTP URL, select **HTTP** from the Storage Type list and click **+Add**.

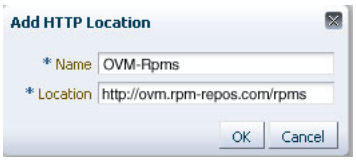

In the Add HTTP Location dialog box, enter a unique name and a HTTP location for the storage that you want to reference, and click **OK**.

A new entry for the storage location is created, with details like **Name**, **Location**, and **Status**.

4. To add an NFS shared location, select **NFS** from the Storage Type list and click **+Add**.

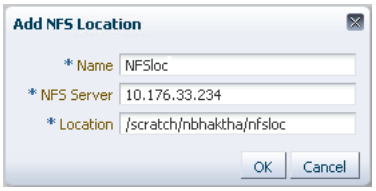

In the Add NFS Location dialog box, do the following:

a. Enter a unique name in the **Name** field for the storage.
b. In **NFS server** field, provide a fully qualified domain name or the IP address of the hosted machine that has NFS services running on them.
c. In the **Location** field, provide the shared location or directory path on the NFS server to define a storage location, then click **OK**.

A new entry for the storage location is created in the table, with details like **Name**, **Location**, and **Status**.

5. To add an Agent location that has read-only privileges set on it, select **Agent** from the Storage Type list and click **+Add**.

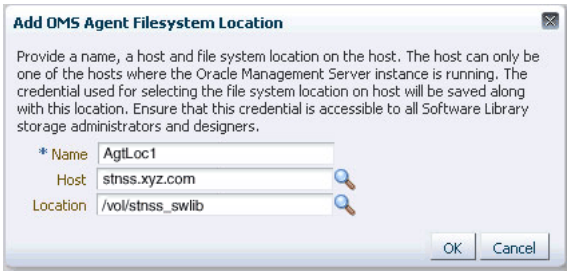

In the Add Agent Location dialog box, enter the following details:

a. In the **Name** field, enter a unique name for the storage.
b. In the **Host** field, click the magnifier icon to select a target from the list available.

   For example, `xyz.company.com`

c. In the **Location** field, click **Login As** to select the credentials and browse the previously selected host.

The credential selected, either Preferred, Named or New, is saved along with the host and selected file system path. The saved credential is used for staging the files to a host target as part of some provisioning or patching activity.

Note: The administrator configuring the Software Library must grant view privilege (at least) on the credential chosen to all designers uploading or staging the files to or from this location.

Note: When you create a new entity, these newly configured Referenced File Locations are available as storage options.

9.6 Using Software Library Entities

To access the Software Library Home Page, in Cloud Control, from the **Enterprise menu,** select **Provisioning and Patching** and then, click **Software Library.** Software Library is a repository that stores certified software binaries such as software patches, virtual appliance images, reference gold images, application software and their associated directive scripts, generally referred to as *Entities.* Accesses and privileges on these entities are decided by the Super Administrators or the owner of the entity.

Entities can broadly be classified as:

Types	Description
Oracle-owned Entities	These entities are available by default on the Software Library Home page, once the Software Library is configured. In the following graphic, all the entities that are owned by **Oracle**, qualify as Oracle-owned entities, and all the folders that appear with a lock icon against them are Oracle-owned folders like Application Server Provisioning, Bare Metal Provisioning, Cloud, and so on.
Custom Entities	These entities are created by the Software Library users. In the following graphic, entities like Components, Directives, Images, Network, and so on are the User-owned entities.

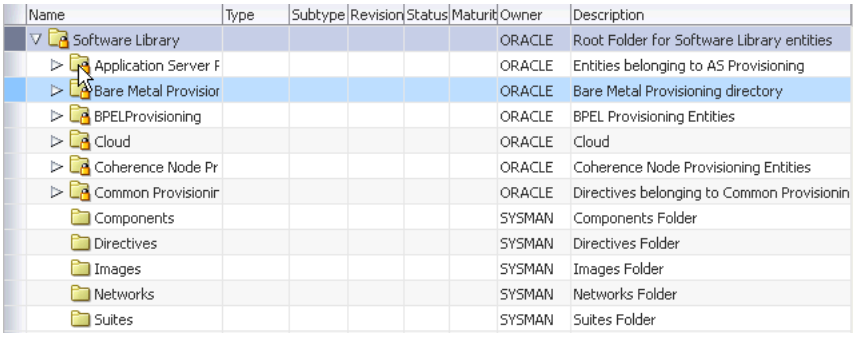

Note: All Oracle-owned folders (and entities) are available on the Software Library Home page by default. The Oracle-owned folders have a read-only privilege, so you cannot select these folders to create an entity. You must create a custom folder to place your entities in them.

A number of lifecycle management tasks like patching and provisioning deployment procedures make use of the entities available in Software Library to accomplish the desired goal. Here is a pictorial representation of how a Provisioning Deployment Procedure and a Patching Deployment Procedure makes use of the entities available in the Software Library:

Figure 9–3 Using Software Library Entities for Provisioning and Patching Tasks

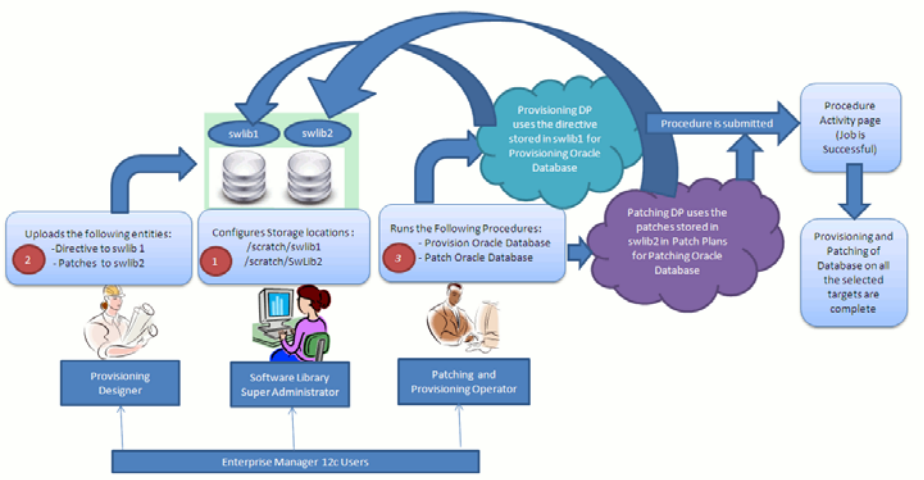

9.7 Tasks Performed Using the Software Library Home Page

From the Software Library Home page, you can do the following:

- Organizing Entities
- Creating Entities
- Customizing Entities
- Managing Entities

9.7.1 Organizing Entities

Only designers who have the privilege to create any Software Library entity, can create folders. To create a custom folder, perform the following steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library.**
2. On the Software Library Home page, from **Actions** menu, click **Create Folder** to create a custom folder of your own.

   The custom folder can contain User-owned folders, entities, and customized entities created by using the *Create Like* option.
3. In the Create Folder dialog box, enter a unique name for the folder. Also, select the parent folder in which you want to create this new custom folder and click **Save.**

   For example, if the root folder is `Software Library` and you created a custom folder in it called `Cloud12gTest`, then the Parent Folder field is populated as follows: `/Software Library/Cloud12gTest`.

   Note: Only the owner of the folder or the Super Administrator has the privilege to delete the folder, nobody else can.

9.7.2 Creating Entities

From the Software Library Home page, you can create the following entities:

- Creating Generic Components
- Creating Directives

9.7.2.1 Creating Generic Components

To create a generic component from the Software Library Home page, 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 a custom folder that is not owned by Oracle.

   Note: You cannot create a generic component in an Oracle Owned Folder. To create a custom folder, see Section 9.7.1.
3. From the **Actions** menu, select **Create Entity** and click **Component**. Alternately, right click the custom folder, and from the menu, select select **Create Entity** and click **Component**.
4. From the Create Entity: Component dialog box, select **Generic Component** and click **Continue**.

   Enterprise Manager Cloud Control displays the Create Generic Component : 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 becuase 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, you can customize the generic component that you are creating by adding some new properties or updating the existing properties of the component.

   Note: Select **Shared Type** to reuse the component property. Shared Type can be stored as a template, which can be used for creating different and more complicated top level types.

   To add a new property, do the following, and click **Next**:

   a. Select **Top Level Type** or **Shared Type**, and click **Add**.

   b. Enter a unique name for the property. Depending on the property type selected, enter an initial or default value for the property.

   c. To add a constraint, specify the Minimum or Maximum value for the selected property type, and click **Add Constraint**.

      The Configured Constraints table lists all the constraints added. To remove a particular constraint from the property, select the property and click **Remove**.

7. On the Select Files page, you can select one or more files to be associated with the entity. Select one of the following options:

   - **Upload Files:** If you want to upload some entity files from the local file system or the agent machine to the selected destination location.

     To select the destination location, in the Specify Destination section, in the **Upload Location** field, click the magnifier icon to select one of the following options:

     – **OMS Shared FileSystem**

     – **OMS Agent FileSystem**

     The corresponding Storage Type and Location Path of the selected location is populated.

     Note:

     To upload the files, the status of the storage location to which the files are to be uploaded should be **Active**.

     If you select OMS Agent Filesystem location, then ensure that you have the necessary privileges to access the location

     ---

     In the Specify Source section, enter the location from where the files are being sourced, these locations can wither be local file system or a remote file system monitored by the Management Agent. Select one of the following options for File Source:,:

     – If you select **Local Machine**, and click **Add**, the Add File dialog box appears. Click **Browse** to select the entity file from the source location, and give a unique name , and click **OK**.

       You can upload the files to any of the configured storage locations available in OMS Shared Filesystem location or OMS Agent Filesystem location

     – If you select **Agent Machine**, select the name of the host machine from where you want to pick up the entity files. Click **+Add** and log into the host machine with the desired credentials. For more information about the different credential types and their setup, see *Oracle Enterprise Manager Lifecycle Management Guide*.

       Once you log into the host machine, browse to the location where the files to be uploaded are present. You can upload the files to any of the configured storage locations available in OMS Shared Filesystem location or OMS Agent Filesystem location.

   - **Refer Files:** If you select the **Refer Files** option, you only need to enter the source location details, since you are not technically uploading anything to the Software Library. In the Specify Source section, select from **HTTP**, **NFS**, or **Agent** Storage types, and click OK. The corresponding Storage Type and Location Path of the selected location is populated.

     Click **+Add** to reference the entity present at the selected Referenced File Location. In the Add Referenced File dialog box, enter a relative path to the file under Base Location. Click **Stage As** to organize the file in a temporary stage location with a unique name.

     For details about each of these storage options, see Section 9.5.3

8. On the Set Directives page, click **Choose Directives** to associate a component with one or more directives. Click **Next**.
9. On the Review page, review all the details, and click **Finish** to create the component and upload it on the Software Library.

9.7.2.2 Creating Directives

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

To create a directive from a Software Library Home page, 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 a custom folder that is not owned by Oracle.

   Note: You cannot create a generic component in an Oracle Owned Folder. To create a custom folder, see Section 9.7.1.
3. From **Actions** menu, select **Create Entity** and click **Directive**. Enterprise Manager Cloud Control displays the Create Entity: Directives wizard.
4. 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, and it may report a conflict, this is becuase 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.
5. On the Configure page, specify the command line arguments that must be passed to the directive to configure it. This command provides the parameters required to execute the directive.

   To add the command line arguments or parameters, click **Add**.

   In the Add Command Line Arguments dialog box, enter the values in the following fields:

   - **Argument Prefix,** is a switch or a constant command line argument.

     The Argument Prefix eliminates the error-prone task of manually specifying the order of the parameter executions in a given directive. This is specially useful when a directive is made of multiple parameters.

     Oracle recommends that you create command line arguments using an Argument Prefix.
   - **Property Name,** is the name of the property, that must be a string value.
   - **Argument Suffix,** is the text that must follow the command line property.

     Though the suffix is rarely used, it determines how the parameters must be executed, based on the suffix value.

All the parameters added appear in the order of addition against the **Command Line** field.

To change the order of the parameter or edit any property of an existing parameter, click **Edit**.

To remove any of the parameters, click **Remove**.

In the Configuration Properties section, select either **Bash** or **Perl** as defined in the script.

Select **Run Privileged** to run the script with `root` privileges.

6. On the Select Files page, you can select one or more files to be associated with the entity. Select one of the following options:

   - **Upload Files:** If you want to upload some entity files from the local file system or the agent machine to the selected destination location.

     To select the destination location, in the Specify Destination section, in the **Upload Location** field, click the magnifier icon to select one of the following options:

     – **OMS Shared FileSystem**

     – **OMS Agent FileSystem**

     The corresponding Storage Type and Location Path of the selected location is populated.

     Note:

     To upload the files, the status of the storage location to which the files are to be uploaded should be **Active**.

     If you select OMS Agent Filesystem location, then ensure that you have the necessary privileges to access the location

     ---

     In the Specify Source section, enter the location from where the files are being sourced, these locations can wither be local file system or a remote file system monitored by the Management Agent. Select one of the following options for File Source:,:

     – If you select **Local Machine**, and click **Add**, the Add File dialog box appears. Click **Browse** to select the entity file from the source location, and give a unique name , and click **OK**.

       You can upload the files to any of the configured storage locations available in OMS Shared Filesystem location or OMS Agent Filesystem location

     – If you select **Agent Machine**, select the name of the host machine from where you want to pick up the entity files. Click **+Add** and log into the host machine with the desired credentials. For more information about the different credential types and their setup, see *Oracle Enterprise Manager Lifecycle Management Guide*.

       Once you log into the host machine, browse to the location where the files to be uploaded are present. You can upload the files to any of the configured storage locations available in OMS Shared Filesystem location or OMS Agent Filesystem location.

- **Refer Files:** If you select the **Refer Files** option, you only need to enter the source location details, since you are not technically uploading anything to the Software Library. In the Specify Source section, select from **HTTP**, **NFS**, or **Agent** Storage types, and click OK. The corresponding Storage Type and Location Path of the selected location is populated.

  Click **+Add** to reference the entity present at the selected Referenced File Location. In the Add Referenced File dialog box, enter a relative path to the file under Base Location. Click **Stage As** to organize the file in a temporary stage location with a unique name.

  For details about each of these storage options, see Section 9.5.3

7. On the Review page, review all the details, and click **Finish** to create the component and upload it on the Software Library.

9.7.3 Customizing Entities

You cannot edit an entity present in an Oracle owned folder. However, to edit an Oracle-owned entity, you can make a copy of the entity and store it in a custom folder. Since you now have full access on the entity, you can customize the entity based on your requirement and may even choose to grant other users access to this entity.

To create a custom entity from an Oracle owned entity, 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 an entity or alternately, search and select an entity.

   For more information about searching for an entity, see Section 9.7.4.7.
3. From **Actions** menu, select **Create Like**.
4. On the Create Like: <Entity Name> dialog box, enter a name that is unique to the parent folder and a description for the entity.

   By default, the root directory Software Library is preselected in the **Parent Folder** field.

   To change the parent folder and organize the entities, click **Change Parent Folder**. and select the desired folder.
5. Click **OK** to apply the changes.

   The new entity appears in the Entities table, under the selected parent folder.

   You as the owner have all the privileges on the entity, and can update the properties as per your requirement.

   To update the properties of the entity, see Section 9.7.4.6.

For more information on Oracle Owned Entities and User Owned Entities, see Section 9.6.

9.7.4 Managing Entities

From the Software Library Home page you can perform the following maintanance tasks on the existing entities:

- Granting or Revoking Privileges
- Moving Entities

- Changing Entity Maturity
- Adding Notes to Entities
- Adding Attachments to Entities
- Viewing, Editing, and Deleting Entities
- Searching Entities
- Exporting Entities
- Importing Entities

9.7.4.1 Granting or Revoking Privileges

An Enterprise Manager user requires, at the very least, a view privilege on an entity for the entity to be visible on the Software Library Home. The owner or super administrator can choose to grant additional privileges like edit (Update notion) or manage (or full) or at a later point of time, revoke the previously granted privilege.

To grant or revoke privileges, perform the following steps:

1. From the **Enterprise** menu, select **Provisioning and Patching,** then select **Software Library.**
2. To grant or revoke fine-grained privileges to the other users on any entity that you own, select the custom entity and from **Actions** menu, click **Grant/Revoke Privileges.**
3. On Grant/Revoke Privileges on: <entity_name> window, you can either grant or revoke Software Library privileges depending on the users roles and responsibilities in the organization.

   Granting Privileges: To grant one or more new privileges, click **+Add** and search for the users. You can grant them one of the following privileges on the entity you own:

   - **View Software Library Entity:** This is normally an operator privilege where the user can only view the entity on the Software Library Home. The user cannot edit or manage the entity. All the Oracle owned entities can be viewed by all Enterprise Manager users.
   - **Edit Software Library Entity:** This is a designer privilege where a user has Create, Update, and Edit privileges on the entity.
   - **Manage Software Library Entity:** This is a super-administrator privilege where the user has complete access on the entity. With this privilege, you can grant or revoke accesses on this entity to other users, or delete the entity.

   Revoking Privileges: To revoke previously granted privileges, select the user and click **Remove**.
4. Click **Update** to apply the selected grants on the entity.

9.7.4.2 Moving Entities

To move all the revisions of an entity from one folder to another, do the following:

1. From the **Enterprise** menu, select **Provisioning and Patching,** then select **Software Library.**
2. On the Software Library Home page, select an entity or alternately, search and select an entity.

   For more information about searching for an entity, see Section 9.7.4.7.

3. From the **Actions** menu, click **Move Entity** and accept the confirmation.
4. From the Move Entity dialog box, select the destination folder for the entities and click **Set New Parent Folder**.

   Note: Ensure that the source and the destination folders are not owned by Oracle, as you cannot move or edit them.

9.7.4.3 Changing Entity Maturity

When an entity is created from the Enterprise Manager Home, it is present in an Untested state. It is the responsibility of a designer to test the entity, and change the maturity level based on the test result.

To manage the lifecycle and indicate the quality (maturity level) of an entity, 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 an entity or alternately, search and select an entity.

   For more information about searching for an entity, see Section 9.7.4.7.
3. From the **Actions** menu, click **Change Maturity** to change the maturity value an entity after testing.

For example, an Oracle Database Clone component would be tested by selecting it in a deployment procedure interview flow that provisions a database. Once the entity is tested, the designer can change the maturity of the entity to either Beta or Production based on test results. Only when the entity is marked with Production level, the Operator can use it.

9.7.4.4 Adding Notes to Entities

To log information about the changes or updates made to an existing entity, 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 an entity or alternately, search and select an entity.

   For more information about searching for an entity, see Section 9.7.4.7.
3. From **Actions** menu, click **Notes** to include any important information related to the entity. You can also add notes while editing an entity.

   The most recent note appears on top of the table, and the older notes appear below.
4. After updating the details, click **Finish** to submit the changes, and return to the Software Library Home page.

9.7.4.5 Adding Attachments to Entities

To add or upload files that are typically documents (like README, installation, configuration) related to the software the entity represents, 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 an entity or alternately, search and select an entity.

   For more information about searching for an entity, see Section 9.7.4.7.

3. From **Actions** menu, click **Attachments** to include one or more files related to the entity. These files contain some important information about the entity. You can also attach files while editing an entity.

   For example, you can attach a readme file to a patch or a component, attach a test script to a directive and so on. However, you must ensure that the file size of each attachment is not more than 2 MB.

4. Click **Finish** to submit the changes, and return to the Software Library Home page.

9.7.4.6 Viewing, Editing, and Deleting Entities

To view, edit, or delete the details of an entity, 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 an entity or alternately, search and select an entity.

   For more information about searching for an entity, see Section 9.7.4.7.

3. To manage an existing entity, select the entity and perform any of the following functions:

   - **View:** Click **View** icon on the table to view the details of an entity. You cannot update the properties of the entity from here.
   - **Edit:** Click **Edit** icon on the table to update the properties of an entity.

     If you are satisfied with the details, click **Save and Upload** to make the changes available on the Software Library Home page.
   - **Delete:** Click **Delete** icon to remove the entity from the Software Library Home page.

     Note: By deleting an entity, the entity is no longer available for selection, viewing, or editing, and will not be displayed on the Software Library Home page. However, the entity continues to exist in the repository and the associated files, if uploaded, continue to exist in the respective disk storage. To delete the entity completely from the repository and the associated files from the file system, you must purge the deleted entities from the administration page. The purge job not only deletes the files associated with the deleted entity, but removes the deleted entities itself from the repository tables. For more information about how to purge the deleted entities from the storage location, see Section 9.8.4

9.7.4.7 Searching Entities

To search for an entity, perform the following steps:

1. From the **Enterprise** menu, select **Provisioning and Patching,** then select **Software Library.**
2. To search for an entity, perform one of the following operations:

a. **Find:** On the Software Library Home page, you can search for an entity by its **Name**, **Description**, or **Type.** Select the search category, enter the desired value and then click the arrow icon.

   On clicking the arrow icon, the result page displays a number of matching results, and allows you to toggle between the result rows by clicking the up and down arrows.

b. **Search:** To perform a detailed search for an entity, click **Search**. The search option, by default, allows you to search by **Name**, **Description**, **Maturity**, **Status**, **Type**, **File Name** and **Revision** to retrieve a more granular search result.

   Note: If you choose entities that have associated subtypes (like Components), then the page is refreshed with **Subtype** as an additional search category.

   Specify appropriate values in **All** or **Any** of the search fields, and click **Search**.

   To add more search parameters, in the Advanced Search section, click **Add Fields** menu and, select the desired search fields. The selected fields appear in the Advanced Search section as new search parameters. This new search feature enables you to refine your search, and drill down to the most accurate and desired search result.

   To revert to the simple search view, click **Close Search**.

9.7.4.8 Exporting Entities

To export selected entities, perform the following steps:

1. From the **Enterprise** menu, select **Provisioning and Patching,** then select **Software Library.**
2. On the Software Library Home page, from the **Actions** menu, click **Export** to export entities present in the Software Library as a Provisioning Archive (PAR) file.

   The PAR file can be used for recreating the entities on an Enterprise Manager with a different repository.
3. On the Export Software Library Entities page, do the following:
   - Click **+Add** to search and select an entity.
   - In **Directory Location**, enter a directory location accessible to OMS for storing the generated PAR files.
   - In **PAR File**, enter the name of the PAR file with a `.par` extension generated during export.
   - To encrypt and securely store all the secret property values of the PAR file being exported, enter a value in the **Oracle Wallet Password** field.

     Note: Specify the same password for importing this PAR file. For more information on importing, see Section 9.7.4.9.
   - Select **Exclude Associated Files**, to exclude the files, binaries, or scripts associated with an entity, from being exported.

   For example, let us consider that you have a separate test and production environment and want to import only the entities that have been tested and certified in the test environment into production. The entities exported from the test system are made available as a Provisioning Archive (PAR) file. You can now

choose to import this PAR file into the production system (which is identical to the test system) and use the tested entities.

4. Click **Submit** to submit an export job. Once the job runs successfully, the selected entities from the Software Library are exported as a PAR file.

Note: Provisioning Archive Files that were exported from any Enterprise Manager version prior to 12*c* can not be imported into Enterprise Manager 12c.

9.7.4.9 Importing Entities

To import PAR (Provisioning Archive) files into the Software Library or deploy them to an OMS, perform the following steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library.**
2. On the Software Library Home page, from **Actions** menu, click **Import** to import the PAR files.
3. On the Import Software Library Entities page, specify the **PAR File** to be imported.

   To import the PAR file successfully, in the **Password** field, enter the same password that was set on the PAR file to secure the secret property values during export.

   For example, let us consider that you have a separate test and production environment and want to import only the entities that have been tested and certified in the test environment into production. The entities exported from the test system are made available as a Provisioning Archive (PAR) file. You can now choose to import this PAR file into the production system (which is identical to the test system) and use the tested entities.
4. If a revision of the entity being imported already exists in Software Library, then you can overwrite the existing entity with a newer revision during import by selecting **Force New Revision**.

   Note: If a revision of the entity being imported already exists in the repository, and you do not select the Force New Revision option, then import process fails.
5. Click **Submit** to submit an import job. On successful completion of the job, the PAR files are imported into the Software Library.

Note: Provisioning Archive Files that were exported from any Enterprise Manager version prior to 12*c* can not be imported into Enterprise Manager 12c.

9.8 Maintaining Software Library

To maintain the health and proper functionality of the Software Library, the administrator who configured the Software Library, or the Designer who has administration access on it must perform the tasks listed here.

This section includes:

- Periodic Maintenance Tasks

- Re-Importing Oracle Owned Entity Files
- Deleting Software Library Storage Location
- Purging Deleted Entity Files
- Backing Up Software Library

9.8.1 Periodic Maintenance Tasks

Periodically, the Administrator must perform the following tasks for proper functioning of the Software Library:

- Refresh the Software Library regularly to compute the free and used disk space.
- Purge deleted entities to conserve disk space. To do so, see Section 9.8.4.
- Check accessibility of the configured Software Library locations to ensure that they are accessible.

9.8.2 Re-Importing Oracle Owned Entity Files

> **Note:** Re-importing metadata applies only to the Oracle owned files, which means all the entity files that ship with the Enterprise Manager product by default. The metadata of User owned entity files cannot be recovered through the Re-import functionality.

Re-Importing the metadata of Oracle owned entity files is not a periodic activity. Re-import helps to recover the metadata files in one of the following situations:

- If you delete the file system location where the metadata was imported. For example, `/scratch/swlib1/`
- If the import job submitted while creating the first upload location fails.

To re-import the metadata of Oracle owned files, do the following:

1. In Cloud Control, from **Setup** menu, select **Provisioning and Patching,** and then click **Software Library.**
2. On the Software Library Administration page, in the Upload File Location tab, from **Actions** menu, select **Re-Import Metadata** option to submit a job that re-initiates the re-import process.

9.8.3 Deleting Software Library Storage Location

Software Library Storage Administrators have the required privileges to delete a storage location. Before removing a storage location, you are prompted to choose an alternate location where the files will be migrated. On selecting an alternate location, a migration job is submitted, and the location is marked **Inactive**. After successful migration of the entity files to the new location, the location configuration is deleted.

> **Note:** To remove a location from OMS Agent File System or Referenced Agent File System storage, you must have a view privilege on the credentials for the location being removed, and the alternate location where the files are migrated.

To delete a configured storage location, perform the following steps:

1. In Cloud Control, from **Setup** menu, select **Provisioning and Patching,** then select **Software Library.**
2. On the Software Library Administration page, select the storage location, and click **Migrate and Remove**.
3. In the confirmation dialog box, click **Remove** to submit a job, which on successful completion deletes the storage entry from the table.

Note: If you have only one Upload File storage location, either OMS Shared File System or OMS Agent File System, you cannot delete this location. Only if there are more than one configured storage location, the **Migrate and Remove** option is enabled. Also, the migration works only when the same storage type is being used. For example, you cannot migrate from OMS upload location to OMS Agent upload location.

9.8.4 Purging Deleted Entity Files

To purge the deleted entities from all the configured Agent Storage locations, you can run a purge job. Starting with Enterprise Manager 12*c* the purge job can be scheduled. To do so follow these steps:

1. In Cloud Control, from **Setup** menu, select **Provisioning and Patching,** then select **Software Library.**
2. On the Software Library Administration page, select the storage location, and from the Actions menu select **Purge**. The following dialog box appears where you can schedule the purge job:

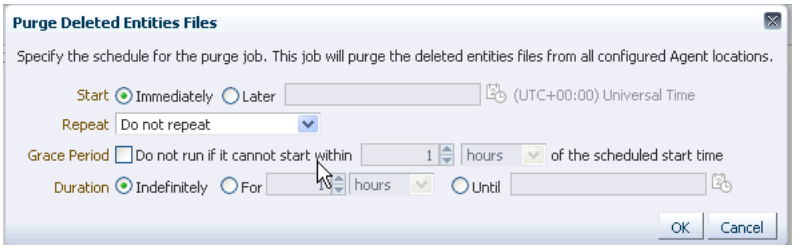

3. Enter allthe details and click OK to submit the job, on successful completion of the job all the deleted entities are removed from the storage location.

9.8.5 Backing Up Software Library

For information about backing up your Software Library, see Chapter 22.

Part II

Security

This section contains the following chapter:

- Configuring Security

10 Configuring Security

This chapter describes how to configure Oracle Enterprise Manager Security. Specifically, this chapter contains the following sections:

- About Oracle Enterprise Manager Security
- Enterprise Manager Authentication
- Enterprise Manager Authorization
- Configuring Secure Communication (SSL) for Cloud Control
- Accessing Managed Targets
- Cryptographic Support
- Setting Up the Auditing System for Enterprise Manager
- Additional Security Considerations

10.1 About Oracle Enterprise Manager Security

Oracle Enterprise Manager provides tools and procedures to help you ensure that you are managing your Oracle environment in a secure manner. The goals of Oracle Enterprise Manager security are:

- To ensure that only users with the proper privileges have access to critical monitoring and administrative data.

  This goal is met by requiring username and password credentials before users can access the Enterprise Manager consoles and appropriate privileges for accessing the critical data.

- To ensure that all data transferred between Enterprise Manager components is transferred in a secure manner and that all data gathered by each Oracle Management Agent can be transferred only to the Oracle Management Service for which the Management Agent is configured.

  This goal is met by enabling Enterprise Manager Framework Security. Enterprise Manager Framework Security automates the process of securing the Enterprise Manager components installed and configured on your network.

- To ensure that sensitive data such as credentials used to access target servers are protected.

  This goal is met by Enterprise Manager's encryption support. The sensitive data is encrypted with an **emkey**. By following the best practice, even the repository owner and the `SYSDBA` will not be able to access the sensitive data.

- To ensure that access to managed targets is controlled through user authentication and privilege delegation.

  This goal is met by configuring the Management Agent with PAM and LDAP for user authentication and using privilege delegation tools like Sudo and PowerBroker.

10.2 Enterprise Manager Authentication

Enterprise Manager authentication is the process of determining the validity of the user accessing Enterprise Manager. The authentication feature is available across the different interfaces such as Enterprise Manager console and Enterprise Manager Command Line Interface (EM CLI).

Enterprise Manager's authentication framework consists of pluggable authentication schemes that let you use the type of authentication protocol best suited to your environment.

The following authentication schemes are available:

- **Oracle Access Manager (OAM) SSO** - Oracle Access Manager is the Oracle Fusion Middleware single sign-on solution. The underlying identity stores will be the Enterprise Directory Identity Stores being supported by Oracle Access Manager. For more information about OAM, see *Oracle® Fusion Middleware Administrator's Guide for Oracle Access Manager 12c Release 1 (11.1.1)*.
- **Repository-Based Authentication**: This is the default authentication option. An Enterprise Manager administrator is also a repository (database) user. By using this option, you can take advantage of all the benefits that this authentication method provides like password control via password profile, enforced password complexity, password life time, and number of failed attempts allowed. During the password grace period, the administrator is prompted to change the password but when the password has expired, it must be changed.For more details, refer to Section 10.2.1, "Repository-Based Authentication".
- **SSO-Based Authentication**: The single sign-on based authentication provides strengthened and centralized user identity management across the enterprise. After you have configured Enterprise Manager to use the Oracle Application Server Single Sign-On, you can register any single sign-on user as an Enterprise Manager administrator. You can then enter your single sign-on credentials to access the Oracle Enterprise Manager console. For more details, refer to Section 10.2.3, "Single Sign-On Based Authentication".
- **Enterprise User Security Based Authentication**: The Enterprise User Security (EUS) option enables you to create and store enterprise users and roles for the Oracle database in an LDAP-compliant directory server. Once the repository is configured with EUS, you can configure Enterprise Manager to use EUS as its authentication mechanism as described in Section 10.2.4, "Enterprise User Security Based Authentication". You can register any EUS user as an Enterprise Manager administrator.

  EUS helps centralize the administration of users and roles across multiple databases. If the managed databases are configured with EUS, the process of logging into these databases is simplified. When you drill down to a managed database, Enterprise Manager will attempt to connect to the database using Enterprise Manager credentials. If successful, Enterprise Manager will directly connect you to the database without displaying a login page.

- **Oracle Internet Directory (OID) Based Authentication** - When using an authentication scheme based on Oracle Internet Directory as the identity store, you can plug in the OID-based authentication scheme to have your applications authenticate users against the OID.
- **Microsoft Active Directory Based Authentication** - When using a Microsoft Active Directory as an identity store, you can plug in this scheme to have your applications authenticate users against the Microsoft Active Directory.

10.2.1 Repository-Based Authentication

Enterprise Manager allows you to create and manage new 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. You can also assign a password profile to the administrator. To create, edit, or view an administrator account:

1. From the **Setup** menu, select **Security**, then select **Administrators**.
2. Click the appropriate task button on the Administrators page. The following screen is displayed:

Figure 10–1 Create / Edit Administrator

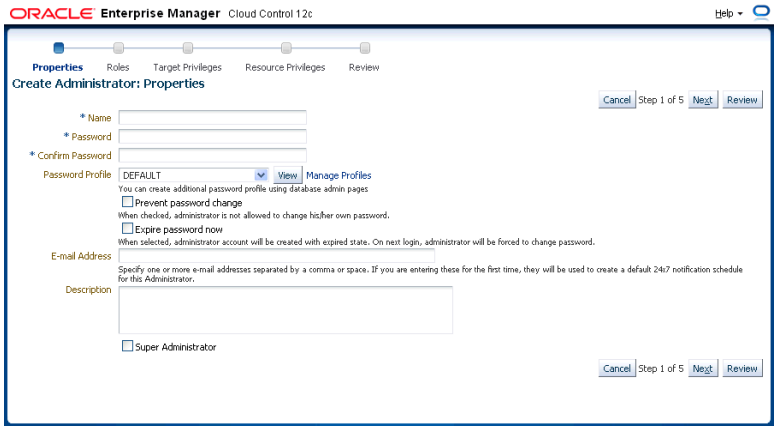

On this page, you can specify the type of administrator account being created and select the password profile. The password cannot be changed by the administrator if the **Prevent Password Change** checkbox is selected.

If you select the **Expire Password Now** checkbox, the password for administrator account will be set to an expired state. If the password has expired, when you login the next time, the following screen is displayed and you are prompted to change the password.

Figure 10–2 Password Expiry Page

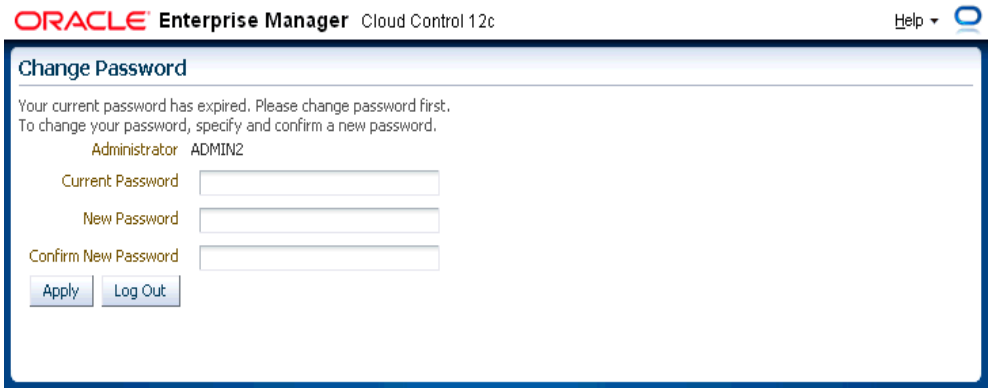

Enter your current password and the new password and click **Apply**. You can now start using Enterprise Manager.

10.2.2 Oracle Access Manager Single Sign-On

When using an Oracle Access Manager Single Sign-On authentication scheme, the underlying identity stores will consist of Enterprise Directory Identity Stores supported by Oracle Access Manager. This section provides instructions on how to configure OAM SSO-based authentication schemes.

Prerequisites

Oracle access manager is installed.

The Oracle Access Manager Single Sign-On server is configured with Oracle HTTP server, Web Gate, and the Oracle Access Manager Identity Store.

1. Run the `emctl config auth` command.

```
emctl config auth oam [-sysman_pwd <pwd>] -oid_host <host> -oid_port <port>
-oid_principal <principal> [-oid_credential <credential>]
-user_base_dn <dn> -group_base_dn <dn>
-oam_host <host< -oam_port <port> [-logout_url <url>] [-is_oam10g] [-user_dn
<dn>] [-group_dn <dn>]
```

   Note: Pass -is_oam10g option only if the OAM version is 10g.

2. Stop each OMS.

```
emctl stop oms -all
```

3. Restart each OMS.

```
emctl start oms
```

10.2.3 Single Sign-On Based Authentication

If you are currently using Oracle Application Server Single Sign-On to control access and authorization for your enterprise, you can extend those capabilities to the Enterprise Manager console.

By default, Enterprise Manager displays the main login page. However, you can configure Enterprise Manager so it uses Oracle Application Server Single Sign-On to authenticate your Enterprise Manager users. Instead of seeing the Enterprise Manager login page, users will see the standard Oracle Application Server Single Sign-On login

page. From the login page, administrators can use their Oracle Application Server Single Sign-On credentials to access the Oracle Enterprise Manager 12c Cloud Control console.

Note:

- You can configure Enterprise Manager to use one of the default Oracle Application Server Single Sign-On or Enterprise User Security features, but not multiple.
- When Enterprise Manager is configured to use Single Sign-On with Server Load Balancer, make sure that the correct monitoring settings have been defined. For details, refer to the chapter on *Cloud Control Common Configurations*.

The following sections describe how to configure Enterprise Manager as an OracleAS Single Sign-On Partner Application:

- Registering Enterprise Manager as a Partner Application
- Removing Single Sign-On Configuration
- Registering Single Sign-On Users as Enterprise Manager Administrators
- Bypassing the Single Sign-On Logon Page

10.2.3.1 Registering Enterprise Manager as a Partner Application

To register Enterprise Manager as a partner application manually, follow these steps:

1. Stop all OMSs by running `emctl stop oms` on each OMS.
2. Enter the following URL to navigate to the SSO Administration page.

   ```
   https://sso_host:sso_port/pls/orasso
   ```

3. Login as `orcladmin` user and click on **SSO Server Administration**.
4. Click **Administer Partner Applications** and then click **Add Partner Application**.
5. Enter the following information on the Add Partner Application page.

   ```
   Name: <EMPartnerName>
   Home URL: protocol://em_host:em_port
   Success URL: protocol://em_host:em_port/osso_login_success
   Logout URL: protocol://em_host:em_port/osso_logout_success
   Administrator Email: user@host.com
   ```

   Note1: host, port, and protocol refer to the Enterprise Manager Host, port and the protocol (http or https) used.

   Note2: The em_host, em_port, email and Enterprise Manager PartnerName need to be replaced appropriately and not typed as shown in this example.

6. Go back to Administer Partner Applications page and click on the Edit icon for <EMPartnerName>.

   Record the values of ID, Token, Encryption Key, Login URL, Single Sign-Off URL, Home URL and write the following in a file osso.txt:

   ```
   sso_server_version= v1.2
   cipher_key=<value of EncryptionKey>
   site_id=<value of ID>
   ```

```
site_token=<value of Token>
login_url=<value of Login URL>
logout_url=<value of Single Sign-Off URL>
cancel_url=<value of Home URL>
sso_timeout_cookie_name=SSO_ID_TIMEOUT
sso_timeout_cookie_key=9E231B3C1A3A808A
```

7. Set the ORACLE_HOME environment variable to WebTier Oracle Home location.

   ```
   setenv ORACLE_HOME /scratch/12c/MWHome/Oracle_WT
   ```

   Then, run the following:

   ```
   $ORACLE_HOME/ohs/bin/iasobf <location of osso.txt> <location
   of osso.conf>
   ```

8. Run the following command on each OMS:

   ```
   emctl config auth sso -ossoconf <osso.conf file loc> -dasurl <DAS URL>
   [-unsecure] [-sysman_pwd <pwd>] [-domain <domain>]-ldap_host <ldap host> -ldap_
   port <ldap port> -ldap_principal <ldap principal> [-ldap_credential <ldap
   credential>] -user_base_dn <user base DN> -group_base_dn <group base DN>
   [-logout_url <sso logout url>]
   ```

   where ldap_host, ldap_port, ldap_principal and ldap_credential are the details of SSO's LDAP.

   The sample output for this command is shown below:

   ```
   Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
   Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
   SSO Configuration done successfully. Please restart Admin & Managed Servers.
   ```

9. Run the following commands on each OMS:

   ```
   emctl stop oms -all
   emctl start oms
   ```

10.2.3.2 Removing Single Sign-On Configuration

To remove the single sign-on configuration, perform the following:

1. Run the following command on each OMS:

   ```
   emctl config auth repos [-sysman_pwd <pwd>]
   ```

   Sample command output:

   ```
   Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
   Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
   Configuring Repos Authentication ... Started
   Configuring Repos Authentication ... Successful
   If you have updated files like httpd.conf (for example, while installing
   WebGate), rollback them.
   If this is a multi-OMS environment, execute this command on remaining servers.
   After that, restart OMS(s) using: 'emctl stop oms -all' and 'emctl start oms'
   ```

2. Bounce all OMSs by issuing the following on each OMS:

   ```
   emctl stop oms -all
   emctl start oms
   ```

10.2.3.3 Registering Single Sign-On Users as Enterprise Manager Administrators

After you have configured Enterprise Manager to use the Single Sign-On logon page, you can register any Single Sign-On user as an Enterprise Manager administrator. You can register single sign-on users using:

- Enterprise Manager Graphical User Interface
- Enterprise Manager Command Line Interface

10.2.3.3.1 Registering Single Sign-On Users Using the Graphical User Interface

You can use the graphical user interface to register single sign-on users by following these steps:

1. Go the Enterprise Manager Console URL.

   The browser is redirected to the standard Single Sign-On Logon page.

2. Enter the credentials for a valid Single Sign-On user. Note: This step requires that an SSO user is already registered with Enterprise Manager.

   If no SSO user is yet registered as Enterprise Manager user, you can create them using the following procedure:

   1. Log in to Enterprise Manager by connecting to Managed Server (MS) directly. eg: https://ms_host:ms_https_port/em.

   2. Log in as a Repository user.

   3. From the **Setup** menu, select **Security** then select **Administrator**

   4. Create SSO users.

3. Log in to Enterprise Manager as a Super Administrator.

4. From the **Setup** menu, select **Security**, then select **Administrators** to display the Administrators page.

   Because Enterprise Manager has been configured to use Single Sign-On, the first page in the Create Administrator wizard now offers you the option of creating an administrator either as an External User or as Repository User.

5. Select **External User Identity Store** and advance to the next page in the wizard.

6. Enter the name and e-mail address of the External User Identity Store user, or click the flashlight icon to search for a user name in the Oracle Internet Directory.

7. Use the rest of the wizard pages to define the roles, system privileges, and other characteristics of the Enterprise Manager administrator and then click **Finish**.

   Enterprise Manager displays a summary page that lists the characteristics of the administrator account.

8. Click **Finish** to create the new Enterprise Manager administrator.

   The External User Identity Store user is now included in the list of Enterprise Manager administrators. You can now verify the account by logging out of the Cloud Control console and logging back in using the External User Identity Store user credentials on the Single Sign-On logon page.

10.2.3.3.2 Registering Single Sign-On Users Using EM CLI s

You can use the following EM CLI command to create Single Sign-On users:

```
emcli create_user -name=ssouser -type=EXTERNAL_USER
```

This command creates a user with the name **ssouser** who is authenticated against the single sign-on user.

Argument	Description
-name	Name of the administrator.
-type	The type of user. The default value for this parameter is EM_USER. The other possible values are: ■ EXTERNAL_USER: Used for single-sign-on based authentication. ■ DB_EXTERNAL_USER: Used for enterprise user based security authentication.
-password	The password for the administrator.
-roles	The list of roles that can be granted to this administrator.
-email	The list of email addresses for this administrator.
-privilege	The system privileges that can be granted to the administrator. This option can be specified more than once.
-profile	The name of the database profile. This is an optional parameter. The default profile used is DEFAULT.
-desc	The description of the user being added.
-expired	This parameter is used to set the password to "expired" status. This is an optional parameter and is set to False by default.
-prevent_change_password	When this parameter is set to True, the user cannot change the password. This is an optional parameter and is set to False by default.
-input_file	This parameter allows the administrator to provide the values for any of these arguments in an input file. The format of value is `name_of_argument:file_path_with_file_name`.

Example 1

```
emcli create_user
        -name="new_admin"
        -email="first.last@oracle.com;joe.shmoe@shmoeshop.com"
        -roles="public"
        -privilege="view_job;923470234ABCDFE23018494753091111"
        -privilege="view_target;<host>.com:host"
```

This example creates an Enterprise Manager administrator named new_admin. This administrator has two privileges: the ability to view the job with ID `923470234ABCDFE23018494753091111` and the ability to view the target `<host>.com:host`. The administrator `new_admin` is granted the PUBLIC role.

Example 2

```
emcli create_user
      -name="User1"
      -type="EXTERNAL_USER"
      -input_file="privilege:/home/user1/priv_file"
```

```
Contents of priv_file are:
  view_target;<host>.com:host
```

This example makes `user1` which has been created externally as an Enterprise Manager user. `user1` will have view privileges on `<host>.com:host`.

Example 3

```
emcli create_user
      -name="User1"
      -desc="This is temp hire."
      -prevent_change_password="true"
      -profile="MGMT_ADMIN_USER_PROFILE"
```

This example sets `user1` as an Enterprise Manager user with some description. The `prevent_change_password` is set to true to indicate that the password cannot be changed by `user1` and the `profile` is set to `MGMT_ADMIN_USER_PROFILE`.

Example 4

```
emcli create_user
      -name="User1"
      -desc="This is temp hire."
      -expire="true"
```

This example sets `user1` as an Enterprise Manager with some description. Since the password is set to expire immediately, when the user logs in for the first time, he is prompted to change the password.

10.2.3.4 Bypassing the Single Sign-On Logon Page

If the OMS is configured with SSO or OAM or some other authentication method, you may want to by-pass the Single Sign-On or OAM authentication under certain circumstances.

To bypass the SSO logon page, connect to the following URL:

1. Connect to `https://ms_host:ms_https_port/em`

   ms_host & ms_https_port are WLS-managed server's hostname & port#. These parameters can be found in the EM_INSTANCE_HOME/emgc.properties file. They are listed as EM_INSTANCE_HOST & MS_HTTPS_PORT in this file.

2. Log in using a repository user's credentials.

10.2.3.5 Restoring the Default Authentication Method

1. Run the following command on each OMS:

   ```
   emctl config auth repos [-sysman_pwd <pwd>]
   ```

   Sample command output:

   ```
   Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
   Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
   Configuring Repos Authentication ... Started
   Configuring Repos Authentication ... Successful
   If you have updated files like httpd.conf (for example, while installing
   WebGate), rollback them.
   If this is a multi-OMS environment, execute this command on remaining servers.
   After that, restart OMS(s) using: 'emctl stop oms -all' and 'emctl start oms'
   ```

2. Run the following commands on each OMS:

```
emctl stop oms -all
emctl start oms
```

10.2.4 Enterprise User Security Based Authentication

Enterprise User Security enables you to create and store Oracle database information as directory objects in an LDAP-compliant directory server. For example, an administrator can create and store enterprise users and roles for the Oracle database in the directory, which helps `centralize t`he administration of users and roles across multiple databases.

> **See Also:** Enterprise User Security Configuration Tasks and Troubleshooting in the *Oracle Database Advanced Security Administrator's Guide*

If you currently use Enterprise User Security for all your Oracle databases, you can extend this feature to Enterprise Manager. Configuring Enterprise Manager for use with Enterprise User Security simplifies the process of logging in to database targets you are managing with the Oracle Enterprise Manager console.

To configure Enterprise Manager for use with Enterprise User Security:

1. Ensure that you have enabled Enterprise User Security for your Oracle Management Repository database, as well as the database targets you will be managing with the Cloud Control console. Refer to *Oracle Database Advanced Security Administrator's Guide* for details.
2. Using the `emctl set property` command, set the following properties:

```
oracle.sysman.emSDK.sec.DirectoryAuthenticationType=EnterpriseUser
oracle.sysman.emSDK.sec.eus.Domain=<ClientDomainName> (For
example:mydomain.com)
oracle.sysman.emSDK.sec.eus.DASHostUrl=<das_url> (For example:
oracle.sysman.emSDK.sec.eus.DASHostUrl=http://my.dashost.com:7777 )
```

   For example:

```
emctl set property -name oracle.sysman.emSDK.sec.DirectoryAuthenticationType
-value EnterpriseUser
```

3. Stop the Oracle Management Service.

> **See Also:** Controlling the Oracle Management Service on page 17-5

4. Start the Management Service.

The next time you use the Oracle Enterprise Manager console to drill down to a managed database, Enterprise Manager will attempt to connect to the database using Enterprise User Security. If successful, Enterprise Manager will connect you to the database without displaying a login page. If the attempt to use Enterprise User Security fails, Enterprise Manager will prompt you for the database credentials.

10.2.4.1 Registering Enterprise Users as Enterprise Manager Users

After you have configured Enterprise Manager to use Enterprise Users, you can register existing enterprise users as Enterprise Manager Users and grant them the necessary privileges so that they can manage Enterprise Manager effectively.

You can register existing enterprise users by using:

- Enterprise Manager Graphic User Interface
- Enterprise Manager Command Line Interface

10.2.4.1.1 Registering Enterprise Users Using the Graphical User Interface

You can use the graphical user interface to register enterprise users by following these steps:

1. Log into Enterprise Manager as a Super Administrator.
2. From the **Setup** menu, select **Security** then select **Administrators** to display the Administrators page. Since Enterprise Manager has been configured to use Enterprise Users, the first page of the Create Administrator wizard will provide the option to create an administrator based on a registered Oracle Internet Directory user or a normal database user.
3. Select Oracle Internet Directory and click **Continue** to go to the next page in the wizard.
4. Enter the name and e-mail address of the Oracle Internet Directory user or click the flashlight icon to search for a user name in the Oracle Internet Directory.
5. Use the rest of the wizard pages to define the roles, system privileges, and other characteristics of the Enterprise Manager administrator and then click **Finish**. Enterprise Manager displays a summary page that lists the characteristics of the administrator account.
6. Click **Finish** to create the new Enterprise Manager administrator.

   The OID user is now included in the list of Enterprise Manager administrators. You can now verify the account by logging out of the Cloud Control console and logging back in using the OID user credentials on the Single Sign-On logon page.

10.2.4.1.2 Registering Enterprise Users Using the Command Line Interface

To register Enterprise Users as Enterprise Manager users using EM CLI, enter the following command:

```
emcli create_user -name=eususer -type=DB_EXTERNAL_USER
```

This command registers the `eususer` as an Enterprise Manager user where `eususer` is an existing Enterprise User. For more details, refer to Registering Single Sign-On Users Using EM CLI.

10.2.5 Microsoft Active Directory Based Authentication

Enterprise Manager uses the authentication capabilities provided by the Oracle WebLogic Server that is part of the OMS. If you are using Microsoft Active Directory as an identity store, you will need to configure it with the Oracle WebLogic Server which is part of the OMS. The following procedure demonstrates how to set up Enterprise Manager authentication using Microsoft Active Directory.

Prerequisites

- Ensure Enterprise Manager Cloud Control 12c is installed and configured properly and that you can log in as a user with Super Admin privileges.
- Ensure Microsoft Active Directory is installed and configured properly.

- Obtain te following from your Microsoft Active Directory administrator. Below is an example of a simple configuration. More complex configurations can be implemented with additional knowledge of LDAP search filters.
  - Active Directory Port
  - Active Directory Prinicpal (User created to authenticate with Active Directory for the Oracle WebLogic Server.
  - Active Directory Principal Password
  - User Base Distinguised Name (DN)
  - Group Base DN

		Example	Your Value
Host	The Active Directory host	server.oracle.com	
Port	The Active Directory Port	389 (LDAP) or 636 (LDAPS)	
Principal User/Password	The Principal User created in Active Directory that will be used to authenticate WebLogic Server. It must be in the Administrators group and belong to the correct Organizational Unit designated in the User base DN. Ensure the "User must change password at next logon" is not checked during setup.		emgcadmin/Welcome11
User Base DN	The User Base Distinguished Name is the container location of valid users who will be granted access to ENTERPRISE MANAGER. Using the default Users container will allow all Active Directory Users to login to ENTERPRISE MANAGER (though they may not have permissions to see/do anything). Using an Organizational Unit will allow you to further restrict access.		
User Base Filter From Name			
User Name Attribute:		sAMAccountName	

	Example	Your Value
User From Name Filter:	(&(sAMAccountName=%u)(objectclass=user))	

10.2.5.1 Configuring WebLogic Server Authentication

Use the following procedure to update

1. As the Weblogic/Enterprise Manager administrator, back up the WLS config.xml file at the following location:

   `../gc_inst/user_projects/domains/GCDomain/config/config.xml`

2. Log in to the WebLogic Admin Console as weblogic. The WebLogic Admin Console URL can be found in the setupinfo.txt file at the following location:

   `$ORACLE_HOME/install/setupinfo.txt`

3. Under Domain Structure, click **Security Realms**.
4. Click **myrealm** and then click the **Providers** tab.
5. Click **Lock & Edit** to enable editing.
6. Click **New** to add a **Provider**.
7. Enter a **Name** for your **Provider** (for example, MS Active Directory).
8. Select **ActiveDirectoryAuthenticator** for *Type* and then click **OK**.

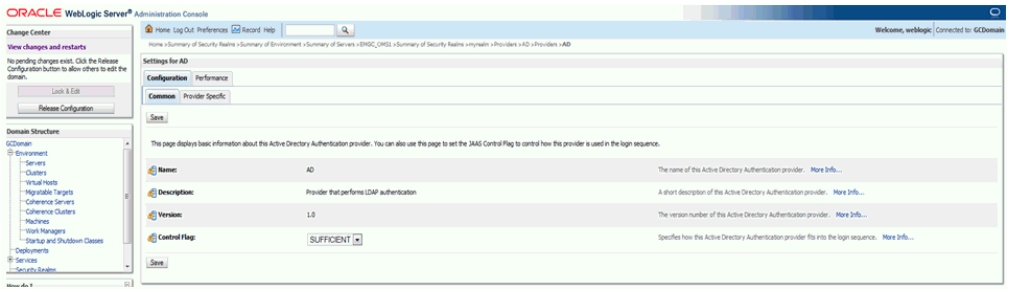

9. On the Providers screen, click the **New Provider** link to begin editing.
10. Set the **Control Flag** to **Sufficient** and then click **Save**.
11. Click the **Provider Specific** tab
12. In the Connection section, enter the following:

    Host: AD Server Host

    Port: 389 (default for LDAP, or 636 for LDAPS)

    Principal:CN=EMGCADMIN,CN=Users,DC=Cloudcontrol,DC=local

    Note: This is the User created in AD steps above and added to Administrators group. The CN/DC string must be confirmed with your Active Directory administrator.

    Credential: pwd for principal

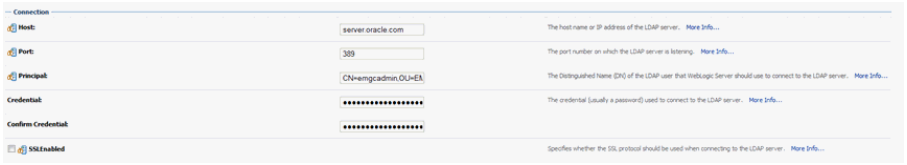

13. In the Users section, set the **User Base DN** to the value provided by your Active Directory administrator. This is the Group or Organization Unit that will have access to Enterprise Manager. To restrict access to a specific set of users, you must use an Organization Unit.

    User Base DN: cn=users,dc=Cloudcontrol,dc=local

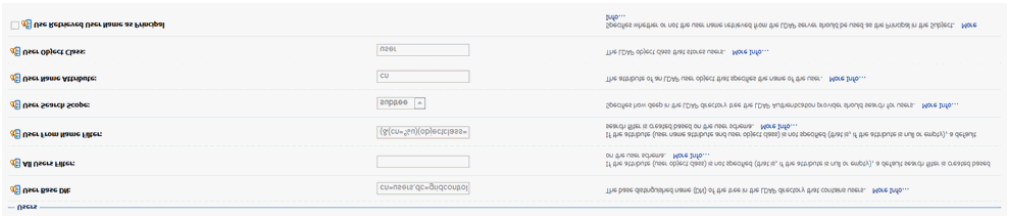

> **Note:** This information must be obtained from the AD Administrator.

14. If you want to use the Login Name instead of the Account Name (which is typically First Last) then you need to set the **User From Name Filter** and **User Name Attribute** as follows:

    User Name Attribute: sAMAccountName

    User From Name Filter: (&(sAMAccountName=%u)(objectclass=user))

15. In the Groups section, enter the following:

    Group Base dn: cn=Users,dc=Cloudcontrol,dc=local

    Note: This information must be obtained from the AD Administrator

16. In the General section, click **Propagate Cause For Login Exception**.

17. Click **Save**.

18. In the Authenticaiton Providers section, click **Reorder** and move your new provider to the top of the list.

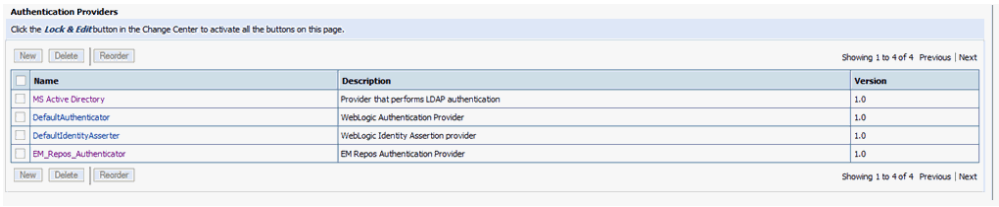

19. Click **Apply & Activate Changes.**

20. There are two options to provision users to Enterprise Manager. You can set a flag to auto-provision all users, or you can manually create them as external users using EM CLI.

    1. To set Auto Provisioning to true, run the following:

    ```
    $ bin/emctl set property -name "em.security.auth.autoprovisioning" -value
    "true"
    Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
    Copyright (c) 1996, 2012 Oracle Corporation.  All rights reserved.
    SYSMAN password:
    ```

    The property *em.security.auth.autoprovisioning* for the OMS *server.oracle.com:4890_Management_Service* has been set to value true

    OMS restart is required to reflect the new property value.

    2. If you do not want all users created automatically, you must manually create them using the EM CLI (after restart).

    ```
    $ bin/emcli create_user -name='TEST' -type='EXTERNAL_USER'
    User "TEST" created successfully
    ```

21. Restart the OMS.

    ```
    $ bin/emctl stop oms -all
    $ bin/emctl start oms
    ```

22. The users will not show up in the Enterprise Manager Administrators UI until they have logged in once.

10.2.5.2 Manage Active Directory Users with External Roles

To assign a group of privileges to the LDAP users, you can create an external role with the same name as the LDAP group. Once the users are authenticated, they will inherit the permissions and privileges granted to the external role automatically.

1. Create a Group in the Active Directory and assign users to the group.
2. From the **Setup** menu, select **Security,** then select **Roles**.
3. Click **Create** .
4. Enter the name of the Active Directory group and a brief description.
5. Check the box for **External**, and then click **Next**.
6. Assign additional **Roles**, and then click **Next**.
7. Assign t**arget privileges,** and then click **Next**.
8. Assign **resource privileges**, and then click **Next**.
9. Review the settings and click **Finish**.

10.2.5.3 Password Management for Active Directory Users

Password management for Active Directory users must be handled through Active Directory. Password changes are not allowed via Enterprise Manager or WebLogic Server.

10.2.5.4 Remove Active Directory Users

An Active Directory user must be deleted from Enterprise Manager to remove access to Cloud Control. If the user remains in Active Directory, they should be removed

from any Groups assigned privileges through External Roles to ensure they cannot login again if auto-provisioning is enabled.

10.2.5.5 Remove Active Directory Authentication

Removing Active Directory authentication will remove all Active Dirctory user accounts from Enterprise Manager.

1. Log in to the WebLogic Server console.
2. Under **Domain Structure**, click on **Security Realms**.
3. Click **myrealm**, then click on the **Providers** tab.
4. Click L**ock & Edit** to enable editing.
5. Click the NT Authenticator provider.
6. Click **Delete**.
7. Click **Save and Activate**.
8. Restart the OMS.

10.3 Enterprise Manager Authorization

Giving the same level of access to all systems to all administrators is dangerous, but individually granting access to tens, hundreds, or even thousands of targets to every new member of the group is time consuming. With Enterprise Manager's administrator privileges and roles feature, this task can be performed within seconds, instead of hours. Authorization controls the access to the secure resources managed by Enterprise Manager via system, target, and object level privileges and roles.

This section describes Enterprise Manager's Authorization model including user classes, roles, and privileges assigned to each user class. The following topics are described:

- Classes of Users
- Privileges and Roles

10.3.1 Authentication Scheme

An authentication scheme is the type of authentication supported by a target type. For example, a host can support a username/password-based authentication, Public Key authentication or Kerberos authentication. In fact, each target type in an enterprise may support different authentication schemes. To accommodate the many authentication schemes that can exist in a managed environment, Enterprise Manger allows you to configure the credentials for these authentication schemes as well.

10.3.2 Classes of Users

Oracle Enterprise Manager supports different classes of Oracle users, depending upon the environment you are managing and the context in which you are using Oracle Enterprise Manager.

The Enterprise Manager administrators you create and manage in the Could Control console are granted privileges and roles to log in to the Cloud Control console and to manage specific target types and to perform specific management tasks. The default super administrator for the Cloud Control Console is the SYSMAN user, which is a database user associated with the Oracle Management Repository. You define the

password for the SYSMAN account during the Enterprise Manager installation procedure.

By restricting access to privileged users and providing tools to secure communications between Oracle Enterprise Manager 12c components, Enterprise Manager protects critical information in the Oracle Management Repository.

The Management Repository contains management data that Enterprise Manager uses to help you monitor the performance and availability of your entire enterprise. This data provides you with information about the types of hardware and software you have deployed, as well as the historical performance and specific characteristics of the applications, databases, applications servers, and other targets that you manage. The Management Repository also contains information about the Enterprise Manager administrators who have the privileges to access the management data.

You can create and manage 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 administrator access categories:

- **Super Administrator**: Powerful Enterprise Manager administrator with full access privileges to all targets and administrator accounts within the Enterprise Manager environment. The Super Administrator, SYSMAN is created by default when Enterprise Manager is installed. The Super Administrator can create other administrator accounts.
- **Administrator**: Regular Enterprise Manager administrator.
- **Repository Owner**: Database administrator for the Management Repository. This account cannot be modified, duplicated, or deleted.

The types of management tasks that the administrator can perform and targets that he can access depends on the roles, system privileges, and target privileges that he is granted. The Super Administrator can choose to let certain administrators perform only certain management tasks, or access only certain targets, or perform certain management tasks on certain targets. In this way, the Super Administrator can divide the workload among his administrators.

10.3.3 Privileges and Roles

User privileges provide a basic level of security in Enterprise Manager. They are designed to control user access to data and to limit the kinds of SQL statements that users can execute. When creating a user, you grant privileges to enable the user to connect to the database, to run queries and make updates, to create schema objects, and more.

When Enterprise Manager is installed, the SYSMAN user (super administrator) is created by default. The SYSMAN Super Administrator then creates other administrator accounts for daily administration work. The SYSMAN account should only be used to perform infrequent system wide, global configuration tasks.

The Super Administrator divides workload among his administrators by filtering target access, or filtering access to management task, or both through the roles, System Privileges, and Target Privileges he grants them. For example, he can allow some administrators to view any target and to add any target in the enterprise and other administrators to only perform specific operations such as maintaining and cloning on a target for which they are responsible.

A role is a collection of Enterprise Manager resource privileges, or target privileges, or both, which you can grant to administrators or to other roles. These roles can be based

upon geographic location (for example, a role for Canadian administrators to manage Canadian systems), line of business (for example, a role for administrators of the human resource systems or the sales systems), or any other model. Administrators do not want to perform the task of individually granting access to tens, hundreds, or even thousands of targets to every new member of their group.

By creating roles, an administrator needs only to assign the role that includes all the appropriate privileges to his team members instead of having to grant many individual privileges. He can divide workload among his administrators by filtering target access, or filtering access to management task, or both.

Out-of-Box Roles: Enterprise Manager Cloud Control 12c comes with predefined roles to manage a wide variety of resource and target types. The following table lists these roles along with their function.

Table 10–1 Out-of-the-Box Roles

Roles	Description
EM_ALL_ ADMINISTRATOR	Role has privileges to perform Enterprise Manager administrative operations. It provides Full privileges on all secure resources (including targets)
EM_ALL_DESIGNER	Role has privileges to design Enterprise Manager operational entities such as Monitoring Templates.
EM_ALL_OPERATOR	Role has privileges to manage Enterprise Manager operations.
EM_ALL_VIEWER	Role has privileges to view Enterprise Manager operations.
EM_CBA_ADMIN	Role to manage Chargeback Objects. It gives the capability to create and view chargeback plans, chargeback consumers, assign chargeback usage, and view any CaT targets.
EM_CLOUD_ ADMINISTRATOR	Enterprise Manager user for setting up and managing the infrastructure cloud. This role could be responsible for deploying the cloud infrastructure (servers, pools, zones) and infrastructure cloud operations for performance and configuration management.
EM_COMPLIANCE_ DESIGNER	Role has privileges for create, modify and delete compliance entities.
EM_COMPLIANCE_ OFFICER	Role has privileges to view compliance framework definition and results.
EM_CPA_ADMIN	Role to manage Consolidation Objects. It gives the capability to create and view consolidation plans, consolidation projects and view any CaT targets.
EM_HOST_DISCOVERY_ OPERATOR	Role has privileges to execute host discovery
EM_INFRASTRUCTURE_ ADMIN	Role has privileges to manage the Enterprise Manager infrastructure such as managing plugin lifecycle or managing self update.
EM_PATCH_ ADMINISTRATOR	Role for creating, editing, deploying, deleting and granting privileges for any patch plan.
EM_PATCH_DESIGNER	Role for creating and viewing for any patch plan
EM_PATCH_OPERATOR	Role for deploying patch plans
EM_PLUGIN_AGENT_ ADMIN	Role to support plug-in lifecycle on Management Agent
EM_PLUGIN_OMS_ ADMIN	Role to support plug-in lifecycle on Management Server

Table 10–1 (Cont.) Out-of-the-Box Roles

Roles	Description
EM_PLUGIN_USER	Role to support view plug-in console
EM_PROVISIONING_ DESIGNER	Role has privileges for provisioning designer
EM_PROVISIONING_ OPERATOR	Role has privileges for provisioning operator
EM_SSA_ ADMINISTRATOR	EM user with privilege to set up the Self Service Portal. This role can define quotas and constraints for self service users and grant them access privileges.
EM_SSA_USER	This role grants EM user the privilege to access the Self Service Portal.
EM_TARGET_ DISCOVERY_OPERATOR	Role has privileges to execute target discovery.
EM_TC_DESIGNER	Role has privileges for creating Template Collections
EM_USER	Role has privilege to access Enterprise Manager Application.
PUBLIC	PUBLIC role is granted to all administrators. This role can be customized at site level to group privileges that need to be granted to all administrators.

Public Role: Enterprise Manager creates one role by default called **Public**. This role is unique in that it is automatically assigned to all new non-super administrators when they are created. By default it has no privileges assigned to it. The Public role should be used to define default privileges you expect to assign to a majority of non-super administrators you create. Privileges need not be assigned to Public initially - they can be added at any time. The role may be deleted if your enterprise does not wish to use it. If deleted, it can be added back in later if you later decide to implement it.

10.3.3.1 Granting Privileges

A privilege is a right to perform management actions within Enterprise Manager. Privileges can be divided into two categories:

- Target Privileges
- Resource Privileges

Target Privileges: These privileges allow an administrator to perform operations on a target. The Target Privileges page shows a list of targets for which privileges can be granted. Select the check box to specify the privileges that are to be granted and click **Next**.

Table 10–2 Target Privileges Applicable to All Targets

Privilege Name	Privilege Display Name	Description
FULL_ANY_TARGET	Full any Target	Ability to do all operations on all the targets, including delete the target
PERFORM_OPERATION_ AS_ANY_AGENT	Execute Command as any Agent	Execute any OS Command as the Agent User at any Agent
PUT_FILE_AS_ANY_ AGENT	Put File as any Agent	Put any File to any Agent's Filesystem as the Agent User

Table 10–2 (Cont.) Target Privileges Applicable to All Targets

Privilege Name	Privilege Display Name	Description
PERFORM_OPERATION_ ANYWHERE	Execute Command Anywhere	Execute any OS Command at any Agent
OPERATOR_ANY_TARGET	Operator any Target	Privilege to grant operator access on all targets
CONNECT_ANY_VIEW_ TARGET	Connect to any viewable target	Ability to connect and manage any of the viewable target
USE_ANY_BEACON	Use any beacon	Ability to register with any Beacon
EM_MONITOR	EM Monitor	Ability to view any EM Repository targets
VIEW_ANY_TARGET	View any Target	Ability to view any target
GRANT_VIEW_ORACLE_ VM_MANAGER	Grant View Oracle VM Manager Privilege	Ability to grant View Oracle VM Manager privilege
GRANT_VIEW_ORACLE_ VM_ZONE	Grant View Zone Privilege	Ability to grant View Zone privilege
GRANT_VIEW_ORACLE_ CLOUD_ZONE	Grant View Database Zone Privilege	Ability to grant view privilege on Database Zone targets
CREATE_PROPAGATING_ GROUP	Create Privilege Propagating Group	Ability to create privilege propagating groups.Privileges granted on a privilege propagating group will be automatically granted on the members of the group
CREATE_TARGET	Create Target	Ability to create a target

Table 10–3 Target Privileges Applicable to Specific Targets

Privilege Name	Privilege Display Name	Description
GROUP_ ADMINISTRATION	Group Administration	Ability to administor groups
FULL_TARGET	Full Target	Ability to do all operations on the target, including delete the target
FMW_DEPLOY_APP_ TARGET	Deploy Fusion Middleware	Ability to deploy Fusion Middleware components
CONNECT_READONLY_ TARGET	Connect Target Readonly	Ability to connect to target in readonly mode
CONNECT_TARGET	Connect Target	Ability to connect and manage target
MANAGE_TARGET_ COMPLIANCE	Manage Target Compliance	Ability to manage compliance of the target
PERFORM_OPERATION_ AS_AGENT	Execute Command as Agent	Execute any OS Command as the Agent User
PUT_FILE_AS_AGENT	Put File as Agent	Put any File to the Agent's Filesystem as the Agent User

Table 10–3 (Cont.) Target Privileges Applicable to Specific Targets

Privilege Name	Privilege Display Name	Description
MANAGE_TARGET_ ALERTS	Manage Target Events	Ability to clear events, re-evaluate metric alert events, create incidents, add events to incidents, and define what actions the administrator can perform on individual incidents, such as acknowledgment or escalation.
PERFORM_OPERATION	Execute Command	Execute any OS Command
CONFIGURE_TARGET	Configure target	Ability to edit target properties and modify monitoring configuration
MANAGE_TARGET_ PATCH	Manage Target Patch	Privilege to Analyze, Apply and Rollback patches on the target
MANAGE_TC_ OPERATION	Manage Template Collection Operations	Ability to associate a template collection to a adminimistration group and Sync targets with the associated template collections.
MANAGE_TARGET_ METRICS	Manage Target Metrics	Ability to edit threshold for metric and policy setting, apply monitoring templates, and manage User Defined Metrics
BLACKOUT_TARGET	Blackout Target	Ability to create, edit, schedule and stop a blackout on the target
OPERATOR_TARGET	Operator Target	Ability to do normal administrative operations on the target, such as configure a blackout and edit the target properties
FMW_OPERATOR_PRIV	Operator Fusion Middleware	"Ability to perform operations, such as start and shutdown and view logs for Fusion Middleware targets
FMW_PROCESS_ CONTROL_TARGET	Process Control Fusion Middleware	Ability to start or shutdown Fusion Middleware target
FMW_VIEW_LOG_DATA_ TARGET	View Fusion Middleware logs	Ability to view Fusion Middleware diagnostics data
VIEW_ORACLE_CLOUD_ ZONE	View Database Zone	Ability to view Database Zone
VIEW_ORACLE_VM_ MANAGER	View Oracle VM Manager	Ability to view Oracle VM Manager
VIEW_ORACLE_VM_ ZONE	View Oracle VM Zone	Ability to view Oracle VM Zone
VIEW_TARGET	View Target	Ability to view properties, inventory and monitor information about a target

Resource: These privileges allow a user to perform operations against specific types of resources. To set Resource Privileges, from the **Setup** menu, choose **Administrators**. Select an administrator from the list and click **Edit**. The Edit Administrator wizard is displayed. Click **Next** to navigate through the wizard to see the System Privileges page. The following table lists all available resource privileges.

Resource Type	Display Name	Description	Privileges Required to Grant
ACCESS	Access Enterprise Manager	Ability to access Enterprise Manager interfaces	ACCESS
AD4J	JVM Diagnostics User	Gives capability to view the JVM Diagnostic data	SUPER_USER
AD4J	JVM Diagnostics Administrator	Gives capability to manage all JVM Diagnostic Administrative operations	SUPER_USER
ASREPLAY_ ENTITY_MGMT	Application Replay Operator	View, create, and edit any Application Replay entity.	SUPER_USER
ASREPLAY_ ENTITY_MGMT	Application Replay Viewer	View any Application Replay entity.	SUPER_USER
BTM	Request Monitoring User	Gives capability to view the Request Monitoring Data	SUPER_USER
BTM	Request Monitoring Administrator	Gives capability to manage all Request Monitoring Administrative Operations	SUPER_USER
CA	Full Corrective Action	Internal privilege, not for granting	
CA	View Corrective Action	Internal privilege, not for granting	VIEW
CCS_SECURE_ CLASS	Manage custom configurations owned by any user	Ability to create new and edit/delete Custom Configuration specification owned by any user	
CCS_SECURE_ CLASS	Manage custom configurations owned by the user	Ability to create new and edit/delete Custom Configuration specification owned by the user	
CHANGE_PLAN	Manage change plans	Create and delete Change Manager Change Plans	FULL
CHANGE_PLAN	Edit change plan	Edit a Change Manager Change Plan	EDIT
CHANGE_PLAN	View change plan	View a Change Manager Change Plan	VIEW
CHARGEBACK_ AND_ CONSOLIDATION	Manage Chargeback Plans	Ability to Create and Modify Chargeback Plans.	SUPER_USER

Resource Type	Display Name	Description	Privileges Required to Grant
CHARGEBACK_ AND_ CONSOLIDATION	Manage Any Consolidation Plan	Ability to Manage any Consolidation Plans.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	View Chargeback and Consolidation Target	Ability to View Chargeback and Consolidation Target.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	View Any Chargeback and Consolidation Target	Ability to View Any Chargeback and Consolidation Target.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Manage Chargeback and Consolidation Target	Ability to Manage a Chargeback and Consolidation Target.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Manage Any Chargeback and Consolidation Target	Ability to Add/Delete Target to Chargeback and Assign Chargeplan to Target or Add Target to Consolidation Project.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Setup Chargeback and Consolidation	Ability to Setup CAT.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	View Any Chargeback Consumers	Ability to View Any Chargeback Consumers.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Manage Chargeback Consumers	Ability to Create and Modify Chargeback Consumers.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Assign Chargeback Usage	Ability to Assign Chargeback Usage to Consumers.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Assign Chargeback Plan	Ability to Assign Chargeback Plan to CAT Targets.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	View Any Chargeback Plan	Ability to view all the Chargeback Plans.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	View Any Consolidation Plan	Ability to view the Consolidation Plans.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	View Any Consolidation Project	Ability to View any Consolidation Project.	SUPER_USER
CHARGEBACK_ AND_ CONSOLIDATION	Manage Any Consolidation Project	Ability to Manage any Consolidation Project.	SUPER_USER
CLOUDPOLICY	Full Policy	Privilege required to View, Modify, Delete a Policy	FULL
CLOUDPOLICY	Modify Policy	Ability to Modify a Policy	EDIT

Resource Type	Display Name	Description	Privileges Required to Grant
CLOUDPOLICY	View Policy	Ability to View a Policy	VIEW
CLOUDPOLICY	View any Policy	Ability to View any Policy	VIEW
CLOUDPOLICY	Create any Policy	Ability to Create any Policy	CREATE
CLOUDPOLICYGROUP	Full Policy Group	Privilege required to View, Modify, Delete a Policy Group	FULL
CLOUDPOLICYGROUP	Modify Policy Group	Ability to Modify a Policy Group	EDIT
CLOUDPOLICYGROUP	View Policy Group	Ability to View a Policy Group	VIEW
CLOUDPOLICYGROUP	View any Policy Group	Ability to View any Policy Group	VIEW
CLOUDPOLICYGROUP	Create Policy Group	Ability to Create Policy Group	CREATE
COMPLIANCE_FWK	Create Compliance Entity	Ability to create compliance framework, standard, rules	CREATE
COMPLIANCE_FWK	Full any Compliance Entity	Ability to edit/delete compliance framework, standard, rules	FULL
COMPLIANCE_FWK	View any Compliance Framework	Ability to view compliance framework definition and results	VIEW
DISCOVERY	Can Scan Network	Privilege to create, edit and delete host discovery configuration	
DISCOVERY	View Any Discovered Hosts	Privilege to view any discovered hosts	
DISCOVERY	View Any Discovered Targets On Host	Privilege to view any discovered targets on host	
DP	Grant full privilege	Ability to grant upto full privilege on deployment procedures.	GRANT
DP	Grant launch privilege	Ability to grant launch privilege on deployment procedures.	GRANT

Resource Type	Display Name	Description	Privileges Required to Grant
DP	Import	Ability to create deployment procedures and ability to import/export customized deployment procedures.	CREATE
DP	Full	Ability to perform launch, create like, edit structure and delete operations on a Deployment Procedure.	GRANT_FULL_DP
DP	Create	Ability to create deployment procedures.	CREATE
DP	Launch	Ability to perform launch and create like operations on a Deployment Procedure.	GRANT_LAUNCH_DP
EMHA_SECURE_CLASS	Enterprise Manager High Availability Administration	Gives access to manage Enterprise Manager High Availability	ADMIN
EVENT	Manage Events	Manage events privilege object	MANAGE_EVENT
EVENT	View Events	View events privilege object	VIEW
FMW_DIAG_SEC_CLASS	Create Object	Ability to manage the offline diagnostic object lifecycle	SUPER_USER
FMW_DIAG_SEC_CLASS	View object	Ability to view the offline diagnostics objects	SUPER_USER
ISSUE	Manage Problems	Manage problems privilege object	MANAGE_PROBLEM
ISSUE	Manage Incidents	Manage incidents privilege object	MANAGE_INCIDENT
ISSUE	View Issues - (Incidents and Problems)	View issues - Incidents and Problems privilege object	VIEW
JOB	Full	Ability to perform all the valid operations on job, library job, deployment procedure configuration and on deployment procedure instance.	FULL
JOB	Grant view privilege	Ability to grant view privilege on jobs.	GRANT

Resource Type	Display Name	Description	Privileges Required to Grant
JOB	Manage	Ability to perform various operations except edit and delete on job, library job, deployment procedure configuration and on deployment procedure instance.	EDIT
JOB	View	Ability to view, do create like on a job, launch deployment procedure configuration and view deployment procedure instance.	GRANT_VIEW_JOB
JOB	Create	Ability to submit jobs, create library jobs, create deployment procedure instance and create deployment procedure configuration.	CREATE
MEXT_SECURE_CLASS	Full MEXT	Gives complete access to edit, and delete metric extension object	
MEXT_SECURE_CLASS	Edit MEXT	Can edit or create the next version of a metric extension object, but cannot delete it	
MEXT_SECURE_CLASS	Create New Metric Extension	Create or import new metric extensions	
NAMED_CREDENTIALS	Create new Named Credential	Ability to create new named credentials	
NAMED_CREDENTIALS	View Credential	View Credential	
NAMED_CREDENTIALS	Edit Credential	User can update credential but cannot delete it.	
NAMED_CREDENTIALS	Full Credential	Full Credential	
PATCH	Privileges for Patch Setup	Privilege to grant privileges any Patching plan object	
PATCH	Manage privileges on any Patching Plan	Privilege to grant or revoke privileges on any Patching plan object	MANAGE

Resource Type	Display Name	Description	Privileges Required to Grant
PATCH	Full privileges on any Patching Plan	Privilege to view, modify, execute and delete any Patching plan object	FULL
PATCH	Manage privileges on a Patching Plan	Privilege to grant or revoke privileges on a Patching plan object	MANAGE
PATCH	View any Patching Plan	Privilege to view any Patching plan object	VIEW
PATCH	Full Patch Plan	Privilege to view, modify, execute and delete a Patching plan object	MANAGE_PRIV_ ANY_PATCH_PLAN
PATCH	View any Patching Plan Template	Privilege to view any Patching Plan Template object	VIEW
PATCH	Create Patch Plan	Privilege for creating a Patching Plan object	
PATCH	View Patching Plan	Privilege to View a Patching Plan Object	MANAGE_PRIV_ ANY_PATCH_PLAN
PATCH	Create Patch Plan Template	Privilege for creating a Patching Plan Template object	
PLUGIN	Plug-in view privilege	Gives access to manage Enterprise Manager plug-in life cycle console	USER
PLUGIN	Plug-in Agent Administrator	Gives access to manage Enterprise Manager plug-in on Agent	ADMIN
PLUGIN	Plug-in OMS Administrator	Gives access to manage Enterprise Manager plug-in on Management Server	ADMIN
REPORT_DEF	View Report	Ability to view report definition and stored reports, generate on demand reports and do a create like	VIEW
REPORT_DEF	Publish Report	Ability to publish reports for public viewing	
RULESET_SEC	Edit Business Ruleset	Edit Business Ruleset	EDIT
RULESET_SEC	Create Business Ruleset	Create Business Ruleset	CREATE
SBRM_BACKUP_ CONFIG	Create Backup Configuration	Ability to create a backup configuration.	SUPER_USER

Resource Type	Display Name	Description	Privileges Required to Grant
SBRM_BACKUP_CONFIG	Use Backup Configuration	Ability to use a backup configuration.	SUPER_USER
SBRM_BACKUP_CONFIG	Edit Backup Configuration	Ability to edit a backup configuration.	SUPER_USER
SBRM_BACKUP_CONFIG	Full Access	Full access to a backup configuration.	SUPER_USER
SELFUPDATE_SECURE_CLASS	Self Update Administrator	Gives access to manage Enterprise Manager Update	FULL
SELFUPDATE_SECURE_CLASS	View any Enterprise Manager Update	Gives access to view any Enterprise Manager Update	VIEW
SSA	Access Cloud Self Service Portal	Users with this privilege have access to Cloud Self Service Portal.	SUPER_USER
SSA	Setup Cloud Self Service Portal	Privilege to perform Cloud Self Service Portal setup like defining quotas for roles, publishing assemblies etc.	SUPER_USER
SWLIB_ADMINISTRATION	Software Library Storage Administration	Ability to manage upload and reference file storage locations, import and export entities, and purge deleted entities	FULL
SWLIB_ENTITY_MGMT	View any Assembly Entity	View any Assembly Entity	SWLIB_GRANT_ANY_ENTITY_PRIV
SWLIB_ENTITY_MGMT	View any Template Entity	View any Template Entity	SWLIB_GRANT_ANY_ENTITY_PRIV
SWLIB_ENTITY_MGMT	Grant Any Entity Privilege	Ability to grant view, edit and delete privilege on any Software Library entity. This privilege is required if the user granting the privilege on an entity is not a super administrator or owner of the entity.	GRANT
SWLIB_ENTITY_MGMT	Manage Entity	Ability to view, edit and delete a Software Library entity	SWLIB_GRANT_ANY_ENTITY_PRIV
SWLIB_ENTITY_MGMT	View Software Library Entity	Ability to view a Software Library entity	SWLIB_GRANT_ANY_ENTITY_PRIV

Resource Type	Display Name	Description	Privileges Required to Grant
SWLIB_ENTITY_MGMT	Edit an Software Library Entity	Ability to edit a Software Library entity	SWLIB_GRANT_ANY_ENTITY_PRIV
SWLIB_ENTITY_MGMT	Create Any Software Library Entity	Ability to create any Software Library entity	CREATE
SWLIB_ENTITY_MGMT	View Any Software Library Entity	Ability to view any Software Library entity	VIEW
SWLIB_ENTITY_MGMT	Edit Any Software Library Entity	Ability to edit any Software Library entity	EDIT
SWLIB_ENTITY_MGMT	Manage Any Software Library Entity	Ability to create, view, edit and delete any Software Library entity	FULL
SWLIB_ENTITY_MGMT	Import Any Software Library Entity	Ability to import any Software Library entity from a Provisioning Archive (PAR) file	IMPORT
SWLIB_ENTITY_MGMT	Export Any Software Library Entity	Ability to view and export any Software Library entity to a Provisioning Archive (PAR) file	EXPORT
SYSTEM	Super User	Provides all the privileges to any target in the system	
TEMPLATE	View Template	Ability to access a template and apply it to any target on which you have Manage Target Metrics	
TEMPLATE	View Template	Ability to view a template and apply it to any target on which you have Manage Target Metrics	VIEW
TEMPLATE	View any Monitoring Template	View any Monitoring Template.	VIEW
TEMPLATECOLLECTION	Full Template Collection	Ability to edit and delete Template Collection	FULL
TEMPLATECOLLECTION	View Template Collection	Ability to view Template Collection	VIEW
TEMPLATECOLLECTION	View any Template Collection	Ability to view any Template Collection	VIEW
TEMPLATECOLLECTION	Create any Template Collection	Ability to create any Template Collection	CREATE

Select the check box to select the resource privilege to be granted to the administrator and click **Next**.

10.4 Configuring Secure Communication (SSL) for Cloud Control

This section contains the following topics:

- About Enterprise Manager Framework Security
- Enabling Security for the Oracle Management Service
- Securing the Oracle Management Agent
- Enabling Security with Multiple Management Service Installations
- Restricting HTTP Access to the Management Service
- Managing Agent Registration Passwords
- Configuring the OMS with Server Load Balance
- Enabling Security for the Management Repository Database

10.4.1 About Enterprise Manager Framework Security

Enterprise Manager Framework Security provides safe and secure communication channels between the components of Enterprise Manager. For example, Framework Security provides secure connections between your Oracle Management Service and its Management Agents.

> **See Also:** *Oracle Enterprise Manager Concepts* for an overview of Enterprise Manager components

Enterprise Manager Framework Security implements the following types of secure connections between the Enterprise Manager components:

- HTTPS and Public Key Infrastructure (PKI) components, including signed digital certificates, for communications between the Management Service and the Management Agents.

  > **See Also:** *Oracle Security Overview* for an overview of Public Key Infrastructure features, such as digital certificates and public keys

- Oracle Advanced Security for communications between the Management Service and the Management Repository.

  > **See Also:** *Oracle Database Advanced Security Administrator's Guide*

10.4.2 Enabling Security for the Oracle Management Service

To enable Enterprise Manager Framework Security for the Management Service, you use the `emctl secure oms` utility, which is located in the following subdirectory of the Management Service home directory:

```
ORACLE_HOME/bin
```

The `emctl secure oms` utility performs the following actions:

- Generates a Root Key within your Management Repository. The Root Key is used during distribution of Oracle Wallets containing unique digital certificates for your Management Agents.

- Modifies your WebTier to enable an HTTPS channel between your Management Service and Management Agents, independent from any existing HTTPS configuration that may be present in your WebTier.
- Enables your Management Service to accept requests from Management Agents using Enterprise Manager Framework Security.

To run the `emctl secure oms` utility you must first choose an Agent Registration Password. The Agent Registration password is used to validate that future installation of Oracle Management Agents are authorized to load their data into this Enterprise Manager installation.

To enable Enterprise Manager Framework Security for the Oracle Management Service:

1. Stop the Management Service, the WebTier, and the other application server components using the following command:

   ```
   OMS_ORACLE_HOME/bin/emctl stop oms
   ```

2. Enter the following command:

   ```
   OMS_ORACLE_HOME/bin/emctl secure oms
   ```

3. You will be prompted for the Enterprise Manager Root Password. Enter the `SYSMAN` password.
4. You will be prompted for the Agent Registration Password, which is the password required for any Management Agent attempting to secure with the Management Service. Specify an Agent Registration Password for the Management Service.
5. Restart the OMS.
6. After the Management Service restarts, test the secure connection to the Management Service by browsing to the following secure URL using the HTTPS protocol:

   ```
   https://hostname.domain:https_console_port/em
   ```

   Note: The Enterprise Manager console URL can be found by running the "emctl status oms -details" command.

   For example:

   ```
   https://mgmthost1.acme.com:7799/em
   ```

   If the Management Service security has been enabled, your browser displays the Enterprise Manager Login page.

Example 10–1 Sample Output of the emctl secure oms Command

```
emctl secure oms
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Securing OMS... Started.
Securing OMS... Successful
```

Example 10–2 Usage of the emctl secure oms Command (II)

```
emctl secure oms [-sysman_pwd <sysman password>] [-reg_pwd <registration
password>] [-host <hostname>] [-slb_port <slb port>] [-slb_console_port <slb
console port>] [-reset] [-console] [-lock] [-lock_console] [-secure_port <secure_
port>] [-upload_http_port <upload_http_port>] [-root_dc <root_dc>] [-root_country
```

```
<root_country>] [-root_email <root_email>] [-root_state <root_state>] [-root_loc
<root_loc>] [-root_org <root_org>] [-root_unit <root_unit>] [-wallet <wallet_loc>
-trust_certs_loc <certs_loc>] [-key_strength <strength>] [-cert_validity
<validity>] [-protocol <protocol>] [-force_newca] [-ms_hostname <Managed Server
hostname>] [-sign_alg <md5|sha1|sha256|sha384|sha512>]
```

Valid values for <protocol> are the allowed values for Apache's SSLProtocol directive

The parameters are explained below:

- `sysman_pwd` - Oracle Management Repository user password.
- `reg_pwd` - The Management Agent registration password.
- `host` - The host name to be used in the certificate used by the Oracle Management Service. You may need to use the SLB host name if there is an SLB before the Management Service.
- `reset` - A new certificate authority will be created. All the Agents and Oracle Management Services need to be resecured.
- `secure_port` - Specify this to change HTTPS Upload port on WebTier
- `upload_http_port` - Specify this to change HTTP Upload port on WebTier
- `slb_port` - This parameter is required when Server Load Balancer is used. It specifies the secure upload port configured in the Server Load Balancer.
- `slb_console_port` - This parameter is required when Server Load Balancer is used. It specifies the secure console port configured in the Server Load Balancer.
- `root_dc` - The domain component used in the root certificate. The default value is com.
- `root_country` - The country to be used in the root certificate. The default value is US.
- `root_state` - The state to be used in the root certificate. The default value is CA.
- `root_loc` - The location to be used in the root certificate. The default value is EnterpriseManager on <hostname>.
- `root_org` - The organization name to be used in the root certificate. The default value is EnterpriseManager on <hostname>.
- `root_unit` - The organizational unit to be used in the root certificate. The default value is EnterpriseManager on <hostname>.
- `root_email` - The email address to be used in the root certificate. The default value is EnterpriseManager@<hostname>.
- `wallet`: This is the location of the wallet containing the third party certificate. This parameter should be specified while configuring third party certificates.
- `trust_certs_loc` - The location of the `trusted_certs.txt` (required when third party certificates are used).
- `key_strength`: The strength of the key to be used. Valid values are 512, 1024, 2048, and 4096.
- `cert_validity`: The number of days for which the self-signed certificate is valid. The valid range is between 1 to 3650.
- `protocol`: This parameter is used to configure Oracle Management Service in TLSv1-only or SSLv3-only or mixed mode (default). Valid values are the allowed values as per **Apache's SSLProtocol** directive.

Note: The `key_strength` and `cert_validity` parameters are applicable only when the `-wallet` option is not used.

- `force_newca` - If specified, any Agents that are still configured with an older Certificate Authority are ignored.
- `ms_hostname` - Managed Server's hostname.
- `sign_alg` - Signature algorithm.
- `lock`: Locks the Upload
- `lock_console`: Locks the Console
- `console`: If specified, certificate is re-created for HTTPS console port as well

10.4.2.1 Creating a New Certificate Authority

You may need to create a new Certificate Authority (CA) if the current CA is expiring or if you want to change the key strength. A unique identifier is assigned to each CA. For instance, the CA created during installation may have an identifier as ID 1, subsequent CAs will have the IDs 2,3, and so on. At any given time, the last created CA is active and issues certificates for OMSs and Agents.

Example 10–3 Creating a New Certificate Authority

```
emctl secure createca [-sysman_pwd <pwd>] [-host <hostname>] [-key_
strength<strength>] [-cert_validity <validity>] [-root_dc <root_dc>] [-root_
country <root_country>] [-root_email <root_email>] [-root_state <root_state>]
[-root_loc <root_loc>] [-root_org <root_org>] [-root_unit <root_unit>]
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Creating CA... Started.
Successfully created CA with ID 2
```

Example 10–4 Viewing Information about a Certificate Authority

```
emcli get_ca_info -ca_id="1;2" -details
Info about CA with ID: 1
CA is not configured
DN: CN=myhost.example.com, C=US
Serial# : 3423643907115516586
Valid From: Tue Mar 16 11:06:20 PDT 2011
Valid Till: Sat Mar 14 11:06:20 PDT 2020
Number of Agents registered with CA ID 1 is 1
myhost.mydomain.com:3872

Info about CA with ID: 2
CA is configured
DN: CN=myhost.example.com, C=US, ST=CA
Serial# : 1182646629511862286
Valid From: Fri Mar 19 05:17:15 PDT 2011
Valid Till: Tue Mar 17 05:17:15 PDT 2020
There are no Agents registered with CA ID 2
```

The WebLogic Administrator and Node Manager passwords are stored in the Administration Credentials Wallet. This is present in the `EM_INSTANCE_HOME/sysman/config/adminCredsWallet` directory. To recreate Administrator

Credentials wallet, run the following command on each machine on which the Management Service is running:

```
emctl secure create_admin_creds_wallet [-admin_pwd <pwd>]
[-nodemgr_pwd <pwd>]
```

10.4.2.2 Viewing the Security Status and OMS Port Information

To view the security status and OMS port information, use the following command

Example 10–5 emctl status oms -details

```
> emctl status oms -details
Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.
Enter Enterprise Manager Root (SYSMAN) Password : ******
Console Server Host : omshost1.example.com
HTTP Console Port : 7802
HTTPS Console Port : 5416
HTTP Upload Port : 7654
HTTPS Upload Port : 4473
OMS is not configured with SLB or virtual hostname
Agent Upload is locked.
OMS Console is locked.
Active CA ID: 1
Console URL: https://omshost1.example.com:5416/em
Upload URL: https://omshost1.example.com:4473/empbs/upload

WLS Domain Information
Domain Name : EMGC_DOMAIN
Admin Server Host: omshost1.example.com

Managed Server Information
Managed Server Instance Name: EMGC_OMS1
Managed Server Instance Host: omshost1.example.com
```

10.4.2.3 Configuring Transparent Layer Security

The Oracle Management Service can be configured in the following modes:

- **TLSv1-only mode**: To configure the OMS to use only TLSv1 connections, do the following:
  1. Stop the OMS by entering the following command:

     ```
     OMS_ORACLE_HOME/bin/emctl stop oms
     ```

  2. Enter the following command:

     ```
     emctl secure oms -protocol TLSv1
     ```

  3. Append `-Dweblogic.security.SSL.protocolVersion=TLS1` to `JAVA_OPTIONS` in `Domain_Home/bin/startEMServer.sh/cmd`. If this property already exists, update the value to TLS1.
  4. Restart the OMS with the following command:

     ```
     OMS_ORACLE_HOME/bin/emctl start oms
     ```

- **SSLv3 Only Mode**: To configure the OMS to use SSLv3 connections only, do the following:

1. Stop the OMS by entering the following command:

   ```
   OMS_ORACLE_HOME/bin/emctl stop oms
   ```

2. Enter the following command:

   ```
   emctl secure oms -protocol SSLv3
   ```

3. Append `-Dweblogic.security.SSL.protocolVersion=SSL3` to `JAVA_OPTIONS` in `Domain_Home/bin/startEMServer.sh` or `startEMServer.cmd` on Windows. If this property already exists, update the value to SSL3.

4. Restart the OMS with the following command:

   ```
   OMS_ORACLE_HOME/bin/emctl start oms
   ```

- **Mixed Mode**: To configure the OMS to use both SSLv3 and TLSv1 connections, do the following:

  1. Stop the OMS by entering the following command:

     ```
     OMS_ORACLE_HOME/bin/emctl stop oms
     ```

  2. Enter the following command:

     ```
     emctl secure oms
     ```

  3. Append `-Dweblogic.security.SSL.protocolVersion=ALL` to `JAVA_OPTIONS` in `Domain_Home/bin/startEMServer.sh`. If this property already exists, update the value to `ALL`.

  4. Restart the OMS with the following command:

     ```
     OMS_ORACLE_HOME/bin/emctl start oms
     ```

Note: By default, the OMS is configured to use the Mixed Mode. To configure the Management Agent in TLSv1 only mode, set `allowTLSOnly=true` in the `emd.properties` file and restart the Agent.

10.4.3 Securing the Oracle Management Agent

When you install the Management Agent on a host, you must identify the Management Service that will be used by the Management Agent. To enable Enterprise Manager Framework Security for the Management Agent, use the `emctl secure agent` utility, which is located in the following directory of the Management Agent home directory:

AGENT_HOME/bin (UNIX)
AGENT_HOME\bin (Windows)

The `emctl secure agent` utility performs the following actions:

- Obtains an Oracle Wallet from the Management Service that contains a unique digital certificate for the Management Agent. This certificate is required in order for the Management Agent to conduct SSL communication with the secure Management Service.
- Obtains an Agent Key for the Management Agent that is registered with the Management Service.

- Configures the Management Agent so it is available on your network over HTTPS and so it uses the Management Service HTTPS upload URL for all its communication with the Management Service.

To enable Enterprise Manager Framework Security for the Management Agent:

1. Ensure that your Management Service and the Management Repository are up and running.
2. Change directory to the following directory:

   ```
   AGENT_HOME/bin (UNIX)
   AGENT_HOME\bin (Windows)
   ```

3. Stop the Management Agent:

   ```
   emctl stop agent
   ```

4. Enter the following command:

   ```
   emctl secure agent (UNIX)
   emctl secure agent (Windows)
   ```

   The `emctl secure agent` utility prompts you for the Agent Registration Password, authenticates the password against the Management Service, and reconfigures the Management Agent to use Enterprise Manager Framework Security.

   Example 10–6 shows sample output of the `emctl secure agent` utility.

5. Restart the Management Agent:

   ```
   emctl start agent
   ```

6. Confirm that the Management Agent is secure by checking the Management Agent home page.

   Note: You can also check if the Agent Management is secure by running the `emctl status agent -secure` command, or by checking the Agent and Repository URLs in the output of the `emctl status agent` command.

   In the Management Agent home page , the **Secure Upload** field indicates whether or not Enterprise Manager Framework Security has been enabled for the Management Agent.

Example 10–6 Sample Output of the emctl secure agent Utility

```
emctl secure agent
Oracle Enterprise Manager 12c Release 1 Cloud Control.
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Securing agent...   Started
Securing agent...   Successful.
```

Example 10–7 Sample Output of the emctl status agent secure Command

```
emctl status agent -secure
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Checking the security status of the Agent at location set in
/private/home/oracle/product/102/em/agent10g/sysman/config/emd.properties...
```

```
Done.
Agent is secure at HTTPS Port 3872.
Checking the security status of the OMS at
https://cloudcontrol.oraclecorp.com:4889/em/upload/...  Done.
OMS is secure on HTTPS Port 4888
```

10.4.4 Enabling Security with Multiple Management Service Installations

Because you have already established at least one Agent Registration Password and a Root Key in your Management Repository, they must be used for your new Management Service. Your secure Management Agents can then operate against either Management Service.

All the registration passwords assigned to the current Management Repository are listed on the Registration Passwords page in the Oracle Enterprise Manager 12c Cloud Control Console.

If you install a new Management Service that uses a new Management Repository, the new Management Service is considered to be a distinct enterprise. There is no way for the new Management Service to partake in the same security trust relationship as another Management Service that uses a different Management Repository. Secure Management Agents of one Management Service will not be able to operate against the other Management Service.

10.4.5 Restricting HTTP Access to the Management Service

Note: The Oracle Management Service is locked (both console & upload) by default begining with Enterprise Manager 12c.

It is important that only secure Management Agent installations that use the Management Service HTTPS channel are able to upload data to your Management Repository and Cloud Control console is accessible via HTTPS only.

To restrict access so Management Agents can upload data to the Management Service only over HTTPS:

1. Stop the Management Service, the WebTier, and the other application server components:

   ```
   cd ORACLE_HOME/opmn/bin
   emctl stop oms
   ```

2. Change directory to the following location in the Management Service home:

   ```
   ORACLE_HOME/bin
   ```

3. Enter the following command to prevent Management Agents from uploading data to the Management Service over HTTP:

   ```
   emctl secure lock -upload
   ```

   To lock the console and prevent HTTP access to the console, enter the following command:

   ```
   emctl secure lock -console
   ```

   To lock both, enter either of the following commands:

   ```
   emctl secure lock or
   emctl secure lock -upload -console
   ```

To lock both the console access and uploads from Agents while enabling security on the Management Service, enter the following command:

```
emctl secure oms -lock [other options]
```

4. Restart the Management Service, the WebTier, and the other application server components:

```
cd ORACLE_HOME/bin
emctl start oms
```

5. Verify that you cannot access the OMS upload URL using the HTTP protocol:

   For example, navigate to the following URL:

```
http://hostname.domain:4889/empbs/upload
```

   You should receive an error message similar to the following:

```
Forbidden
You are not authorised to access this resource on the server.
```

6. Verify that you can access the Management Agent Upload URL using the HTTPS protocol:

   For example, navigate to the following URL:

```
https://hostname.domain:4888/empbs/upload
```

   You should receive the following message, which confirms the secure upload port is available to secure Management Agents:

```
Http XML File receiver
Http Recceiver Servlet active!
```

To allow the Management Service to accept uploads from unsecure Management Agents, use the following command:

```
emctl secure unlock -upload
```

Note:

- The OMS need to be stopped before running 'secure unlock', and then restarted afterwards.
- To unlock the console and allow HTTP access to the console, enter the following command:

```
emctl secure unlock -console
```

- To unlock both, enter either of the following command:

```
emctl secure unlock
emctl secure unlock -console -upload
```

Example 10–8 Sample Output of the emctl secure lock Command

```
emctl secure lock
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
OMS Console is locked. Access the console over HTTPS ports.
Agent Upload is locked. Agents must be secure and upload over HTTPS port.
```

```
Restart OMS
```

Example 10–9 Sample Output of the emctl secure unlock Command

```
emctl secure unlock
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
OMS Console is unlocked. HTTP ports too can be used to access console.
Agent Upload is unlocked. Unsecure Agents may upload over HTTP.
Restart OMS
```

10.4.6 Managing Agent Registration Passwords

Enterprise Manager uses the Agent Registration password to validate that installations of Oracle Management Agents are authorized to load their data into the Oracle Management Service.

The Agent Registration password is created during installation when security is enabled for the Oracle Management Service. You can add/edit/delete registration passwords directly from the Enterprise Manager console.

Note: If you want to avoid new Agents from being registered with the OMS, delete all registration passwords.'

10.4.6.1 Using the Cloud Control Console to Manage Agent Registration Passwords

You can use the Cloud Control Console to manage your existing registration passwords or create additional registration passwords:

1. From the **Setup** menu, select **Security**, then select **Registration Passwords**.
2. Enterprise Manager displays the Registration Passwords page (Figure 10–3). The registration password you created when you ran the `emctl secure oms` command appears in the Registration Passwords table.
3. Use the Registration Passwords page to change your registration password, create additional registration passwords, or remove registration passwords associated with the current Management Repository.

Figure 10–3 Managing Registration Passwords in the Cloud Control Console

When you create or edit an Agent Registration Password on the Registration Passwords page, you can determine whether the password is persistent and available for multiple Management Agents or to be used only once or for a predefined period of time.

For example, if an administrator requests to install a Management Agent on a particular host, you can create a one-time-only password that the administrator can use to install and configure one Management Agent.

On the other hand, you can create a persistent password that an administrator can use for the next two weeks before it expires and the administrator must ask for a new password.

10.4.6.2 Using emctl to Add a New Agent Registration Password

To add a new Agent Registration Password, use the following `emctl` command on the machine on which the Management Service has been installed:

```
emctl secure setpwd [sysman pwd] [new registration pwd]
```

The `emctl secure setpwd` command requires that you provide the password of the Enterprise Manager super administrator user, `sysman`, to authorize the addition of the Agent Registration Password.

If you change the Agent Registration Password, you must communicate the new password to other Enterprise Manager administrators who need to install new Management Agents, enable Enterprise Manager Framework Security for existing Management Agents, or install additional Management Services.

As with other security passwords, you should change the Agent Registration Password on a regular and frequent basis to prevent it from becoming too widespread.

10.4.7 Configuring the OMS with Server Load Balance

When you deploy a Management Service that is available behind a Server Load Balancer (SLB), special attention must be given to the DNS host name over which the Management Service will be available. Although the Management Service may run on a particular local host, for example `myhost.mycompany.com`, your Management Agents will access the Management Service using the host name that has been assigned to the Server Load Balancer. For example, `oracleoms.mycompany.com`.

As a result, when you enable Enterprise Manager Framework Security for the Management Service, it is important to ensure that the Server Load Balancer host name is embedded into the Certificate that the Management Service uses for SSL communications. To do so, enter the following commands:

This may be done by using `emctl secure oms` and specifying the host name in the with an extra `-host` parameter as follows:

- Enable security on the Management Service by entering the following command:

  ```
  emctl secure oms -host <slb_hostname> [-slb_console_port <slb
  UI port>] [-slb_port <slb upload port>] [other params]
  ```

  Run this command on each OMS. You will need to restart each OMS after running the 'emctl secure oms' command.

- Create virtual servers and pools on the Server Load Balancer.
- Verify that the console can be accessed using the following URL:

  ```
  https://slbhost:slb_console_port/em
  ```

- Re-secure the Agents with Server Load Balancer by using the following command:

  ```
  emctl secure agent -emdWalletSrcUrl <SLB Upload or UI URL>
  ```

  For example:

  ```
  Agent_Home/bin/emctl secure agent -emdWalletSrcUrl
  https://slbost:slb_upload_port/em
  ```

10.4.8 Enabling Security for the Management Repository Database

This section describes how to enable Security for the Oracle Management Repository. This section includes the following topics:

- About Oracle Advanced Security and the sqlnet.ora Configuration File
- Configuring the Management Service to Connect to a Secure Management Repository Database
- Enabling Oracle Advanced Security for the Management Repository
- Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database

10.4.8.1 About Oracle Advanced Security and the sqlnet.ora Configuration File

You enable security for the Management Repository by using Oracle Advanced Security. Oracle Advanced Security ensures the security of data transferred to and from an Oracle database.

> **See Also:** *Oracle Database Advanced Security Administrator's Guide*

To enable Oracle Advanced Security for the Management Repository database, you must make modifications to the `sqlnet.ora` configuration file. The `sqlnet.ora` configuration file is used to define various database connection properties, including Oracle Advanced Security parameters.

The sqlnet.ora file is located in the following subdirectory of the Database home:

ORACLE_HOME`/network/admin`

After you have enabled Security for the Management Repository and the Management Services that communicate with the Management Repository, you must also configure Oracle Advanced Security for the Management Agent by modifying the `sqlnet.ora` configuration file in the Management Agent home directory.

> **See Also:** "Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database"

It is important that both the Management Service and the Management Repository are configured to use Oracle Advanced Security. Otherwise, errors will occur when the Management Service attempts to connect to the Management Repository. For example, the Management Service might receive the following error:

```
ORA-12645: Parameter does not exist
```

To correct this problem, be sure both the Management Service and the Management Repository are configured as described in the following sections.

Note: The procedures in this section describe how to manually modify the sqlnet.ora configuration file to enable Oracle Advanced Security. Alternatively, you can make these modifications using the administration tools described in the *Oracle Database Advanced Security Administrator's Guide*.

10.4.8.2 Configuring the Management Service to Connect to a Secure Management Repository Database

If you have enabled Oracle Advanced Security for the Management Service database—or if you plan to enable Oracle Advanced Security for the Management Repository database—use the following procedure to enable Oracle Advanced Security for the Management Service:

1. Stop the Management Service:

   ```
   ORACLE_HOME/bin/emctl stop oms
   ```

2. Set Enterprise Manager operational properties by using the `emctl set property` command.
3. Restart the Management Service.

   ```
   ORACLE_HOME/bin/emctl start oms
   ```

Table 10–4 Oracle Advanced Security Properties in the Enterprise Manager Properties File

Property	Description
oracle.sysman.emRep.dbConn.enableEncryption	Defines whether or not Enterprise Manager will use encryption between Management Service and Management Repository. Possible values are TRUE and FALSE. The default value is TRUE. For example: `oracle.sysman.emRep.dbConn. enableEncryption=true`
oracle.net.encryption_client	Defines the Management Service encryption requirement. Possible values are REJECTED, ACCEPTED, REQUESTED, and REQUIRED. The default value is REQUESTED. In other words, if the database supports secure connections, then the Management Service uses secure connections, otherwise the Management Service uses insecure connections. For example: `oracle.net. encryption_client=REQUESTED`

Table 10–4 (Cont.) Oracle Advanced Security Properties in the Enterprise Manager Properties File

Property	Description
oracle.net.encryption_types_client	Defines the different types of encryption algorithms the client supports. Possible values should be listed within parenthesis. The default value is `( DES40C )`. For example: `oracle.net.encryption_types_client= ( DES40C )`
oracle.net.crypto_checksum_client	Defines the Client's checksum requirements. Possible values are REJECTED, ACCEPTED, REQUESTED, and REQUIRED. The default value is REQUESTED. In other words, if the server supports checksum enabled connections, then the Management Service uses them, otherwise it uses normal connections. For example: `oracle.net.crypto_checksum_client=REQUESTED`
oracle.net.crypto_checksum_types_client	This property defines the different types of checksums algorithms the client supports. Possible values should be listed within parentheses. The default value is (MD5). For example: `oracle.net.crypto_checksum_types_client= ( MD5 )`

10.4.8.3 Enabling Oracle Advanced Security for the Management Repository

To be sure your database is secure and that only encrypted data is transferred between your database server and other sources, review the security documentation available in the Oracle Database documentation library.

See Also: *Oracle Database Advanced Security Administrator's Guide*

The following instructions provide an example of how you can confirm that Oracle Advanced Security is enabled for your Management Repository database and its connections with the Management Service:

1. Locate the `sqlnet.ora` configuration file in the following directory of the database Oracle Home:

   ```
   ORACLE_HOME/network/admin
   ```

2. Using a text editor, look for the following entries (or similar entries) in the `sqlnet.ora` file:

   ```
   SQLNET.ENCRYPTION_SERVER = REQUESTED
   SQLNET.CRYPTO_SEED = "abcdefg123456789"
   ```

> **See Also:** "Configuring Network Data Encryption and Integrity for Oracle Servers and Clients in the *Oracle Application Server 10g Administrator's Guide.*

3. Save your changes and exit the text editor.

10.4.8.4 Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database

After you have enabled Oracle Advanced Security for the Management Repository, you must also enable Advanced Security for the Management Agent that is monitoring the Management Repository:

1. Locate the `sqlnet.ora` configuration file in the following directory inside the home directory for the Management Agent that is monitoring the Management Repository:

   AGENT_HOME`/network/admin (UNIX)`
   AGENT_HOME`\network\admin (Windows)`

2. Using a text editor, add the following entry to the `sqlnet.ora` configuration file:

   ```
   SQLNET.CRYPTO_SEED = "abcdefg123456789"
   ```

   The `SQLNET.CRYPTO_SEED` can be any string between 10 to 70 characters.

   > **See Also:** "Configuring Network Data Encryption and Integrity for Oracle Servers and Clients in the *Oracle Application Server Administrator's Guide.*

3. Save your changes and exit the text editor.
4. Restart the Management Agent.

10.4.9 Configuring Third Party Certificates

You can configure third party certificates for:

- HTTPS Upload Virtual Host
- HTTPS Console Virtual Host

> **Note:** Only Single Sign-On wallets are supported.

10.4.9.1 Configuring Third Party Certificate for HTTPS Upload Virtual Host

You can configure the third party certificate for the HTTPS Upload Virtual Host in two ways:

Method I

1. Create a wallet for each OMS in the Cloud.
2. While creating the wallet, specify the host name of the machine where the OMS is installed or the Load Balancer Name if the OMS is behind the Load Balancer for Common Name.
3. Write the certificates of all the Certificate Authorities in the certificate chain (like the Root Certificate Authority, Intermediate Certificate Authority) into a file named `trusted_certs.txt`.

4. Download or copy the `trusted_certs.txt` file to the host machines on which each Agent that is communicating with the OMS is running.
5. Run the add_trust_cert command on each Agent and then restart that Agent.

```
emctl secure add_trust_cert -trust_certs_loc <location of the trusted_certs.txt
file>
```

6. Secure the OMS and restart it.

```
emctl secure oms -wallet <location of wallet> -trust_certs_loc <loc of trusted_
certs.txt> [any other options]
```

Method 2

1. Create a wallet for each OMS in the Cloud.
2. Specify the host name of the machine where the OMS is installed or the Load Balancer Name if the OMS is behind the Server Load Balancer for Common Name (CN).
3. Write the certificates of all the Certificate Authorities in the certificate chain (like the Root Certificate Authority, Intermediate Certificate Authority) into a file named `trusted_certs.txt`.
4. Restart the OMS after it has been secured.

```
emctl secure oms -wallet <location of wallet> -trust_certs_loc <loc of trusted_
certs.txt> [any other options]
```

5. Either re-secure the Agent by running the `emctl secure agent` command (should be run on all Agents) or import the trust points by running the `emctl secure add_trust_cert -trust_certs_loc <location of the trusted_certs.txt file>` command. The `-trust_certs_loc` parameter must contain the path and the filename of the `trusted_certs.txt` file.

> **Note:** This file must only contain certificates in base64 format and no special characters or empty lines.

10.4.9.2 Configuring Third Party Certificate for HTTPS Console Virtual Host

To configure the third party certificate for HTTPS WebTier Virtual Host:

1. Create a wallet for each OMS in the Cloud. Specify the host name of the machine where the OMS is installed or the Load Balancer Name if the OMS is behind the Server Load Balancer for Common Name.
2. Run the following command on each OMS and the restart that OMS:

```
emctl secure console -wallet <location of wallet>
```

> **Note:** Only single-sign-on wallets are supported.

10.5 Accessing Managed Targets

The following topics are discussed in this section:

- Credential Subsystem
- Pluggable Authentication Modules (PAM) Support

- Sudo and Powerbroker Support

10.5.1 Credential Subsystem

Credentials like user names and passwords are typically required to access targets such as databases, application servers, and hosts. Credentials are encrypted and stored in Enterprise Manager. Beginning with Enterprise Manager 12c, the credential subsystem supports, in addition to basic username-password, strong authentication schemes such as PKI, SSH-keys and Kerberos. SSH-key based host authentication, used by jobs, deployment procedures and other Enterprise Manger subsystems, is now supported.

By using appropriate credentials, you can:

- Collect metrics in the background as well as real-time
- Perform jobs such as backup, patching, and cloning
- Perform real-time target administration such as start, and stop
- Connect to My Oracle Support

Based on their usage, credentials can be classified into the following categories:

- Named Credential
- Job Credentials
- Monitoring Credentials
- Collection Credentials
- Preferred Credentials

10.5.1.1 Named Credential

Credentials are stored within Enterprise Manager as "named" entities. Administrators define and store credentials within Enterprise Manager and refer to the credential by a credential name. Named credentials can be a username/password, or a public key-private key pair. An Enterprise Manager administrator can then use the named credential for performing operations like running jobs, patching and other system management tasks. For example, an administrator can store the username and password they want to use for patching as "MyPatchingCreds". He can later submit a patching job that uses "MyPatchingCreds" to patch a production databases.

There are two categories of named credentials:

- **Global Named Credential**

  A global named credential is an entity, which is not associated with any Enterprise Manager object. Global named credentials consist of the authentication scheme along with any authentication parameters. Because these are independent entities, an Enterprise Manger administrator can associate these credentials with objects at a later time.

- **Target Named Credentials**

  Target named credential is an entity which are associated with individual targets at the time of creation. This entity will also contain authentication scheme along with authentication parameters for a specific target.

Access Control for Named Credentials

The access control model for credentials adhere to the following rules:

- Only credential owners can grant privileges on their credential objects to other users.
- Enterprise Manager Super Administrators cannot obtain any privileges on a newly created credential until he is explicitly granted privileges on the credential object.
- Enterprise Manager administrators, regardless of privilege level, cannot see the sensitive fields such as passwords and private keys from the console UI.
- Credentials privileges cannot be assigned to a role. This eliminates back door entry by Enterprise Manager Super Administrators to grant themselves privileges on the credentials for which they do not have explicit access.
- An Enterprise Manager administrator cannot view other administrators' credentials unless an explicit grant is provided. Even Enterprise Manager Super Administrators cannot view other users' credentials.
- Any Enterprise Manager administrator can create his own credentials and have FULL privileges on the credentials owned.

Enterprise Manager Administrators will be able to grant privileges to other administrators while creating the credential or by granting the privileges when editing the credential.

All the credentials owned by an Enterprise Manager administrator will be deleted if that administrator is deleted from Enterprise Manager. Since access to shared credentials is not automatically granted to Super Administrators, re-assigning named credentials belonging to a regular Enterprise Manager administrator by a Super Administrator is not allowed.

Credential Privilege Levels

The following privilege levels are available for all credentials:

- **VIEW**: An administrator with VIEW privileges on other administrator's credentials will be able to view the structure and username of the credential. Sensitive information of the credential such as the password will never be shown. Administrators with VIEW privilege on a credential will also be able to use the credentials for running jobs, patching and other system management operations within Enterprise Manager.
- **EDIT**: Allows an Enterprise Manager administrator to change a sensitive information such as the password, or the public/private key pair of the credential. The administrator will not be able to change the Authentication Scheme of the credential. The username for the credential cannot be changed.
- **FULL**: Allows an Enterprise Manager administrator to change the credential username, sensitive information such as the password or the public/private key pair, and authentication scheme. An administrator with FULL privilege on a named credential will be able to delete the named credential.

To create or edit a named credential, from the **Setup** menu, choose **Security** and then **Named Credential**. The Named Credential page displays as shown in the following figure.

Figure 10–4 Named Credentials Page

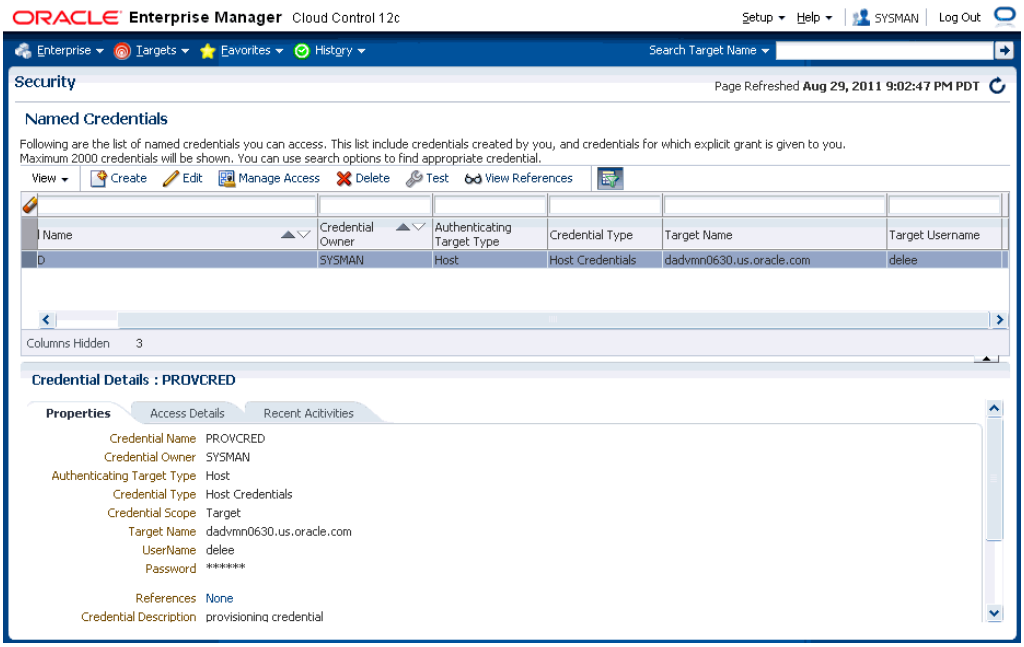

From the Named Credential page, you can **Create** a new named credential, **Edit** an existing credential, **Manage Access** (grant/revoke privileges), **Delete**, **Test**, **View References**, or click the *Query by Example* icon to filter the list of named credentials.

10.5.1.2 Job Credentials

The job system uses the credential subsystem to retrieve the credentials required to submit a job on the targets. The administrator can define their preferred and default credentials from the **Setup** menu, choose **Security** and then **Preferred Credentials** page. As an administrator, you can:

1. Use Preferred Credentials
2. Use Named Credentials
3. Create new credentials

while submitting the job.

> **Note:** If the user chooses to use preferred credentials, these credentials will be used when the user submits the job. If the preferred credentials are not available, the default credentials will be used. If default credentials are not present, the job cannot be submitted.

10.5.1.3 Monitoring Credentials

These credentials are used by the Management Agent to monitor certain types of targets. For example, most database monitoring involves connecting to the database, which requires a username, password, and optionally, a role. Monitoring credentials, if

stored in the repository, can also be potentially used by management applications to connect directly to the target from the OMS.

To create or edit a monitoring credentials, from the **Setup** menu, choose **Security** and then **Monitoring Credentials**. The Monitoring Credentials page displays as shown in the following figure.

Figure 10–5 Monitoring Credentials

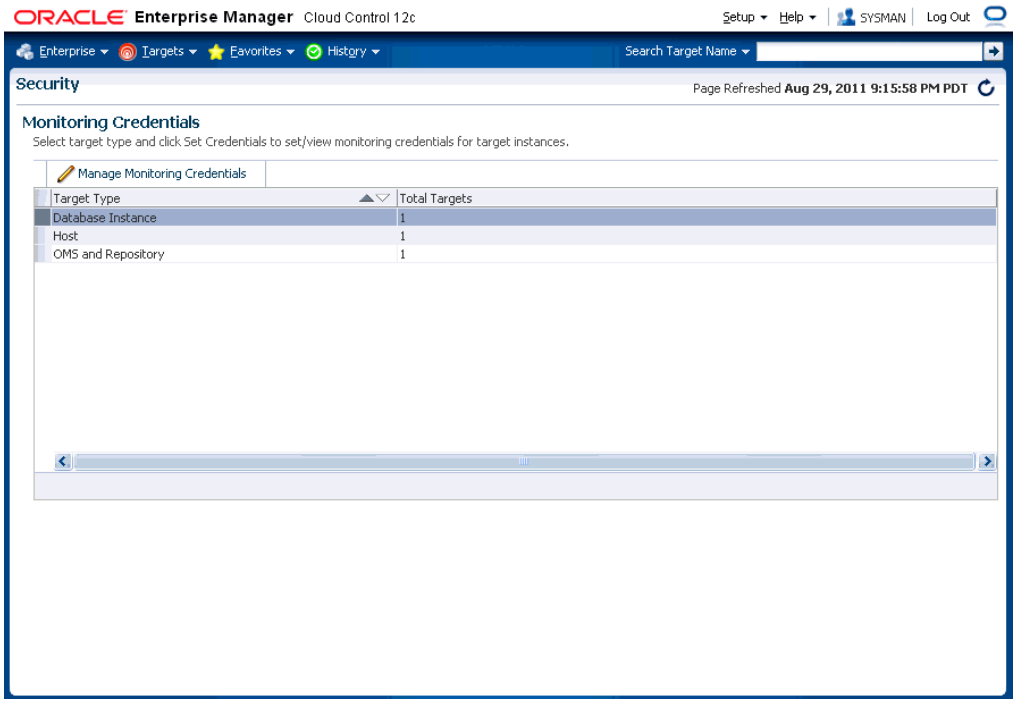

To modify monitoring credentials, select the desired target type and click **Manage Monitoring Credentials**. The monitoring credentials page for the selected target type displays.

10.5.1.4 Collection Credentials

These credentials are associated with metric extensions and older user-defined metrics.

10.5.1.5 Preferred Credentials

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

- **Default Credentials**: Default credentials can be set for a particular target type and will be available for all the targets of the target type. It will be overridden by target preferred credentials.
- **Target Credentials**: Target credentials are preferred credentials set for a particular target. They could be used by applications such as the job system, notifications, or

patching. For example, if the user chooses to use preferred credentials while submitting a job, then the preferred credentials set for the target (target credentials) will be used. If the target credentials are not present, the default credentials (for the target type) will be used. If the default credentials are not present, the job will fail. If not specified, by default, preferred credentials refer to preferred target credentials"

For example, to set the host preferred credentials, from the **Setup** menu, choose **Security** and then **Preferred Credential**. In the Preferred Credentials page, select the **Host** target type from the table and click **Manage Preferred Credentials**. The Host Preferred Credentials are displayed.

Figure 10–6 Host Preferred Credentials

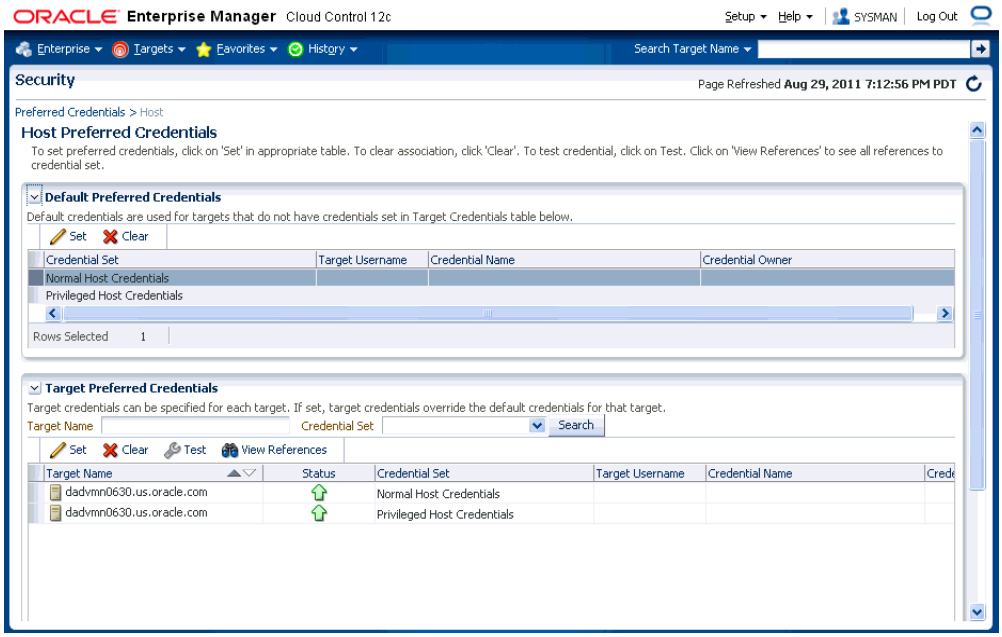

On this page, you can set both default and explicit preferred credentials for the host target types.

10.5.1.6 Managing Credentials Using EM CLI

You can manage passwords using EM CLI verbs. Using EM CLI, you can:

- Change the database user password in both the target database and Enterprise Manager.

  ```
  emcli update_db_password -change_at_target=Yes|No -change_all_reference=Yes|No
  ```

- Update a password which has already been changed at the host target.

  ```
  emcli update_host_password -change_all_reference=Yes|No
  ```

- Set preferred credentials for given users.

  ```
  emcli set_preferred_credential
  -set_name="set_name"
  -target_name="target_name"
  -target_type="ttype"
  -credential_name="cred_name"
  [-credential_owner ="owner]"
  ```

And

```
emcli set_preferred_credential
-set_name="set_name"
-target_name="target_name"
-target_type="ttype"
-credential_name="cred_name"
[-credential_owner ="owner]"
```

For detailed descriptions of these verbs, refer to *Enterprise Manager Command Line Interface* guide.

10.5.2 Setting Up SSH Key-based Host Authentication

Secure Shell or SSH allows data to be exchanged over the network using a secure channel between two devices. SSH is used primarily on Linux and Unix based systems. SSH was designed as a replacement for FTP, telnet and other unsecure remote shells, which send information, notably passwords in plaintext, leaving them open for interception. The encryption used by SSH provides confidentiality and integrity of data over an insecure network. SSH also protects the system against DNS spoofing attacks. This makes SSH a better choice in production environments over telnet/FTP and other username/password based authentications.

You can configure Enterprise Manager to use SSH while performing management operations, thus allowing Enterprise Manager administrators to leverage the security features provided by SSH along with the management capabilities of Enterprise Manager. When authenticating in this mode, the Agent acts as a Java SSH client and connect to the host using the username/password provided in the credential.

Enterprise Manager allows you to store a public-private key pair for administrators and allows them to view and install the public key on the hosts. Administrators can then submit jobs/patching operations in which they specify the credential that refers to the private key to perform the operation. The OMS passes the private key to the Agent along with the commands and the command parameters. Agent invokes the Java SSH client and attempts to connect to the host using the private key. Since the host already has the public key installed, it identifies the private key and successfully authenticates the Agent's Java SSH client. The Agent can now run the commands via the SSH client on the host to perform the requested operations.

Setup Example Session

> **Note:** The procedure shown in this example assumes that you have a firm understanding of SSH setup procedures and user and host equivalence using public private key pair using SSH.

To generate, manage, or convert SSH authentication keys, you use the *SSH-keygen* utility available on UNIX systems. This utility SSH-keygen tool provides different options to create with different strengths RSA keys for SSH protocol version 1 and RSA or DSA keys for use by SSH protocol version 2.

Example 10–10 Setting Up SSH key-based Authentication

```
$ ssh-keygen -t rsa
```

The command options instruct the utility to generate SSH keys (RSA key pair).

```
Generating public/private rsa key pair.
```

```
Enter file in which to save the key (/home/myhome/.ssh/id_rsa):
```
The path specified is the standard path to the location where SSH keys are stored ($HOME/.ssh).

```
Enter passphrase (empty for no passphrase)
Enter same passphrase again:
Your identification has been saved in /home/admin1/.ssh/id_rsa.
Your public key has been saved in /home/admin1/.ssh/id_rsa.pub.
The key fingerprint is:
bb:da:59:7a:fc:24:c6:9a:ee:dd:af:da:1b:1b:ed:7f admin1@myhost2170474
```

The ssh-regkey utility has now generated two files in the .ssh directory.

```
$ ls
id_rsa  id_rsa.pub
```
To permit access to the host without having SSH prompt for a password, copy the public key to the authorized_keys file on that system.

```
$ cp id_rsa.pub  authorized_keys
```

From this point, all keys listed in that file are allowed access.

Next, perform a remote login using SSH. The system will not prompt you for a password.

```
$ ssh myhost
The authenticity of host 'myhost (10.229.147.184)' can't be established.
RSA key fingerprint is de:a0:2a:d5:23:f0:8a:72:98:74:2c:6f:bf:ad:5b:2b.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'myhost,10.229.147.184' (RSA) to the list of known
hosts.
Last login: Mon Aug 29 16:48:45 2011 from anotherhost.example.com
$
```

You are now ready to add the credential to Enterprise Manager.

1. From the **Setup** menu, select **Security**, then select **Named Credentials**.
2. On the Named Credentials page, click Create.
3. Select **Host** from the **Authenticating Target Type** drop-down menu.
4. Select **SSH Key Credentials** from the **Credential Type** drop-down menu as shown in the following figure.

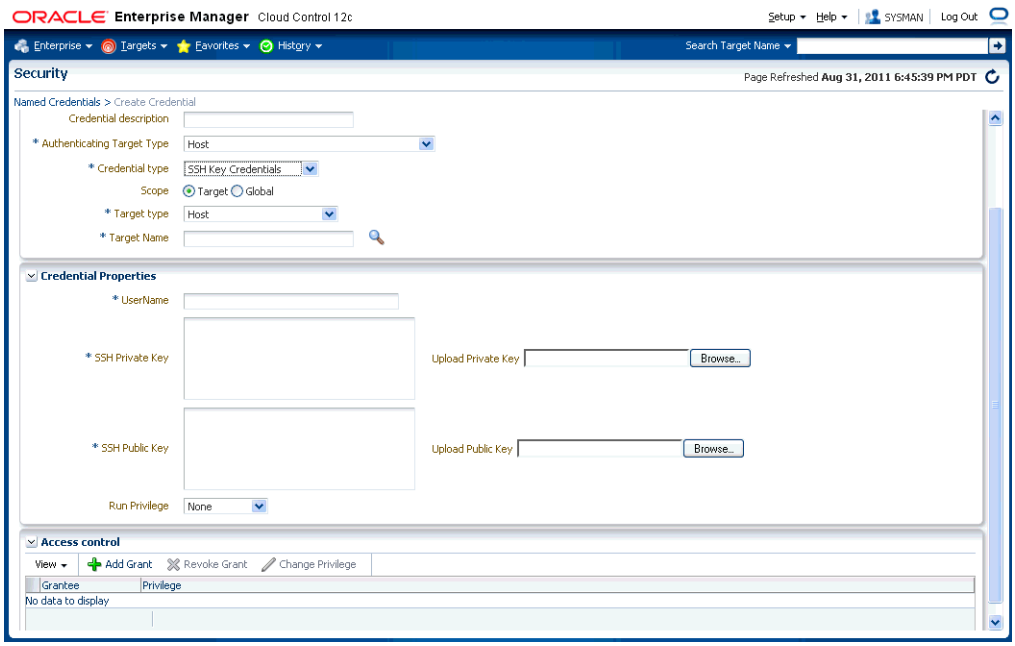

5. Ensure that the SSH private key/public key files have been copied to the host on which the browser is running.
6. From the **Credential Properties** region, click **Browse** for **Public Key** and **Private Key** to upload the generated public key/private key files.
7. Click **Test and Save** to verify the credentials and save them.

10.5.3 Pluggable Authentication Modules (PAM) Support for Hosts

Pluggable authentication modules, or PAM, is a mechanism to integrate multiple low-level authentication schemes into a high-level application programming interface (API). It allows programs that rely on authentication to be written independently of the underlying authentication scheme. By using PAM, instead of using the local password file to authenticate the user accessing the host, you can take advantage of other authentication mechanisms such as LDAP, RADIUS and Kerberors. If your host authentication is configured over PAM, the Management Agent needs to be configured accordingly to enable PAM Authentication. Refer to note 422073.1 for deployment details.

> **Note:** The local password file (usually `/etc/passwd`) will be checked and used first. This should be synchronized with the LDAP password if it is being used. If this fails, the Management Agent will switch to the external authentication module.

10.5.3.1 Configuring PAM for RHEL4 Users

For users on RHEL4, the PAM file configuration is as follows:

```
#%PAM-1.0
auth required pam_ldap.so
account required pam_ldap.so
password required pam_ldap.so
session required pam_ldap.so
```

For more details, see
https://www.redhat.com/docs/manuals/enterprise/RHEL-4-Manual/ref-guide/s1-pam-format.html

10.5.3.2 Configuring PAM for AIX Users

For AIX users, use the `edit/etc/pam.conf` file and add the following lines:

```
emagent auth    required        /usr/lib/security/pam_aix
emagent account required        /usr/lib/security/pam_aix
emagent password  required      /usr/lib/security/pam_aix
emagent session required        /usr/lib/security/pam_aix
```

After editing the file, apply patch **5527130** and run `root.sh`

10.5.4 Sudo and PowerBroker Support

Privilege delegation allows a logged-in user to perform an activity with the privileges of another user. Sudo and PowerBroker are privilege delegation tools that allow a logged-in user to be assigned these privileges. Typically, the privileges that are granted to a specific user are administered centrally. For example, the sudo command can be used to run a script that requires root access:

```
sudo root root.sh
```

In the invocation of sudo in the example above, an administrator can use the sudo command to run a script as root provided he has been granted the appropriate privileges by the system administrator.

Enterprise Manager preferred credentials allow you to use two types of privilege delegation tools: Sudo and PowerBroker. You can use EM CLI or the Manage Privilege Delegation Settings page to set/edit privilege delegation settings for a host. See the *Enterprise Manager Command Line Interface* guide for more information on using the command line.

Sudo: sudo allows a permitted user to execute a command as the super user or another user, as specified in the sudoers file. If the invoking user is root or if the target user is the same as the invoking user, no password is required. Otherwise, sudo requires that users authenticate themselves with a password by default. Once a user has been authenticated, a timestamp is updated and the user may then use sudo without a password for a short period of time (5 minutes unless overridden in sudoers). sudo determines who is an authorized user by consulting the file `/etc/sudoers` file. For more information, see the manual page on sudo (man sudo) on Unix. Enterprise Manager authenticates the user using sudo, and executes the script as sudo. For example, if the command to be executed is foo -arg1 -arg2, it will be executed as sudo -S foo -arg1 -arg2.

PowerBroker: BeyondTrust Powerbroker enables UNIX system administrators to specify the circumstances under which other people may run certain programs such as root (or other important accounts). The result is that responsibility for such actions as adding user accounts, fixing line printer queues, and so on, can be safely assigned to the appropriate people, without disclosing the root password. The full power of root is thus protected from potential misuse or abuse-for example, modifying databases or file permissions, erasing disks, or more subtle damage.

BeyondTrust PowerBroker can access existing programs as well as its own set of utilities that execute common system administration tasks. Utilities being developed to run on top of BeyondTrust PowerBroker can manage passwords, accounts, backups, line printers, file ownership or removal, rebooting, logging people out, killing their

programs, deciding who can log in to where from where, and so on. They can also provide TCP/IP, Load Balancer, cron, NIS, NFS, FTP, rlogin, and accounting subsystem management. Users can work from within a restricted shell or editor to access certain programs or files as root.

See your Sudo or PowerBroker documentation for detailed setup and configuration information.

10.5.4.1 Creating a Privilege Delegation Setting

Enterprise Manager allows you to create privilege delegation settings either by creating the setting directly on a host target, or by creating a Privilege Delegation Setting Template that you can apply to multiple hosts.

To create a privilege delegation setting directly on a host:

1. From the **Setup** menu, select **Security**, then select **Privilege Delegation**. The following screen is displayed:

Figure 10–7 Manage Privilege Delegation Settings

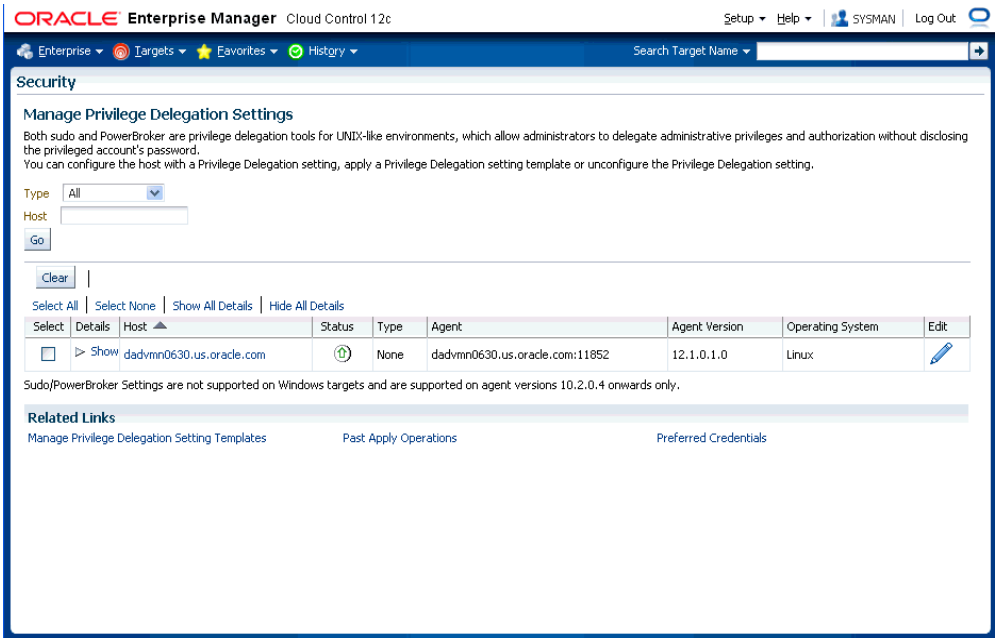

2. For any host target appearing in the table, click **Edit**. Enterprise Manager takes you to the Host Privilege Delegation Setting page.
3. Select a privilege delegation type (Sudo or PowerBroker).
4. Enter the privilege delegation command to be used and, in the case of PowerBroker, the optional Password Prompt.
5. Click **Update** to apply the settings to the host. The following figure shows the Host Privilege Delegation Setting window that you can use to create a PowerBroker setting.

Figure 10–8 Host Privilege Delegation Setting - PowerBroker

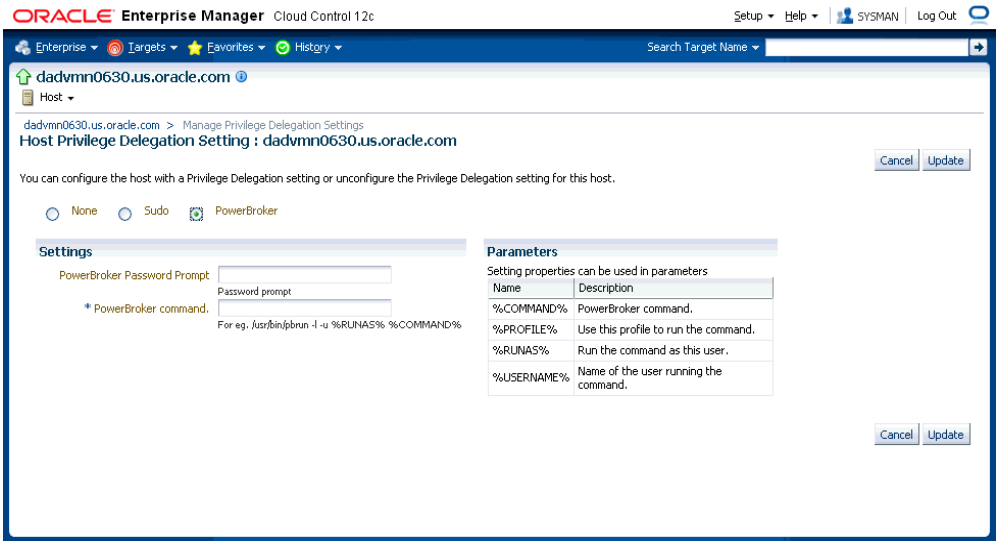

Once you have created a privilege delegation setting, you must apply this setting to selected targets. This setting can be applied to one more hosts or to a composite (Group) target (the group must contain at least one host target). You can apply a Privilege Delegation setting using the Cloud Control console. From the **Setup** menu, choose **Security** and then **Privilege Delegation**.

10.6 Cryptographic Support

To protect the integrity of sensitive data in Enterprise Manager, a signing on verification method known as the `emkey` is used. Encryption key is the master key that is used to encrypt/decrypt sensitive data, such as passwords and preferred credentials that are stored in the Repository. The key is originally in stored in repository. It is removed from repository and copied to the Credential Store during installation of the first OMS. (the `emkey` is secured out-of-the-box). A backup is created in `OMS_ORACLE_HOME/sysman/config/emkey.ora`. It is recommended to create a backup of this file on some other machine. When starting up, OMS reads the `emkey` from Credential Store and repository. If the `emkey` is not found or is corrupted, it fails to start. By storing the key separately from Enterprise Manager schema, we ensure that the sensitive data such as Named Credentials in the Repository remain inaccessible to the schema owner and other SYSDBA users (Privileged users who can perform maintenance tasks on the database) in the Repository. Moreover, keeping the key from the schema will ensure that sensitive data remain inaccessible while Repository backups are accessed. Further, the schema owner should not have access to the OMS/Repository Oracle homes.

10.6.1 Configuring the emkey

The `emkey` is an encryption key that is used to encrypt and decrypt sensitive data in Enterprise Manager such as host passwords, database passwords and others.

WARNING: If the `emkey.ora` file is lost or corrupted, all of the encrypted data in the Management Repository becomes unusable. Maintain a backup copy of this file on another system.

During startup, the Oracle Management Service checks the status of the emkey. If the emkey has been properly configured, it uses it encrypting and decrypting data. If the emkey has not been configured properly, the following error message is displayed.

Example 10–11 emctl start oms Command

```
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
emctl start oms
Starting HTTP Server ...
Starting Oracle Management Server ...
Checking Oracle Management Server Status ...
Oracle Management Server is not functioning because of the following reason:
The Enterprise Manager Key is not configured properly. Run "emctl status emkey"
for more details.
```

10.6.2 emctl Commands

The emctl commands related to emkey are given below:

- emctl status emkey [-sysman_pwd <pwd>]
- emctl config emkey -copy_to_credstore [-sysman_pwd <pwd>]
- emctl config emkey -copy_to_repos [-sysman_pwd <pwd>]
- emctl config emkey -remove_from_repos [-sysman_pwd <pwd>]
- emctl config emkey -copy_to_file_from_credstore -admin_host <host> -admin_port <port> -admin_user <username> [-admin_pwd <pwd>] [-repos_pwd <pwd>] -emkey_file <emkey file>
- emctl config emkey -copy_to_file_from_repos (-repos_host <host> -repos_port <port> -repos_sid <sid> | -repos_conndesc <conn desc>) -repos_user <username> [-repos_pwd <pwd>] [-admin_pwd <pwd>] -emkey_file <emkey file>
- emctl config emkey -copy_to_credstore_from_file -admin_host <host> -admin_port <port> -admin_user <username> [-admin_pwd <pwd>] [-repos_pwd <pwd>] -emkey_file <emkey file>
- emctl config emkey -copy_to_repos_from_file (-repos_host <host> -repos_port <port> -repos_sid <sid> | -repos_conndesc <conn desc>) -repos_user <username> [-repos_pwd <pwd>] [-admin_pwd <pwd>] -emkey_file <emkey file>

10.6.2.1 emctl status emkey

This command shows the health or status of the emkey. Depending on the status of the emkey, the following messages are displayed:

- When the emkey has been correctly configured in the Credential Store, the following message is displayed.

Example 10–12 emctl status emkey - Example 1

```
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EmKey is configured properly, but is not secure. Secure the EmKey by running
"emctl config emkey -remove_from_repos"
```

- When the emkey has been correctly configured in the Credential Store and has been removed from the Management Repository, the following message is displayed.

Example 10–13 emctl status emkey - Example 2

```
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey is configured properly.
```

- When the emkey is corrupted in the Credential Store and removed from the Management Repository, the following message is displayed.

Example 10–14 emctl status emkey - Example 3

```
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey is not configured properly or is corrupted in the credential store and
does not exist in the Management Repository. To correct the problem:
1) Get the backed up emkey.ora file.
2) Configure the emkey by running "emctl config emkey -copy_to_credstore_from_
file"
```

10.6.2.2 emctl config emkey -copy_to_credstore

This command copies the emkey from the Management Repository to the Credential Store.

Example 10–15 Sample Output of the emctl config emkey -copy_to_credstore Command

```
emctl config emkey -copy_to_credstore [-sysman_pwd <pwd>]
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been copied to the Credential Store.
```

10.6.2.3 emctl config emkey -copy_to_repos

This command copies the emkey from the Credential Store to Management Repository.

Example 10–16 Sample Output of the emctl config emkey -copy_to_repos Command

```
emctl config emkey -copy_to_repos [-sysman_pwd <pwd>]
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been copied to the Management Repository. This operation will cause
the EMKey to become unsecure.
After the required operation has been completed, secure the EMKey by running
"emctl config emkey -remove_from_repos".
```

10.6.2.4 emctl config emkey -copy_to_file_from_credstore

This command copies the emkey from the Credential Store to a specified file.

Example 10–17 Sample Output of the emctl config emkey -copy_to_file_from_credstore Command

```
emctl config emkey -copy_to_file_from_credstore -admin_host <host> -admin_port
<port> -admin_user <username> [-admin_pwd <pwd>] [-repos_pwd <pwd>] -emkey_file
<emkey file>
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been copied to file.
```

10.6.2.5 emctl config emkey -copy_to_file_from_repos

This command copies the emkey from the Management Repository to a specified file.

Example 10–18 Sample Output of the emctl config emkey -copy_to_file_from_repos Command

```
emctl config emkey -copy_to_file_from_repos (-repos_host <host> -repos_port <port>
-repos_sid <sid> | -repos_conndesc <conn desc>) -repos_user <username> [-repos_pwd
<pwd>] [-admin_pwd <pwd>] -emkey_file <emkey file>
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been copied to file.
```

10.6.2.6 emctl config emkey -copy_to_credstore_from_file

This command copies the emkey from a specified file to the Credential Store.

Example 10–19 Sample Output of the emctl config emkey -copy_to_credstore_from_file Command

```
emctl config emkey -copy_to_credstore_from_file -admin_host <host> -admin_port
<port> -admin_user <username> [-admin_pwd <pwd>] [-repos_pwd <pwd>] -emkey_file
<emkey file>
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been copied to the Credential Store.
```

10.6.2.7 emctl config emkey -copy_to_repos_from_file

This command copies the emkey from a specified file to the repository.

Example 10–20 Sample Output of the emctl config emkey -copy_to_repos_from_file Command

```
emctl config emkey -copy_to_repos_from_file (-repos_host <host> -repos_port <port>
-repos_sid <sid> | -repos_conndesc <conn desc>) -repos_user <username> [-repos_pwd
<pwd>] [-admin_pwd <pwd>] -emkey_file <emkey file>
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been copied to the Management Repository. This operation will cause
the EMKey to become unsecure.
After the required operation has been completed, secure the EMKey by running
"emctl config emkey -remove_from_repos".
```

10.6.2.8 emctl config emkey -remove_from_repos

This command removes the emkey from the repository.

Example 10–21 Sample Output of emctl config emkey -remove_from_repos Command

```
emctl config emkey -remove_from_repos [-sysman_pwd <pwd>]
Oracle Enterprise Manager 12c Release 1 Cloud Control
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
The EMKey has been removed from the Management Repository.
```

Note: If the emkey is corrupted in the Credential Store, you will not be able to remove it from the Management Repository.

10.6.3 Install and Upgrade Scenarios

This section explains the install and upgrade scenarios for emkey.

10.6.3.1 Installing the Management Repository

A new emkey is generated as a strong random number when the Management Repository is installed.

10.6.3.2 Installing the First Oracle Management Service

When the Oracle Management Service is installed, the Installer copies the emkey to Credential Store and removes it from repository (emkey is secured out-of-box).

10.6.3.3 Upgrading from 10.2 or 11.1 to 12.1

The Management Repository is upgraded as usual. When upgrading the OMS, the `omsca` (OMS Configuration Assistant) copies the emkey to Credential Store and removes from repository. If the emkey is already secured before upgrade or has been removed from repository, then `omsca` reads the emkey from emkey.ora file present in old OMS Oracle Home and copies it to Credential Store.

Note: After all the Oracle Management Service have been upgraded, you can secure the emkey, that is, remove it from the Management Repository by running the following command:

```
emctl config emkey -remove_from_repos
```

10.6.3.4 Recreating the Management Repository

When the Management Repository is recreated, a new emkey is generated. This new key will not be in synchronization with the existing emkey in the Credential Store.

1. Copy the new emkey to Credential Store by using the `emctl config emkey -copy_to_credstore` command.
2. Take a backup by entering the `emctl config emkey -copy_to_file_from_repos` command or the `emctl config emkey -copy_to_file_from_credstore` command.
3. Secure the emkey by using the `emctl config emkey -remove_from_repos` command.

10.7 Setting Up the Auditing System for Enterprise Manager

All operations performed by Enterprise Manager users such as creating users, granting privileges, starting a remote job like patching or cloning, need to be audited to ensure compliance with the Sarbanes-Oxley Act of 2002 (SAS 70). This act defines standards an auditor must use to assess the contracted internal controls of a service organization. Auditing an operation enables an administrator to monitor, detect, and investigate problems and enforce enterprise wide security policies.

Irrespective of how the user has logged into Enterprise Manager, when auditing is enabled, each user action is audited and the audit details are stored in a record.

For Enterprise Manager 12c, BASIC auditing is enabled by default, thus creating an audit trail of credentials being created, edited, accessed, associated and deleted. Named credentials are first-class security objects on which privileges can be granted or revoked privileges. This means that multiple Enterprise Manager administrators will be able to use and modify the credential objects. Because credentials are sensitive data that can be used to perform various operations on the systems, there is a need to audit the operations on credentials.

Enterprise Manger audits all the operations performed on credentials. The auditing information includes, but is not be limited to, the current username, credential name, operation performed, operation status success or failure. The audit logs contain information about the credential owner, action initiator, credential name, user name, and target name, job names along with the date time of the operation. Credential fields like password, private keys are never logged.

The following operations are audited:

- **Creating a Named Credential**: Creating new Enterprise Manager credentials will be audited.
- **Editing a Named Credential**: Editing a credential may consist of changing the username and/or the sensitive credential attributes. Credential edits may also include changing the authentication scheme for the credential.
- **Delete a Named Credential**: Deleting a credential from Enterprise Manager will be audited.
- **Associating a Named Credential**: A named credential can be set as a preferred credential for a credential set at the target level or at target type level. The named credential can also be reference directly from a job. All operations involving the setting of the named credentials as preferred credentials and using it in a job or deployment procedure will be audited.
- **Accessing a Named Credential**: Enterprise Manager subsystems request credentials from the credential store to perform various system management tasks

10.7.1 Configuring the Enterprise Manager Audit System

You can configure the Enterprise Manager Audit System by using the following EM CLI commands:

- `enable_audit`: Enables auditing for all user operations.
- `disable_audit`: Disables auditing for all user operations.
- `show_operations_list`: Shows a list of the user operations being audited.
- `show_audit_settings`: Shows the audit status, operation list, externalization service details, and purge period details.

10.7.2 Configuring the Audit Data Export Service

Audit data needs to be protected and maintained for several years. The volume of audit data may become very large and impact the performance of the system. To limit the amount of data stored in the repository, the audit data must be externalized or archived at regular intervals. The archived audit data is stored in an XML file complying with the ODL format. To externalize the audit data, the `EM_AUDIT_EXTERNALIZATION` API is used. Records of the format <file-prefix>.NNNNN.xml, where NNNN is a number are generated. The numbers start with 00001 and continue to 99999.

You can set up the audit externalization service for exporting audit data into the file system by using the `update_audit_setting -externalization_switch` command.

10.7.3 Updating the Audit Settings

The `update_audit_settings` command updates the current audit settings in the repository and restarts the Management Service.

Example 10–22 Usage of the update_audit_setting command

```
emcli update_audit_settings
      -audit_switch="ENABLE/DISABLE"
      -operations_to_enable="name of the operations to enable, for all oprtations
       use ALL"
      -operations_to_disable="name of the operations to disable, for all
       oprtations use ALL"
      -externalization_switch="ENABLE/DISABLE"
      -directory_name="directory_name (DB Directory)"
      -file_prefix="file_prefix"
      -file_size="file_size (Bytes)"
      -data_retention_period="data_retention_period (Days)"
```

- `-audit_switch`: Enables auditing across Enterprise Manager. The possible values are `ENABLE/DISABLE`. Default value is `DISABLE`.
- `-operations_to_disable`: Enables auditing for specified operations. Enter **All** to enable all operations.
- `-operations_to_disable`: Disables auditing for specified operations. Enter **All** to disable all operations.
- `-externalization_switch`: Enables the audit data export service. The possible values are `ENABLE/DISABLE`. Default value is `DISABLE`.
- `-directory`: The database directory that is mapped to the OS directory where the export service archives the audit data files.
- `-file_prefix`: The file prefix to be used by the export service to create the file in which audit data is to be stored.
- `-file_size`: The size of the file on which the audit data is to be stored. The default value is 5000000 bytes.
- `data_retention_period`: The period for which the audit data is to be retained inside the repository. The default value is 365 days.

10.7.4 Searching the Audit Data

You can search for audit data that has been generated over a specified period. You can also search for the following:

- Audit details of a specific user operation or all user operations.
- Audit details of operations with a Success or Failure status or All operations.

From the **Setup** menu, select **Management Services and Repository**. The Overview page is displayed. Click the **Audit Data** link under the Audit section. The Audit Data page is displayed. Specify the search criteria in the fields and click **Go**. The results are displayed in the Summary table.

Figure 10–9 Audit Data Search Page

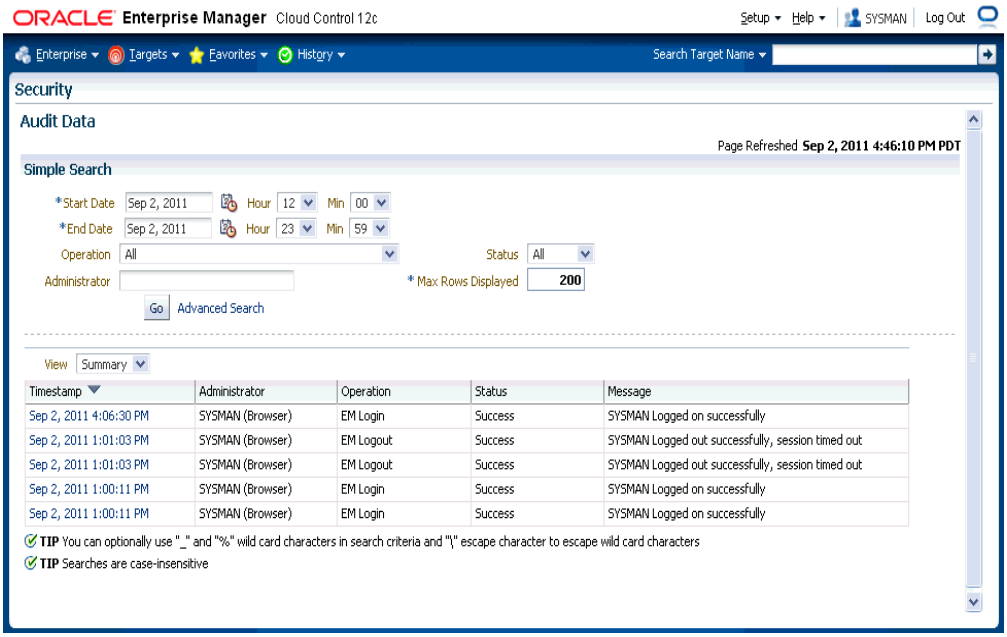

To view the details of each record that meets the search criteria, select Detailed in the View drop-down list. To drill down to the full record details, click on the **Timestamp**. The Audit Record page is displayed.

Figure 10–10 Audit Record Details Page

Field Name	Description
General	

Field Name	Description
Operation Timestamp	The date and time on which the operation took place.
Administrator	The id of the administrator who has logged into Enterprise Manager.
Operation	The type of operation being audited.
Status	The status of the operation which can be success or failure.
Message	A descriptive message indicating the status of the operation.
Normalized Timestamp	This is the UTC timestamp.
Client Information	
Session	This can either be the HTTP Session ID or the DBMS Session ID.
IP Address	The IP address of the client's host machine.
Hostname	The name of the client's host machine.
Upstream Component Type	The type of client, Console, Web Service, EM CLI, being used.
Authentication Type	The nature of the session (HTTP Session, DB Session).
Upstream Component Name	The name of the client being used.
OMS Information	
Hostname	The host name of the Oracle Management Service.
IP Address	The IP address of the Oracle Management Service.
Instance ID	The Instance ID of the Oracle Management Service.
Operation Specific Information	
Object Name	The operation being performed on an object

10.7.5 List of Operations Audited

The following table lists the names of operation and their description.

Table 10–5 List of Operations Audited

Operation Name	Description
TCAUD_ADD_TEMPLATE_ENTITY	Add entity to Template Collection
ADD_AGENT_REGISTRATION_ PASSWORD	Add Registration Password
SWLIBADDLOCATION	Configuring a new storage location in Software Library
ADD_CS_TARGET_ASSOC	Add Standard-Target Association
APPLY_TEMPLATE	APPLY_TEMPLATE
APPLY_UPDATE	Apply the update
TCAUD_ASSOC_TO_AG	Associate Template Collection to AG
ATTACH_MEXT	Attach Metric Extension

Table 10–5 (Cont.) List of Operations Audited

Operation Name	Description
AUDIT_EXPORT_SETTINGS	Audit Export Settings to externalize audit data
AUDIT_SETTINGS	Audit Settings to enable or Disable auditing
CHANGE_CONNECTOR_SETTINGS	enable/disable a Connector
CHANGE_PASSWORD	Change Password
CHANGE_PREFERRED_CREDENTIAL	change_pref_cred
CONFIG_CONNECTOR	Configure a Connector Instance
CREATE_CHANGE_MANAGEMENT_SETTING	Create the change management settings for the Real-time Monitoring rule.
CREATE_CONNECTOR	Create a Connector Instance
CCS_CREATE_MD	Create (or import) Custom Configuration Specification
CCS_CREATE_PARSER	Create Custom Configuration Specification Parser
CCS_CREATE_CUSTOM_TARGET_TYPE	Create Custom Target Type
CREATE_FACET	Create a new facet.
CREATE_FACET_PARAMETER	Create a new facet parameter.
CREATE_FACET_PATTERN	Create a new facet pattern.
CREATE_CSG	Create Framework
CREATE_MEXT	Create Metric Extension
CREATE_TEMPLATE	CREATE_TEMPLATE
CREATE_NAMED_CREDENTIAL	Create Named Credential
CREATE_PG_SCHED	Create Policy Group Schedule
CREATE_CCC_RULE	Create a Real-time Monitoring rule.
RES_STATE_CREATE_OP	Resolution State created
CREATE_ROLE	Create Role
CREATE_RULE	Create Rule
CREATE_CS	Create Standard
TCAUD_CREATE	Create Template Collection
CREATE_USER	Create User
CREATE_UDP	Create User Defined Policy
CREATE_UDPG	Create User Defined Policy Group
DB_LOGIN	Audit Database user Login
DB_LOGOUT	Audit Database user Logout
DELETE_CONNECTOR	Delete a Connector Instance
CCS_DELETE_MD	Delete Custom Configuration Specification
CCS_DELETE_PARSER	Delete Custom Configuration Specification Parser
DELETE_FACET	Delete a facet.

Table 10–5 (Cont.) List of Operations Audited

Operation Name	Description
DELETE_FACET_PARAMETER	Delete a facet parameter.
DELETE_FACET_PATTERN	Delete a facet pattern.
DELETE_CSG	Delete framework
DELETE_JOB	Delete job.
DELETE_MEXT	Delete Metric Extension
DELETE_TEMPLATE	Delete a monitoring template.
DELETE_NAMED_CREDENTIAL	Delete Named Credential
DELETE_PG_EVAL	Delete Policy Group Evaluation Results
DELETE_PG_SCHED	Delete Policy Group Schedule
DELETE_CCC_RULE	Delete a Real-time Monitoring rule.
DELETE_AGENT_REGISTRATION_ PASSWORD	Delete Registration Password
RES_STATE_DELETE_OP	Resolution State deleted
DELETE_ROLE	Drop Role
DELETE_RULE	Delete Rule
SWLIBDELETEFOLDER	Deleting a directory in Software Library
SWLIBDELETEENTITY	Deleting an entity in Software Library
DELETE_CS	Delete Standard
TCAUD_DELETE	Delete Template Collection
DELETE_UPDATE	Delete the update
DELETE_USER	Delete User
DELETE_UDP	Delete User Defined Policy
DELETE_UDPG	Delete User Defined Policy Group
CCS_DEPLOY	Deploy Custom Configuration Specification
DETACH_MEXT	Detach Metric Extension
DISABLE_CS_TARGET_ASSOC	Disable Standard-Target Association
TCAUD_DEASSOC_FROM_AG	Disassociate Template Collection from AG
DOWNLOAD_UPDATE	Download an available update
EDIT_CSG	Edit Framework
EDIT_JOB	edit_job
EDIT_TEMPLATE	EDIT_TEMPLATE
EDIT_PG_SCHED	Edit Policy Group Schedule
EDIT_AGENT_REGISTRATION_ PASSWORD	Edit Registration Password
EDIT_RULE	Edit Rule
EDIT_CS	Edit Standard
EDIT_CS_TARGET_ASSOC	Edit Standard-Target Association
TCAUD_EDIT	Edit Template Collection

Table 10–5 (Cont.) List of Operations Audited

Operation Name	Description
EDIT_UDP	Edit User Defined Policy
EDIT_UDPG	Edit User Defined Policy Group
LOGIN	logon
LOGOUT	logoff
ENABLE_CS_TARGET_ASSOC	Enable Standard-Target Association
EVALUATE_UDP	Evaluate User Defined Policy
PERFORM_OPERATION_AS_AGENT	Execute any OS Command as the Agent User (uncredentialed)
FILE_TRANSFER	file_transfer
GET_FILE	get_file
GET_NAMED_CREDENTIAL	Get Named Credential
GRANT_JOB_PRIVILEGE	Grant Job Privilege
GRANT_PRIVILEGE	Grant Privilege
GRANT_ROLE	Grant Role
GRANT_SYSTEM_PRIVILEGE	Grant System Privilege
GRANT_TARGET_PRIVILEGE	Grant Target Privilege
IMPORT_FACET	Import a facet.
IMPORT_CSG	Import Framework
IMPORT_CCC_RULE	Import a Real-time Monitoring rule.
IMPORT_RULE	Import Rule
IMPORT_CS	Import Standard
IMPORT_UDP	Import User Defined Policy
INCLUDE_ACTION_TO_MONITOR	Include an action to monitor for the Real-time Monitoring rule.
INCLUDE_FILTER_FACET	Include a filter facet into the Real-time Monitoring rule.
INCLUDE_MONITORING_FACET	Include a monitoring facet into the Real-time Monitoring rule.
JOB_OUTPUT	Job output obtained after job execution
MODIFY_CHANGE_MANAGEMENT_SETTING	Modify the change management settings for the Real-time Monitoring rule.
MODIFY_FACET	Update a facet.
MODIFY_FACET_CONTENT	Update the basic facet information.
MODIFY_FACET_PARAMETER	Update a facet parameter.
MODIFY_FACET_PATTERN	Update a facet pattern.
MODIFY_METRIC_SETTINGS	MODIFY_METRIC_SETTINGS
UPDATE_NAMED_CREDENTIAL	Modify Named Credential
MODIFY_POLICY_SETTINGS	Modify Policy Settings
MODIFY_CCC_RULE	Update a Real-time Monitoring rule.

Table 10–5 (Cont.) List of Operations Audited

Operation Name	Description
RES_STATE_MODIFY_OP	Resolution State modified
MODIFY_ROLE	Modify Role
MODIFY_USER	Modify User
SWLIBMOVEENTITY	Moving all revisions of an entity in Software Library to another directory
PUBLISH_MEXT	Publish Metric Extension
SWLIBPURGELOCATION	Purging a storage location in Software Library
PUT_FILE_AS_AGENT	Put any File to the Agent's Filesystem as the Agent User (uncredentialed)
PUT_FILE	put_file
REFRESH_UPDATE	Refresh from EM Store
AGENT_REGISTRATION_PASSWORD_USAGE	Registration Password Usage
REMOTE_OPERATION_JOB	remote_op
REMOVE_ACTION_FROM_MONITOR	Remove an action from monitor for the Real-time Monitoring rule.
REMOVE_CHANGE_MANAGEMENT_SETTING	Remove the change management settings for the Real-time Monitoring rule.
TCAUD_REMOVE_TEMPLATE_ENTITY	Remove entity from Template Collection
REMOVE_FILTER_FACET	Remove a filter facet from the Real-time Monitoring rule.
REMOVE_MONITORING_FACET	Remove a monitoring facet from the Real-time Monitoring rule.
REMOVE_PRIVILEGE_DELEGATION_SETTING	Remove Privilege Delegation Setting
SWLIBDELETELOCATION	Removing a storage location in Software Library
REMOVE_CS_TARGET_ASSOC	Remove Standard-Target Association
REMOVE_UPDATE	Remove the update
TCAUD_RENAME	Rename Template Collection
AGENT_RESYNC	Agent resynchronization operation
REPOSITORY_RESYNC	Repository resynchronization operation
RETRY_JOB	
REVOKE_JOB_PRIVILEGE	Revoke Job Privilege
REVOKE_PRIVILEGE	Revoke Privilege
REVOKE_ROLE	Revoke Role
REVOKE_SYSTEM_PRIVILEGE	Revoke System Privilege
REVOKE_TARGET_PRIVILEGE	Revoke Target Privilege
SAVE_MONITORING_SETTINGS	SAVE_MONITORING_SETTINGS
SET_PRIVILEGE_DELEGATION_SETTING	Set Privilege Delegation Setting

Table 10–5 (Cont.) List of Operations Audited

Operation Name	Description
STOP_JOB	
SUBMIT_JOB	submit_job
SUBSCRIBE_UPDATE	Subscribe to an Update Type
SUSPEND_JOB	Suspend job
CCS_UNDEPLOY	Undeploy Custom Configuration Specification
UNSUBSCRIBE_UPDATE	Unsubscribe an Update Type
CCS_UPDATE_MD	Update Custom Configuration Specification
UPDATE_DB_PASSWORD	Update Database Password
INSERT_UPDATE	Show the update on self update home
UPDATE_MEXT	Update Metric Extension
UPDATE_PASSWORD	Update Password

10.8 Additional Security Considerations

After you enable security for the Enterprise Manager components and framework, there are additional security considerations. This section provides the following topics:

- Changing the SYSMAN and MGMT_VIEW Passwords
- Responding to Browser-Specific Security Certificate Alerts
- Configuring Beacons to Monitor Web Applications Over HTTPS
- Patching Oracle Homes When the User is Locked
- Cloning Oracle Homes

10.8.1 Changing the SYSMAN and MGMT_VIEW Passwords

This section describes the commands used to change the SYSMAN and MGMT_VIEW passwords.

10.8.1.1 Changing the SYSMAN User Password

To change the password of the `SYSMAN` user, you use the following command:

```
emctl config oms -change_repos_pwd [-old_pwd <old_pwd>] [-new_pwd <new_pwd>]
[-use_sys_pwd [-sys_pwd <sys_pwd>]]
```

Parameter	Description
-change_repos_pwd	.
`-old_pwd`	This is the current SYSMAN password.
`-new_pwd`	This is the new password.
`-use_sys_pwd`	This parameter is optional and is used to connect to the database as a SYS user. Use this option if SYSMAN account on the database has expired/locked.
`-sys_pwd`	This is the password for the SYS user.

1. Stop all OMS instances.

   ```
   emctl stop oms
   ```

2. For each OMS, run the following command:

   ```
   emctl config oms -change_repos_pwd'
   ```

3. Restart the Administration Server and all OMS instances.

   ```
   emctl stop oms -all
   emctl start oms
   ```

10.8.1.2 Changing the MGMT_VIEW User Password

To change the password of the `MGMT_VIEW` user, you use the following command:

```
emctl config oms -change_view_user_pwd [-sysman_pwd <sysman_pwd>] [-user_pwd
<user_pwd>] [-auto_generate]
```

Parameter	Description
`-change_view_user_pwd`	Used to change MGMT_VIEW user's password.
`-sysman_pwd`	The password for the SYSMAN user.
`-user_pwd`	The new password for theMGMT_VIEW user..
`-auto_generate`	If this option is specified, the password is auto generated.

Important: In order to change the MGMT_VIEW password, you must ensure that the password of the WebLogic administrative user in the credential store matches the actual password of the user SYSMAN. If the credentials do not match, the connection to the Repository Database fails and the SYSMAN password cannot be modified

1. Stop all OMS instances.

   ```
   emctl stop oms -all
   ```

2. Execute the following command to update the WebLogic and nodemanager passwords in the Credential store:

   ```
   cd <OMS_HOME>/bin
   emctl secure create_admin_creds_wallet -admin_pwd <existing weblogic pwd>
   -nodemgr_pwd <existing nodemanager pwd>
   ```

3. Log into the Management Repository database as a DBA user and execute the following to manually modify the password of the sysman_mds schema to the new password that will be set for the sysman user:

   ```
   SQL> alter user sysman_mds identified by <new_pwd of sysman user>;
   ```

4. For **ONE** of the OMSs, run the following command to modify the SYSMAN password::

   ```
   cd <OMS_HOME>/bin
   emctl config oms -change_repos_pwd -change_in_db -old_pwd <new_pwd> -new_pwd
   <new_pwd>
   ```

5. Restart the AdminServer and all the OMSs.

   ```
   emctl stop oms -all
   ```

```
emctl start oms
```

10.8.2 Responding to Browser-Specific Security Certificate Alerts

This section describes how to respond to browser-specific security alert dialog boxes when you are using Enterprise Manager in a secure environment.

The security alert dialog boxes described in this section appear because Enterprise Manager Framework Security is enabled, but you have not secured your web tier properly.

This section contains the following topics:

- Responding to the Internet Explorer Security Alert Dialog Box
- Responding to Mozilla Firefox New Site Certificate Dialog Box

10.8.2.1 Responding to the Internet Explorer Security Alert Dialog Box

Security is enabled by default for the Management Service. However, if you have not enabled the more extensive security features of your web tier, you will likely receive the following warning: "There is a problem with this Web site's security certificate." This occurs when you first attempt to display the Cloud Control console using the HTTPS URL in Internet Explorer.

Figure 10–11 Internet Explorer Security Alert

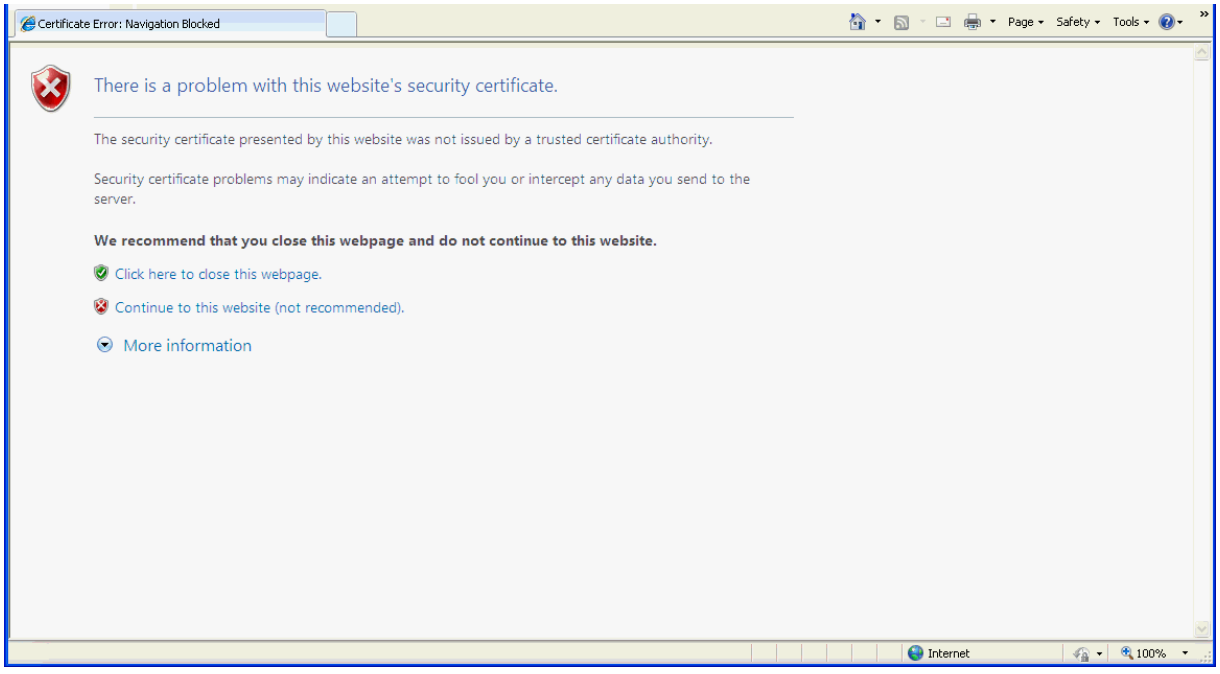

When Internet Explorer displays the certificate warning page, use the following instructions to install the certificate and avoid viewing this page again in future Enterprise Manager sessions:

1. From the certificate warning page, click **Continue to this Web site (not recommended).**

   Internet Explorer displays a Security Warning dialog.

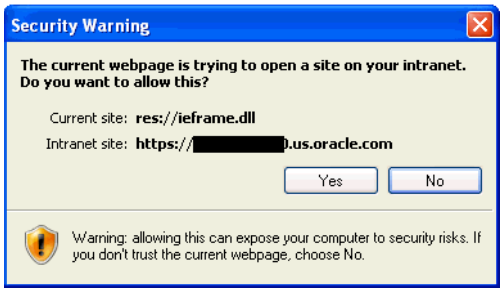

2. Click **Yes**. Internet Explorer may display a Security Alert dialog if you have not selected "In the future, do not show this warning." in a previous Internet Explorer session. Click OK to dismiss the dialog.

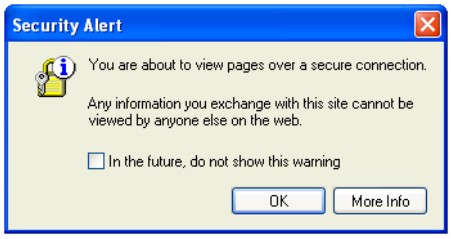

3. The Enterprise Manager console login page displays.
4. At the top of the browser, Click **Certificate Error** to display the **Security Report** pop-up.

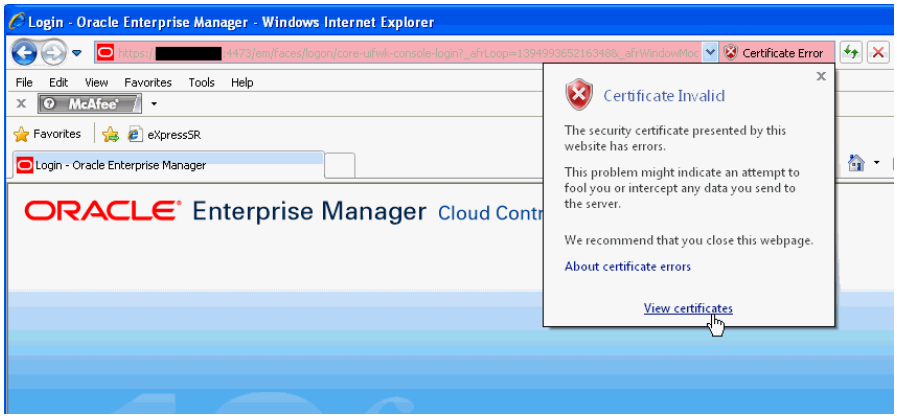

5. Click **View Certificates**. The Certificates dialog appears.

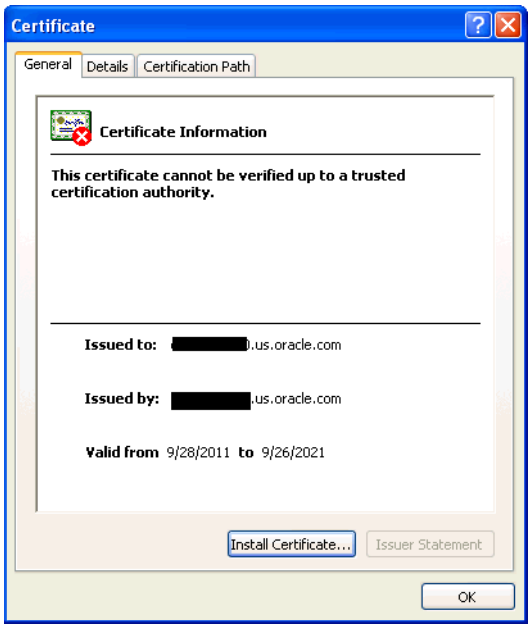

6. Click the **Certificate Path** tab and select the first entry in the list of certificates as shown in the following graphic.

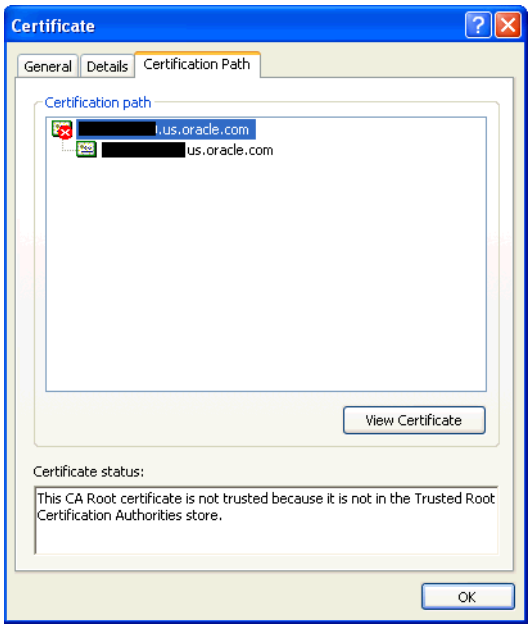

7. Click **View Certificate** to display a second Certificate dialog box.
8. Click **Install Certificate** to display the Certificate Import wizard.
9. Accept the default settings in the wizard, click **Finish** when you are done.

   Internet Explorer displays a Security Warning asking if you want to install the certificate. Click **Yes.** Internet Explorer will display a message stating that the certificate was imported successfully.
10. Click **OK** to close each of the security dialog boxes and click **Yes** on the Security Alert dialog box to continue with your browser session.

You should no longer receive the Security Alert dialog box in any future connections to Enterprise Manager when you use this browser.

10.8.2.2 Responding to Mozilla Firefox New Site Certificate Dialog Box

Firefox will also issue a connection warning if you have not enabled its more extensive security features. When you first attempt to display the Cloud Control console using the HTTPS URL in Mozilla Firefox, you will recieve a warning because the connection is untrusted.

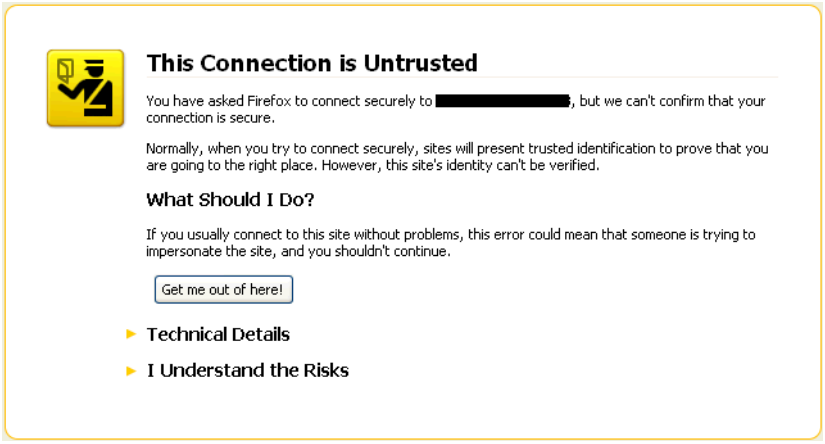

When Firefox displays the Untrusted Connection page, use the following instructions to install the certificate and avoid viewing this page again in future Enterprise Manager sessions:

1. Review the instructions and information. Click **I Understand the Risks**. Firefox displays additional information and the opportunity to add the certificate.
2. Click **Add Exception...** . Firefox displays the **Add Security Exception** dialog.

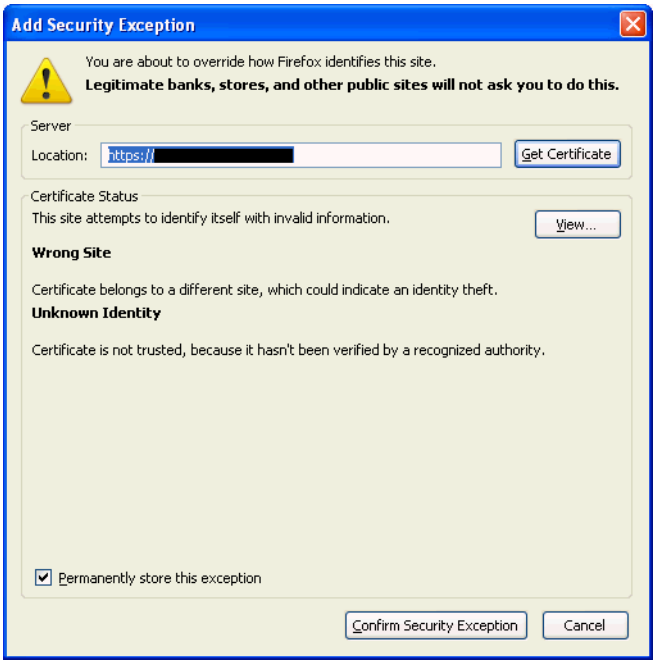

3. Ensure that the **Permanently store this exception option** is selected.

   You should no longer receive the New Site Certificate dialog box when using the current browser.

4. Click **Confirm Security Exception**. The Enterprise Manager console displays.

You will no longer receive the untrusted connection warning in any future connections to Enterprise Manager when you use this browser

10.8.3 Configuring Beacons to Monitor Web Applications Over HTTPS

Oracle Beacons provide application performance availability and performance monitoring. They are part of the Application Service Level Management features of Enterprise Manager.

> **See Also:** "About Application Service Level Management" in the Enterprise Manager Online Help

When a Beacon is used to monitor a URL over Secure Sockets Layer (SSL) using an HTTPS URL, the Beacon must be configured to recognize the Certificate Authority that has been used by the Web site where that URL resides.

> **See Also:** "The Public Key Infrastructure Approach to Security" in the *Oracle Security Overview* for an overview of Public Key Infrastructure features, such as Certificate Authorities

The Beacon software is preconfigured to recognize most commercial Certificate Authorities that are likely to be used by a secure Internet Web Site. However, you may encounter Web Sites that, although available over HTTPS, do not have a Certificate that has been signed by a commercial Certificate Authority recognized by the Beacon. The following are out-of-box certificates recognized by Beacons:

- Class 1 Public Primary Certification Authority by VeriSign, Inc.
- Class 2 Public Primary Certification Authority by VeriSign, Inc.
- Class 3 Public Primary Certification Authority by VeriSign, Inc.
- Secure Server Certification Authority by RSA Data Security, Inc.
- GTE CyberTrust Root by GTE Corporation
- GTE CyberTrust Global Root by GTE CyberTrust Solutions, Inc.
- Entrust.net Secure Server Certification Authority by Entrust.net ((c) 1999
- Entrust.net Limited, www.entrust.net/CPS incorp. by ref. (limits liab.))
- Entrust.net Certification Authority (2048) by Entrust.net ((c) 1999
- Entrust.net Limited, www.entrust.net/CPS_2048 incorp. by ref. (limits liab.))
- Entrust.net Secure Server Certification Authority by Entrust.net ((c) 2000
- Entrust.net Limited, www.entrust.net/SSL_CPS incorp. by ref. (limits liab.))

In those cases, for example, if you attempt to use the Test section of the Beacon Performance page to test the HTTP Response of the secure URL, the following error appears in the **Status Description** column of the Response Metrics table on the URL Test Page:

```
javax.net.ssl.SSLException: SSL handshake failed:
X509CertChainIncompleteErr--https://mgmtsys.acme.com/OracleMyPage.Home
```

See Also: "Using Beacons to Monitor Remote URL Availability" in the Enterprise Manager online help

To correct this problem, you can either set the service test "Authenticate SSL Certificates" property to "No"; or add the certificate authority to the list of Certificate Authorities recognized by Beacon.

Setting the Service Test "Authenticate SSL Certifications" Property:

1. From the **Target** menu, choose **Services**.
2. Select the desired service and click **Configure**. The Monitoring Configuration tab displays by default.
3. Click **Service Tests** and **Beacons**. The Service Tests and Beacons page displays.
4. Select the desired service test and click **Edit**. The Edit Service Test page displays.
5. In the **Transaction** region, select the desired service test.
6. Click on the **Advanced Properties** tab.
7. In the **Test Parameters** region, expand **Validation**.
8. For the **Authenticate SSL Certificates** property, choose **No** from the drop-down menu.

Configuring the Beacon to Recognize the Certificate Authority:

1. Obtain the Certificate of the Web Site's Certificate Authority, as follows:
   a. In Microsoft Internet Explorer, connect to the HTTPS URL of the Web Site you are attempting to monitor.
   b. Double-click the lock icon at the bottom of the browser screen, which indicates that you have connected to a secure Web site.

      The browser displays the Certificate dialog box, which describes the Certificate used for this Web site. Other browsers offer a similar mechanism to view the Certificate detail of a Web Site.
   c. Click the **Certificate Path** tab and select the first entry in the list of certificates.
   d. Click **View Certificate** to display a second Certificate dialog box.
   e. Click the **Details** tab on the Certificate window.
   f. Click **Copy to File** to display the Certificate Manager Export wizard.
   g. In the Certificate Manager Export wizard, select **Base64 encoded X.509 (.CER)** as the format you want to export and save the certificate to a text file with an easily-identifiable name, such as `beacon_certificate.cer`.
   h. Open the certificate file using a text editor.

      The content of the certificate file will look similar to the content shown in Example 10–23.
2. Update the list of Beacon Certificate Authorities as follows:
   a. Locate the `b64InternetCertificate.txt` file in the following directory of Agent Home of the Beacon host:

      agent_instance_home`/sysman/config/`

This file contains a list of Base64 Certificates.

b. Edit the `b64InternetCertificate.txt` file by appending the certificate text you exported using the Certificate Manager Export wizard to the end of the file.

Add only the certificate and not the peripheral text generated by the wizard. Do not include: "...base64 certificate content...", "-----BEGIN CERTIFICATE-----", or "-----END CERTIFICATE-----".

3. Restart the Management Agent.

After you restart the Management Agent, the Beacon detects your addition to the list of Certificate Authorities recognized by Beacon and you can successfully monitor the availability and performance of the secure Web site URL.

Example 10–23 Sample Content of an Exported Certificate

```
-----BEGIN CERTIFICATE-----
MIIDBzCCAnCgAwIBAgIQTs4NcImNY3JAs5edi/5RkTANBgkqhkiG9w0BAQQFADCB
... base64 certificate content...
-----END CERTIFICATE-----
```

10.8.4 Patching Oracle Homes When the User is Locked

To patch an Oracle Home used by a user "Oracle" and the user is locked:

1. Edit the default patching script and prepend sudo or sudo -u or pbrun -u to the default patching step. You need to set a policy (by editing the sudoers file) to allow the user submitting the job (who must be a valid operating system user) to be able to run sudo or pbrun without being prompted for password.

Note: You cannot patch Oracle Homes without targets. This must be done by using the Patching wizard.

10.8.5 Cloning Oracle Homes

The cloning application is wizard-driven. The source of the Oracle Home being cloned may be either an installed Oracle Home or a Software Library. Following are the steps in the cloning process:

1. If the source is an installed Oracle Home, then, after selecting the Oracle Home, a user will need to specify the Oracle Home credentials. These credentials once specified for an Oracle Home are stored in the repository. The next time a user clones the same Oracle Home, these credentials are automatically populated. Other parameters queried from the user at this point is a temporary location (on the source computer) and the list of files to be excluded from the Oracle Home. If the cloning source is a Software Library, the source Oracle Home credentials will not be queried for.

2. The user needs to specify the target location and provide the required credentials for each target location. These credentials will be the Oracle Home credentials for each of these target locations. Subsequently, if a user selects any of these cloned Oracle Homes as a source, the Oracle Home credentials are automatically populated.

3. Depending on the product being cloned, the user can view the Enterprise Manager page where query parameters required for the particular product being cloned are displayed.

4. The user can, then, view the execution of user-supplied pre-cloning and post-cloning scripts and the root.sh script. The root.sh script will always be run with sudo privileges, but the user has the option to decide if the pre-cloning and post-cloning scripts run with sudo privileges.
5. Finally, the user can schedule the cloning job at a convenient time.

For more information on cloning, refer to the Enterprise Manager Online Help.

Part III

Generating Reports

This section contains the following chapters:

- Using Information Publisher

11

Using Information Publisher

Information Publisher, Enterprise Manager's reporting framework, makes information about your managed environment available to audiences across your enterprise. Strategically, reports are used to present a view of enterprise monitoring information for business intelligence purposes, but can also serve an administrative role by showing activity, resource utilization, and configuration of managed targets. IT managers can use reports to show availability of sets of managed systems. Executives can view reports on availability of applications (such as corporate email) over a period of time.

> **Note:** The Information Publisher (IP) reporting framework is still supported for Enterprise Manager 12c, however, new report development using this framework has been deprecated for Enterprise Manager 12c.

The reporting framework allows you to create and publish customized reports: Intuitive HTML-based reports can be published via the Web, stored, or e-mailed to selected recipients. Information Publisher comes with a comprehensive library of predefined reports that allow you to generate reports out-of-box without additional setup and configuration.

This chapter covers the following topics:

- About Information Publisher
- Out-of-Box Report Definitions
- Custom Reports
- Scheduling Reports
- Sharing Reports

11.1 About Information Publisher

Information Publisher provides powerful reporting and publishing capability. Information Publisher reports present an intuitive interface to critical decision-making information stored in the Management Repository while ensuring the security of this information by taking advantage of Enterprise Manager's security and access control.

Information Publisher's intuitive user-interface allows you to create and publish reports with little effort. The key benefits of using Information Publisher are:

- Provides a framework for creating content-rich, well-formatted HTML reports based on Management Repository data.

- Out-of-box reports let you start generating reports immediately without any system configuration or setup.
- Ability to schedule automatic generation of reports and store scheduled copies and/or e-mail them to intended audiences.
- Ability for Enterprise Manager administrators to share reports with the entire business community: executives, customers, and other Enterprise Manager administrators.

Information Publisher provides you with a feature-rich framework that is your central information source for your enterprise.

11.2 Out-of-Box Report Definitions

The focal point of Information Publisher is the report definition. A report definition tells the reporting framework how to generate a specific report by defining report properties such as report content, user access, and scheduling of report generation.

Information Publisher comes with a comprehensive library of predefined report definitions, allowing you to generate fully formatted HTML reports presenting critical operations and business information without any additional configuration or setup. .

Generating this HTML report involved three simple steps:

Step 1: Click **Availability History** (Group) in the report definition list.

Step 2: Select the group for which you want to run the report.

Step 3: Click **Continue** to generate the fully-formed report.

Supplied report definitions are organized by functional category with each category covering key areas.

To access the Information Publisher home page, from the **Enterprise** menu, choose **Reports** and then **Information Publisher**.

11.3 Custom Reports

Although the predefined report definitions that come with Information Publisher cover the most common reporting needs, you may want to create specialized reports. If a predefined report comes close to meeting your information requirements, but not quite, you can use Information Publisher's Create Like function to create a new report definition based on one of the existing reports definitions.

11.3.1 Creating Custom Reports

To create custom reports:

1. Choose whether to modify an existing report definition or start from scratch. If an existing report definition closely matches your needs, it is easy to customize it by using the Create Like function.
2. Specify name, category, and sub-category. Cloud Control provides default categories and sub-categories that are used for out-of-box reports. However, you can categorize custom reports in any way you like.
3. Specify any time-period and/or target parameters. The report viewer will be prompted for these parameters while viewing the report.

4. Add reporting elements. Reporting elements are pre-defined content building blocks, that allow you to add a variety of information to your report. Some examples of reporting elements are charts, tables, and images.
5. Customize the report layout. Once you have assembled the reporting elements, you can customize the layout of the report.

11.3.2 Report Parameters

By declaring report parameters, you allow the user to control what data is shown in the report. There are two types of parameters: target and time-period.

Example: If you are defining a report that will be used to diagnose a problem (such as a memory consumption report), the viewer will be able to see information for their target of interest.

By specifying the time-period parameter, the viewer will be able to analyze historical data for their period of interest.

Analyzing Historical Data

Information Publisher allows you to view reports for a variety of time-periods:

- Last 24 Hours/ 7 Days/ 31 Days
- Previous X Days/ Weeks/ Months/ Years (calendar units)
- This Week/ This Month/ This Year (this week so far)
- Any custom date range.

11.3.3 Report Elements

Report elements are the building blocks of a report definition. In general, report elements take parameters to generate viewable information. For example, the Chart from SQL element takes a SQL query to extract data from the Management Repository and a parameter specifying whether to display the data in the form of a pie, bar, or line chart. Report elements let you "assemble" a custom report definition using the Information Publisher user interface.

Information Publisher provides a variety of reporting elements. Generic reporting elements allow you to display any desired information, in the form of charts, tables or images. For example, you can include your corporate Logo, with a link to your corporate Web site. Monitoring elements show monitoring information, such as availability and alerts for managed targets. Service Level Reporting elements show availability, performance, usage and achieved service levels, allowing you to track compliance with Service Level Agreements, as well as share information about achieved service levels with your customers and business executives.

11.4 Scheduling Reports

Enterprise manager allows you to view reports interactively and/or schedule generation of reports on a flexible schedule. For example, you might want to generate an "Inventory Snapshot" report of all of the servers in your environment every day at midnight.

11.4.1 Flexible Schedules

Cloud Control provides the following scheduling options:

- One-time report generation either immediately or at any point in the future
- Periodic report generation
  - Frequency: Any number of Minutes/ Hours/ Days/ Weeks/ Months/ Years
  - You can generate copies indefinitely or until a specific date in the future.

11.4.2 Storing and Purging Report Copies

Enterprise Manager allows you to store any number of scheduled copies for future reference.

You can delete each stored copy manually or you can set up automated purging based on either the number of stored copies or based on retention time. For example, you can have Enterprise Manager purge all reports that are more than 90 days old.

11.4.3 E-mailing Reports

You can choose for scheduled reports to be e-mailed to any number of recipients. You can specify reply-to address and subject of the e-mail.

11.5 Sharing Reports

Information Publisher facilitates easy report sharing with the entire user community. Enterprise Manager administrators can share reports with other administrators and roles. However, there may be cases when you need to share reports with non-Enterprise Manager administrators, such as customers and/or business executives. To facilitate information sharing with these users, Enterprise Manager renders a separate reporting Web site that does not require user authentication.

Note: To ensure that no sensitive information is compromised, only Enterprise Manager administrators with a special system privilege are allowed to publish reports to the Enterprise Manager reports Web site.

Information Publisher honors Enterprise Manager roles and privileges, ensuring that only Enterprise Manager administrators can create reports on the information they are allowed to see.

When sharing reports, administrators have an option of allowing report viewers to see the report with the owner's privileges. For example, as a system administrator you might want to share a host's performance information with a DBA using your server, but you do not want to grant the DBA any privileges on your host target. In this case, you could create a host performance report, and allow the DBA to view it with your privileges. This way, they only see the information you want them to see, without having access to the host homepage.

Part IV

Accessing Enterpise Manager via Mobile Devices

This section contains the following chapter:

- Cloud Control Mobile

12

Cloud Control Mobile

This chapter describes how to set up and use an iDevice to remotely connect to Enterprise Manager for the purpose of managing incidents and problems in Cloud Control.

The chapter contains the following sections:

Reviewing System Requirements
Performing Initial Setup
Connecting the First Time
Encountering the Login Screen
Managing Settings
Using Cloud Control Mobile in Incident Manager
Working in Cloud Control Mobile
Learning Tips and Tricks

12.1 Reviewing System Requirements

Cloud Control Mobile can be deployed to the following Enterprise Manager Cloud Control 12c minimum configurations:

- A new installation of the Enterprise Manager Cloud Control 12c (12.1.0.1) patched release (released February 2012 or later)
- An existing Enterprise Manager Cloud Control 12c (12.1.0.1) installation with Bundle Patch 1 (BP1) applied

Additional requirements are as follows:

- iDevice (iPhone, iPod touch, or iPad) running iOS 4.2.x or later
- A Wi-Fi or 3G connection to a network that has access to Enterprise Manager (Cloud Control Mobile supports connections over VPN)
- An Apple account with which to download the app from the iTunes App Store

12.2 Performing Initial Setup

Initial setup involves the following tasks:

- Connect to a Wi-Fi or 3G network
- Install and configure VPN
- Download the Cloud Control Mobile app and sync with your iDevice
- Add a Cloud Control URL to connect to the installed Enterprise Manager

12.3 Connecting the First Time

When you first install the app, there is no default Enterprise Manager connection, so you must supply a Cloud Control URL. There are two ways to do this:

- Use the iDevice Settings app
- Launch the Cloud Control Mobile app

In either case, first log in to VPN if required before proceeding with the instructions below. Without the VPN connection, the login screen will not appear.

iDevice Settings

Define a default Cloud Control URL as follows:

1. Tap the Settings icon on the Home screen.
2. Tap Cloud Control in the apps list.
3. On the Cloud Control screen, enter a name to identify the site and type the Cloud Control URL to which to connect. The URL should be of the form:

   ```
   https://www.yoursite.com/em
   ```

4. Tap **Settings** to store the information and return to the list of apps.

You can now launch the Cloud Control Mobile app to log in.

Initial App Launch

Define a default Cloud Control URL as follows:

1. Tap the Cloud Control Mobile icon on the Home screen.
2. On the Add Site screen, enter a name to identify the site and type the Cloud Control URL to which to connect. The URL should be of the form:

   ```
   https://www.yoursite.com/em
   ```

   Before you can type in the name field you may first have to clear the field by tapping the X at the right.
3. Tap **Done** to store the information.
4. Tap **Done** on the Sites screen. Note that you also have the option to add additional sites before exiting this screen.
5. Tap **Settings** on the Sites navigation bar to close the Sites list screen.
6. Tap **Save** on the Settings navigation bar to complete the action.

Proceed with the login.

12.4 Encountering the Login Screen

You encounter the login screen under the following conditions:

- After supplying a default Cloud Control URL upon initial launch
- Anytime you subsequently launch the app
- When you change the default site
- When you log out

Tap the **Settings** icon to see a list of sites or to change the default login site. See Section 12.5, "Managing Settings" for more information.

Specify your credentials and tap **Login**; the Incident Manager opens, displaying the my open incidents and problems view.

If your Enterprise Manager installation does not have a site certificate signed by a valid certificate authority, an alert overlays the login screen noting an invalid certificate. You have the option to continue with the login or change the URL to which you are trying to connect.

Note: If your installed Enterprise Manager uses single sign-on, the SSO process supplants site login. Upon completion of single sign-on, the workflow proceeds to the my open incidents and problems view.

12.5 Managing Settings

Cloud Control Mobile has its own settings interface apart from the iDevice Settings app that you use to manage all apps.

In managing Cloud Control Mobile app settings, you perform the following actions:

- Add a site
- Edit a site
- Delete a site
- Change the default site

Each action starts with the same basic steps:

1. Tap the actions icon on the right of the navigation bar.
2. Tap **Settings** in the action sheet.
3. Tap the Edit Sites table row.
4. Tap **Edit**. The Sites management screen appears:

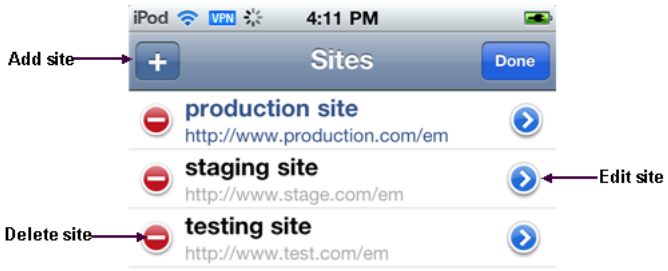

Then proceed as described below for each individual action.

Add a Site

1. Tap the + sign on the left of the Sites navigation bar.
2. Type a name and a URL for the site to be added.
3. Tap **Done** on the Add Site navigation bar to close the screen.
4. Tap **Done** on the Sites navigation bar to exit edit mode.
5. Tap **Settings** on the Sites navigation bar to close the Sites list screen.

6. Tap **Save** on the Settings navigation bar to complete the action.

Edit a Site

1. Tap the blue arrow to the right of the URL to be edited.
2. Change the values as appropriate.
3. Tap **Done** on the Edit Site navigation bar to close the screen.
4. Tap **Done** on the Sites navigation bar to exit edit mode.
5. Tap **Settings** on the Sites navigation bar to close the Sites list screen.
6. Tap **Save** on the Sites navigation bar to complete the action.

Delete a Site

1. Tap the red circle to the left of the URL to be deleted.
2. Tap **Delete** that appears on the right in the table row.
3. Tap **Done** on the Sites navigation bar to close the screen.
4. Tap **Settings** on the Sites navigation bar to close the Sites list screen.
5. Tap **Save** on the Settings navigation bar to complete the action.

Change the Default Site

1. Tap the site to be the new default. The check mark to the right in the table row confirms your selection.
2. Tap **Settings** to close the Sites list screen.
3. Tap **Save** to complete the action.
4. After a brief moment, the Login screen appears. Specify credentials to log in to the new site.

Note that you also can change the default site in iDevice Settings for Cloud Control.

12.6 Using Cloud Control Mobile in Incident Manager

Connecting to Cloud Control remotely, you can do the following in Incident Manager:

- View your open incidents and problems; drill down to incident and problem details, including associated updates and events
- See the list of incidents for a given problem; link to these incidents and their details
- Acknowledge incidents and problems
- Manage incident and problem workflow for better tracking (change status, assign owner, escalate, and so forth)
- See who has been notified about an issue and what comments have been added by administrators

In addition, The FAQ that follows may help in understanding differences, subtleties, and nuances between the mobile app and its browser-based counterpart.

How do I view more issues?

The app displays five rows at a time. Use the next and previous controls at the bottom of the display to scroll the list. Or tap the number range itself (1 - 5) to pick from a list of increments.

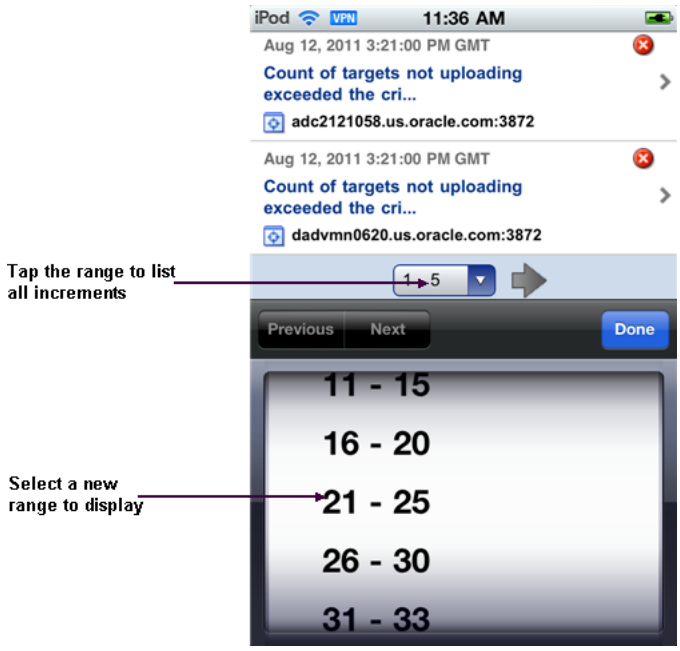

How do I view issue details?

Simply tap the line that identifies the incident or problem. On the Details screen, you can continue to drill down to updates and events as well as expand the summary description.

Are the views the same in the mobile app as in the browser-based product?

Yes, except for event-related views (standard or custom), which are not available in the app. Any custom views created in the browser-based product augment the list of standard views.

Can I refresh the view?

Yes, just tap the refresh icon on the left of the navigation bar. When you do, the timestamp reflects the date and time of the refresh.

Can I view target details?

Yes, first drill down to incident details, then tap the target name to jump to target details. Tap **Incident** to return to the incident details.

Can I set search criteria or create a custom view?

No, you cannot set search criteria or create a custom view in the mobile app, but you can create and manage views in the browser-based product, which are then available in the mobile app.

Can I invoke the incident rules feature?

No, but you can receive notifications generated by incident rules on your mobile device, provided your Enterprise Manager account has the appropriate notification preferences.

Can I connect to My Oracle Support?

If a problem has an assigned SR number, you can click the number to view the SR details in the My Oracle Support (MOS) Mobile app.

Can I access guided resolution information and diagnostics?

No.

Do all iDevices work the same way with the mobile app?

Pretty much. The one difference you will note on an iPad is that if you tap a link to an issue or a target in an external source such as Safari or an e-mail message, Safari launches, pointing to the relevant page in the browser-based product, where you are greeted with the usual Cloud Control login screen. With the other devices, tapping a link in an external source launches the mobile app.

12.7 Working in Cloud Control Mobile

This section covers the following operational tasks in Cloud Control Mobile:

- Viewing Incidents and Problems
- Changing Views
- Performing Actions

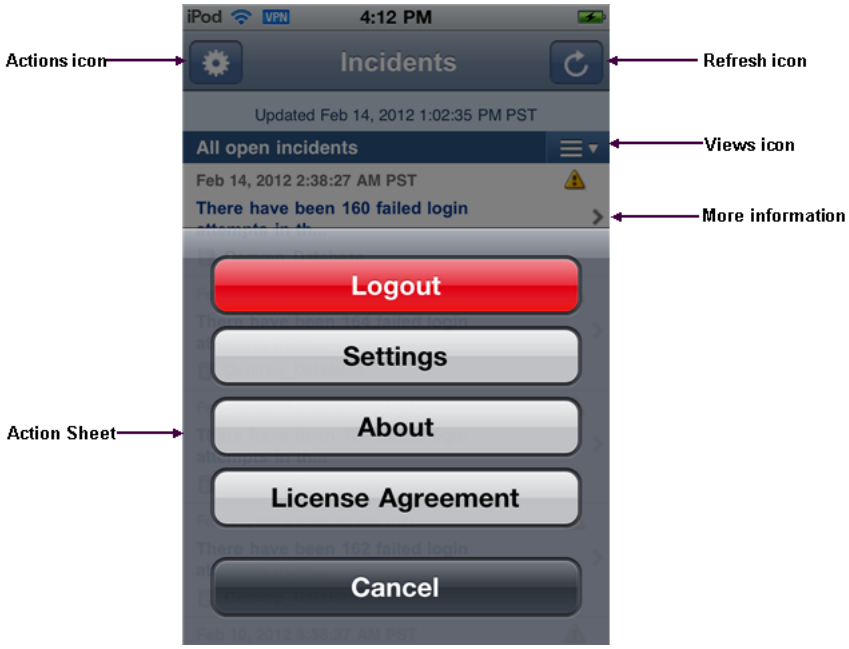

12.7.1 Viewing Incidents and Problems

Although navigation is intuitive, the following sections offer guidance on viewing incidents and problems. As the interactions are slightly different, there is a separate section for each type of issue.

Viewing Incidents

Use the following guidelines as you view incidents in the list:

- An arrow on the right indicates availability of additional information.
- Tap anywhere in the incident row to drill down to incident details.

- The summary appears at the top. As summaries can be lengthy, you may need to tap the opening lines of the summary to view the complete summary. Tap **Incident** to return to incident details.
- Problem ID in incident details is a link to problem details. If you follow the link, tap **Incidents** there to return to the starting point; that is, the original list view where you first opened the incident.
- In the incident details view, target name is a link to target details. Tap **Incident** there to return to incident details.
- Tracking information appears below target name in the incident details view.
- Scroll down in incident details and tap **All Updates** to see the equivalent of the **Updates** tab in the browser-based product. Tap **Incident** there to return to incident details.
- Scroll further and tap **Event List** to see the equivalent of the **Events** tab in the browser-based product. Tap **Incident** there to return to incident details.

Viewing Problems

Use the following guidelines as you view problems in the list:

- An arrow on the right indicates availability of additional information.
- Tap anywhere in the problem row to drill down to problem details.
- The summary appears at the top. As summaries can be lengthy, you may need to tap the opening lines of the summary to view the complete summary. Tap **Problem** to return to problem details.
- If the problem has an SR number assigned, tap the number to log in to the My Oracle Support (MOS) Mobile app using your MOS credentials. You can then view the SR details and take appropriate action. Go to the Home screen and tap the Cloud Control Mobile app icon to return to problem details.
- In the problem details view, target name is a link to target details. Tap **Problem** there to return to problem details.
- Tracking information appears below target name in the problem details view.
- Scroll down in problem details and tap **All Updates** to see the equivalent of the **Updates** tab in the browser-based product. Tap **Problem** there to return to problem details.
- Scroll further to see additional details such as first and last incident and number of incidents in which the problem has occurred.
- Scroll further and tap **Incident List** to see the equivalent of the **Incidents** tab in the browser-based product. Tap **Problem** there to return to problem details.
- Each incident summary in the list is a link to the details of the incident. If you follow the link, tap **Incidents** there to return to the starting point; that is, the original list view where you first opened the issue.

12.7.2 Changing Views

When you first log in, the my open incidents and problems view appears by default. You can change the view as follows:

1. Open the Views menu by tapping the views icon (three horizontal lines to the right of the current views title).

2. Tap the view you want. The check mark to the right confirms your selection.
3. Tap **Incidents** to display the new view.

12.7.3 Performing Actions

You can perform the following actions while viewing incident or problem details:

- Acknowledge the issue
- Manage the workflow of the issue

To acknowledge an incident or problem:

1. While viewing the details, tap **Actions**.
2. Tap **Acknowledge** in the action sheet.

   A message confirms the update on the Details screen.

Note that the Acknowledge action may not appear in the action sheet for a variety of reasons; for example, the issue has already been acknowledged or closed, or you do not have the right permissions to acknowledge the issue.

To manage the workflow of an incident or problem for better tracking:

1. While viewing the details, tap **Actions**.
2. Tap **Manage** in the action sheet.
3. Complete the Manage dialog the same as you would in the browser-based product. The only thing missing is the ability to add styles and formatting to the comment.
4. Tap **Save** to complete the action.

12.8 Learning Tips and Tricks

Use a touch-and-hold gesture at any time within the app to display on the bottom of the screen the current site to which you are logged in or about to log in. If the site name is unavailable, the URL appears. With the site identity displayed, tap the information icon on the right to access the Settings screen, where you can manage your sites and change the default site. Repeat the touch-and-hold gesture to remove the site display from the bottom of the screen.

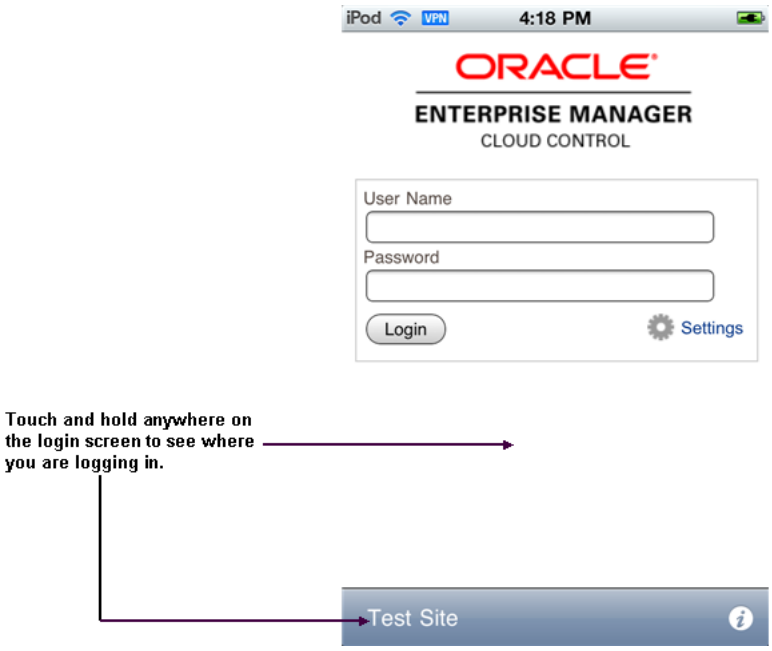

If you have logged out of the app and find that you are stuck on a page trying to return to Cloud Control Mobile, you have a couple of options to resolve the issue:

- Open iDevice Settings and change the Cloud Control default URL.
- Force quit the app and restart Cloud Control Mobile. For example, press and hold the On/Off button on top of the device until the power off slider appears, and then press the Home button until the app closes.

Part V

Administering Cloud Control

This section contains the following chapters:

- Personalizing Cloud Control
- Maintaining Enterprise Manager
- Updating Cloud Control
- Patching Enterprise Manager
- Starting and Stopping Enterprise Manager Compone nts
- Locating and Configuring Enterprise Manager Log Files
- Maintaining and Troubleshooting the Management Repository

13

Personalizing Cloud Control

You can personalize the page layout and data displayed in certain Cloud Control pages, including target home pages such as the Oracle WebLogic Server target home page. The changes you make are persisted for the currently logged in user, enabling you to create customized consoles for monitoring various target types.

Personalization support provided in the current release allows you to:

- Customize the layout of regions on a page
- Add a region to, or remove a region, from a page
- Specify what data should be displayed within each region
- Set your own homepage

Note that not all pages in Cloud Control can be personalized. The page edit mode will only be enabled for those pages or page regions that can be modified.

This chapter contains the following sections:

- Personalizing a Cloud Control Page
- Customizing a Region
- Setting Your Homepage

13.1 Personalizing a Cloud Control Page

Pages in Cloud Control are laid out in a columnar format. Each column contains one or more *regions*, each of which contains data rendered as a bar chart, graph or other visual component.

You can modify the layout of columns within a page, as well as select the regions to display within each column, enabling you to personalize how the data on a page is arranged and displayed.

To personalize a page:

1. Navigate to the page you want to personalize.
2. Do one of the following:
   - Select **Personalize Page** from the menu item that displays the username of the currently logged-in user, just to the left of the Log Out menu item. In Figure 13–1, the menu item displays the SYSMAN user name.

Figure 13–1 Personalize Page Menu

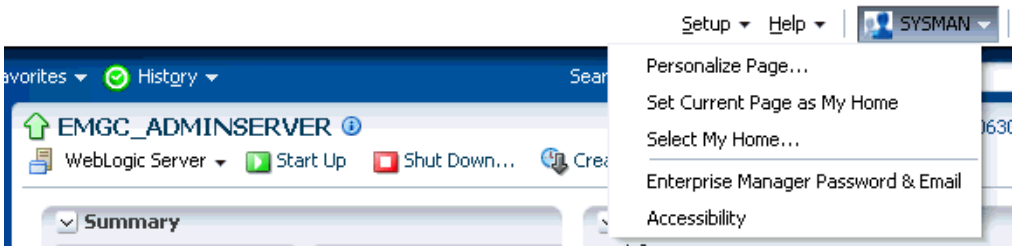

- Or, click the "Personalize Page" icon on the right-hand side of the page, shown just above the "Page Refreshed" time stamp, as shown in Figure 13–2.

Figure 13–2 Personalize Page Icon

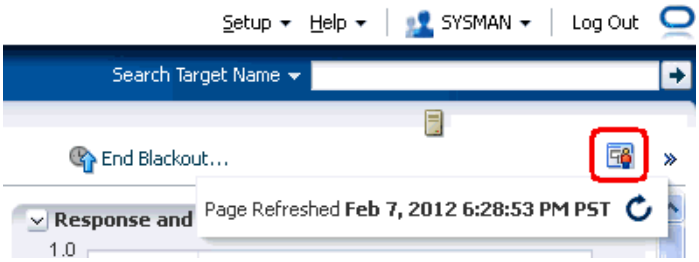

Note that the menu item or icon will only be enabled if the page you are currently on can be personalized.

3. You are now in page edit mode. Click the **Change Layout** button. A graphical menu of column layout options opens.
4. Select the column layout you want to use.
5. Next, add a region to each column. Click the **Add Content** button for a specific column. The Resource Catalog, which contains available components used to display data, opens.
6. Select a region, then click **Add** to add it to the column. Note that you can "stack" regions on top of one another.
7. Once a region has been added to a column, you can:
   - Customize the region. See Section 13.2, "Customizing a Region" for details.
   - Click the **View Actions** icon to move the region up or down within the column.
   - Drag the region from one column to another.
8. Click **Close** to save your changes.

13.2 Customizing a Region

A *region* contains business data rendered as a bar chart, graph or other visual component. You can select the component to display within a specific region.

To customize a region within a page column:

1. Navigate to the page containing the region you want to customize and enable the page editing mode as described in Section 13.1, "Personalizing a Cloud Control Page".
2. Click the "ratchet" icon next to the "X" icon within a region, as shown in Figure 13–3. Note that the icon will only be enabled if the region can be customized.

Figure 13–3 Customize Region Icon

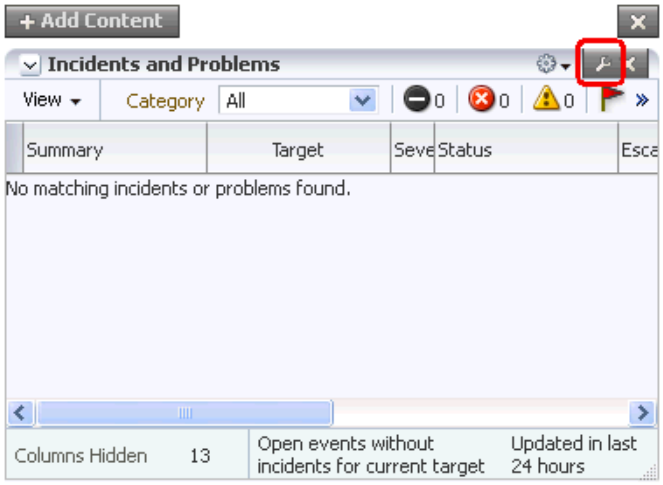

For most resources, you will specify the target host from which to collect data.

Other configurable parameters and customization options vary between regions. When you click the icon, a dialog opens to enable you to specify parameters, such as target type, target name and metric name.

3. If at a later time you want to remove the region from the page, click the "X" icon in the region.
4. Click **Close** to save your changes.

13.3 Setting Your Homepage

Cloud Control allows you to choose the page that will serve as your homepage - the first page you see after logging in to Cloud Control. You can either:

- Choose your own page, such as a target homepage that you view frequently or have customized to suit your specific needs
- Select from a pre-designed homepage templates created for specific types of Cloud Control users

Choosing Your Own Homepage

1. Navigate to the page you want to set as your homepage.
2. Select **Set Current Page As My Home** from the menu item that displays the username of the currently logged-in user, just to the left of the Log Out menu item, as shown in Figure 13–1.

Selecting a Pre-Designed Homepage

1. Select **Select My Home** from the menu item that displays the username of the currently logged-in user, just to the left of the Log Out menu item, as shown in Figure 13–1.
2. Click the **Preview** button to preview a page design template you are interested in.
3. Click the **Select As My Home** button to select a template as your homepage. Once you have selected a page, you can customize it to suit your needs.

De-selecting Your Homepage

Your homepage is saved as a "favorite" page. To de-select your current homepage:

1. From the **Favorites** menu, select **Manage Favorites**.
2. Select your homepage from the list, then click the**Remove Favorite** button.
3. Click **OK** when finished.

14

Maintaining Enterprise Manager

Enterprise Manager provides extensive monitoring and management capabilities for various Oracle and non-Oracle products. Used to manage your heterogeneous IT infrastructure, Enterprise Manager plays an integral role in monitoring and maintaining the health of your IT resources. It is therefore essential to make sure Enterprise Manager itself is operating at peak efficiency.

To help you maintain your Enterprise Manager installation, a variety of enhanced self-monitoring and diagnostic functionality is available from the Enterprise Manager console. These functions are designed to help you understand and monitor various components of Enterprise Manger, monitor/measure the quality of services Enterprise Manager provides, diagnose failures quickly, and manage Agents more easily.

This chapter covers the following topics:

- Overview: Managing the Manager
- Management Services and Repository
- Viewing Enterprise Manager Topology and Charts
- Viewing Enterprise Manager Services
- Controlling and Configuring Management Agents

14.1 Overview: Managing the Manager

Although Enterprise Manager functions as a single entity to manage your IT infrastructure, in reality it is composed of multiple components working in concert to provide a complete management framework from a functional standpoint. Beginning with Enterprise Manager 12c, the functions used to ensure reliability and performance for your monitored targets can be used to maintain Enterprise Manager itself. All major components of Enterprise Manager have been grouped into a single system. A special set of services has been created (based on the system) to model Enterprise Manager functions.

Management Features

- Topology view that allows you to see all major components of Enterprise Manager and their current status.
- Enterprise Manager dashboard displaying the overall health of Enterprise Manager.
- Full control of the Agent directly from the Enterprise Manager console. Functions include:
  - View/edit Agent configuration properties.

- View Agent(s) configuration history and compare the results against other Agents.
- Perform Agent control operations (start/stop/secure).

14.2 Management Services and Repository

The Management Services and Repository home page provides a detailed overview of the OMS and Repository.

Accessing the Management Services and Repository Home Page

From the **Setup** menu, click **Management Services and Repository**.

Figure 14–1 Management Services and Repository Home Page

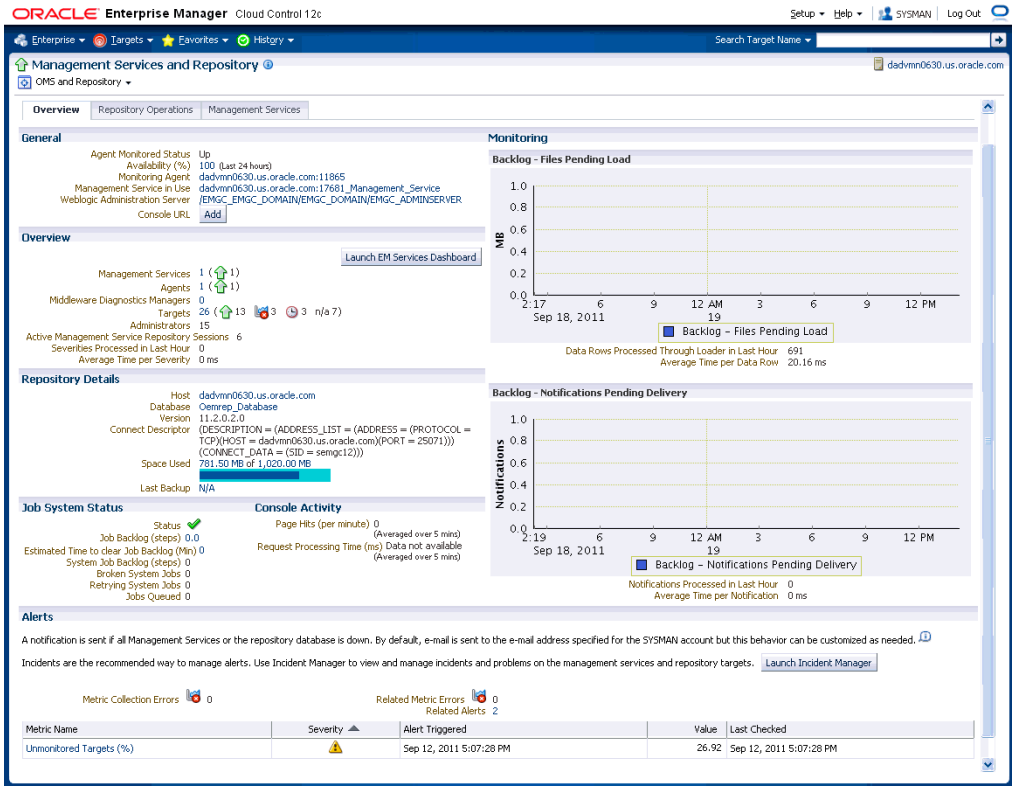

Overview Page

The Overview page displays a summary of the current status of the OMS and Repository. The Overview page also provides details on the status of the Job system and the console activity and two monitoring charts that indicate the backlog in terms of the files pending and the backlog on notifications pending delivery.

Each region provides specific information on the various operational areas of the OMS and Repository.

- **General**: provides the Agent monitored status, availability and the host details of the Management Service in use, the monitoring Agent and the WebLogic administration server.

- **Overview**: Displays the number and breakdown in term of status for the OMS, Agents, targets, active OMS repository sessions and severities detected within the past hour.
- **Repository Details**: Provides physical information about the Management Repository and the host on which the database is located.
- **Alerts**: Provides details on the metric errors recorded and when an alert was triggered. In-context links to Incident Manager are also provided.

Repository Operations Page

The Repository Operations page provides you with an overview of the status and performance of the Repository DBMS Jobs that handle part of Enterprise Manager's maintenance and monitoring functionality. These DBMS jobs run within the Management Repository and require no user input. Charts showing the **Data Upload** to **Repository** and the **Backlog** in Repository collection are provided The **Scheduler Status** region provides the status of the scheduler and the number of Job Queue Processes.

Figure 14–2 Repository Operations Page

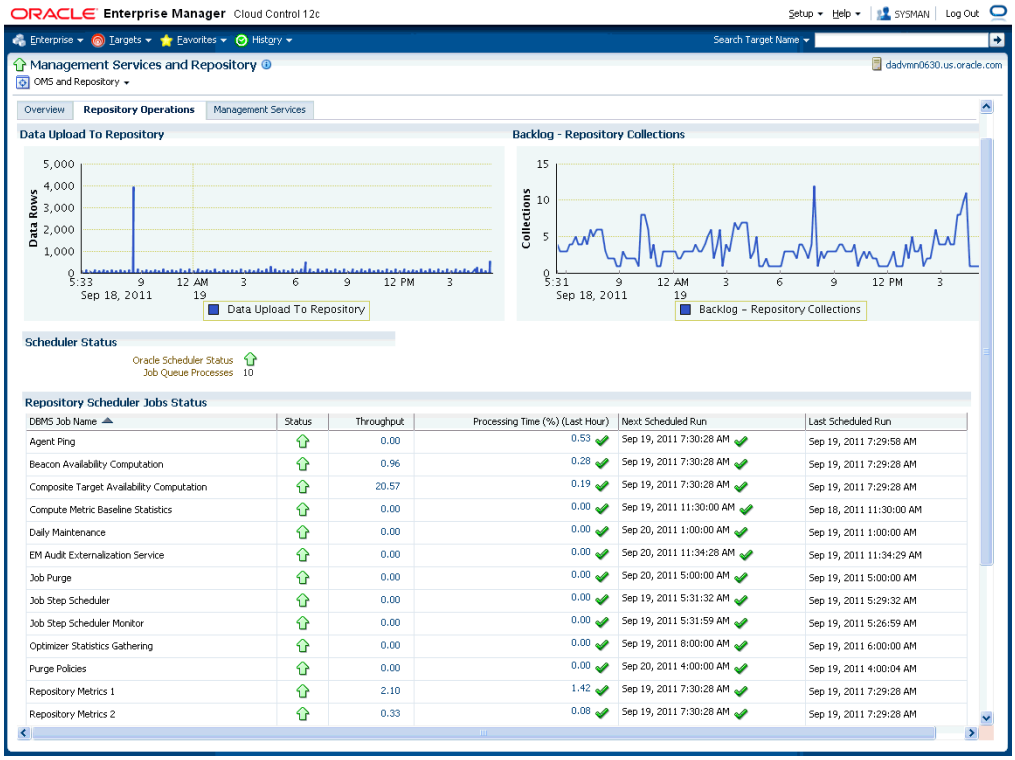

The **Repository Scheduler Jobs Status** region provides details of the DBMS Jobs regarding their status, throughput, processing time, the next scheduled run and the last scheduled run.

Use the Repository Operations page to view the performance of the Repository DBMS jobs. If you want more information, clicking a link brings you to a metrics detail page.

To determine how well the Management Repository is handling its share of the Enterprise Manager functionality, view the Throughput per second and Processing

Time Percent (Last Hour) columns. If the Processing Time Percent (Last Hour) is large and the Throughput is low, there may be problems in that area of management.

Note: Processing Time (%) Last Hour may exceed 100% if a job runs continuously for more than an hour. For example, if you see 125.00 in this column, it means that the job ran for 75 minutes (125 % of one hour).

Management Services Page

The **Management Services** page lists the names of the Management Services and their status. When an OMS is decommissioned, status will be shown as down and delete button will be displayed, allowing you to remove the decommissioned OMS from EM.

Figure 14–3 Management Services Page

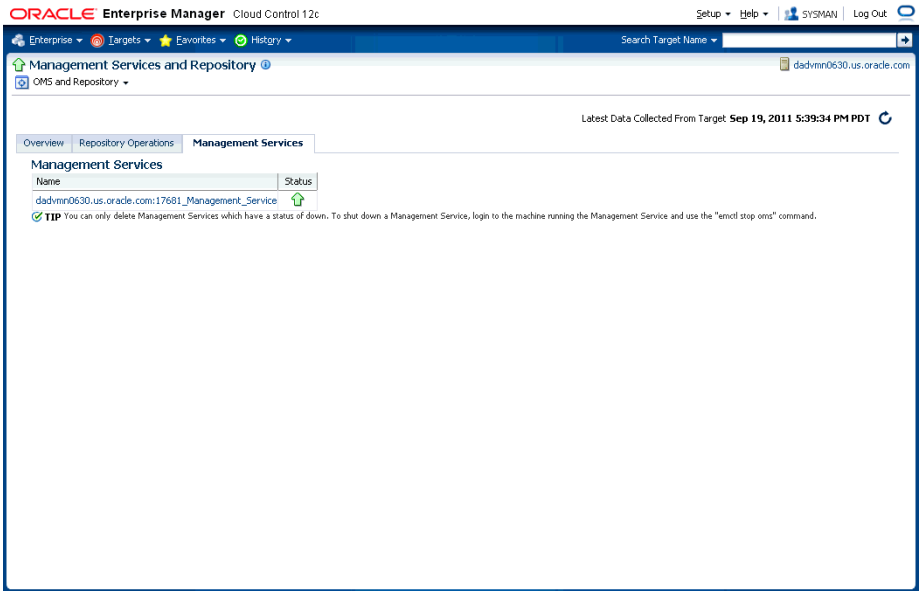

By clicking on the individual Management Service, you can drill down to that OMS' target home page for explicit information and status.

14.3 Viewing Enterprise Manager Topology and Charts

The Enterprise Manager Topology page provides a graphical representation of the Enterprise Manager infrastructure components and their association. Each node in the hierarchy displays key information about the member type, the host on which it resides, and the number of incidents, if any. The incident icons on each of the nodes expand to display a global view of current status for each node in the hierarchy.

Note: In order for the Enterprise Manager repository database to appear in the Topology page, you must first manually discover the database. Manual discovery is also required in order to have the database's metric data (Database Time (centiseconds per second)) displayed in the charts.

Accessing the Enterprise Manager Topology

1. From the Setup menu, choose Management Services and Repository.
2. On the Management Services and Repository page, click on the **OMS and Repository** drop-down menu.
3. Choose **Members** and then **Topology**.

Figure 14–4 Enterprise Manager Topology

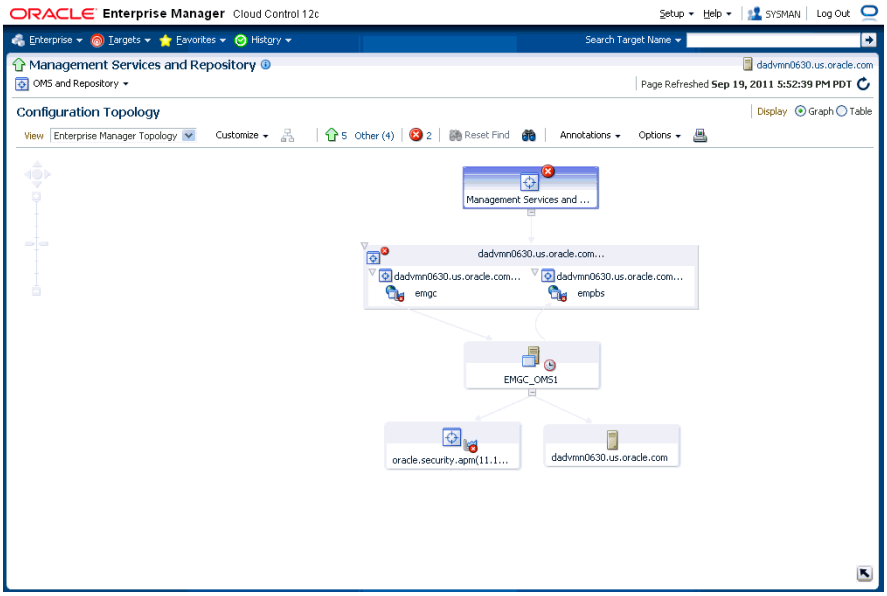

Enterprise Manager Charts

The Enterprise Manager Charts page displays eight charts representing key areas that together indicate the overall health of Enterprise Manager. These are Overall Files Pending Load -Agent, Job Step Backlog, Job Step Throughput (per second), Request Processing Time (ms), Database Time (centiseconds per second), CPU Utilization (%), Pages Paged-in (per second), Pages Paged-out (per second). Data can be viewed for the Last 24 hours, last 7 days or last 31 days.

Accessing the Enterprise Manager Charts

1. From the Setup menu, select **Management Services and Repository**.
2. On the Management Services and Repository page, click on the **OMS and Repository** drop-down menu.
3. Select **Monitoring**, the select **Charts**.

Figure 14–5 Enterprise Manager Charts

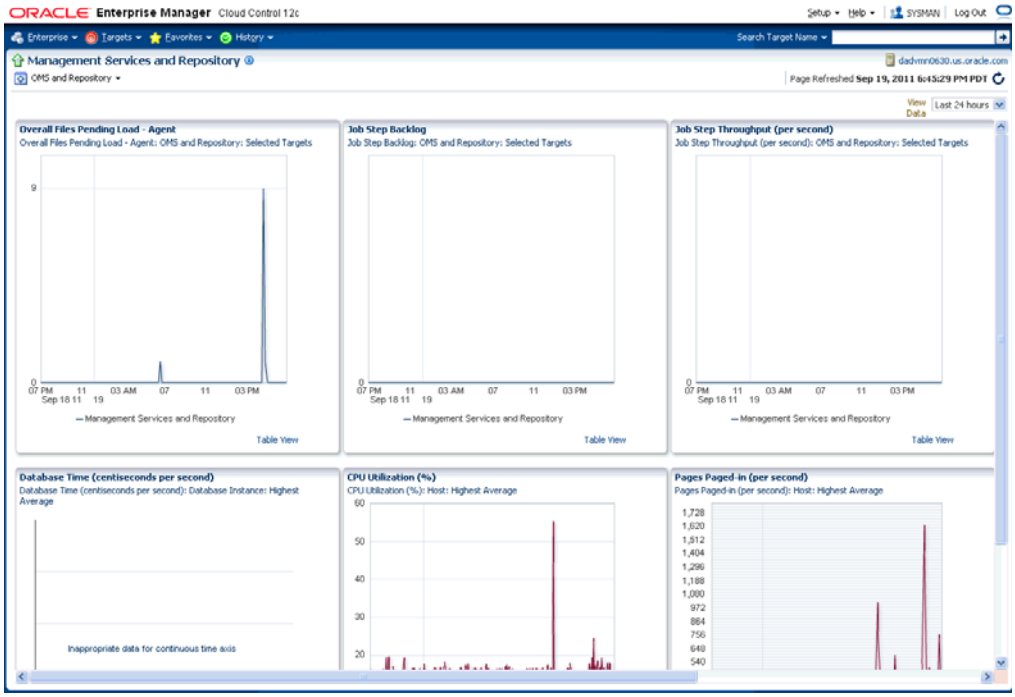

14.4 Viewing Enterprise Manager Services

The Enterprise Manager Services dashboard provides details about the two primary services in-context: Enterprise Manager Jobs Service and the Enterprise Manager Console Service. The **Status**, **Performance** (gauged by means of specific metrics), **Usage and Business Indicators** (again through specific metrics), **Component Status** and the **Service Level** (%) for the two primary services for the Last 24 hours, last 7 days and last 31 days are detailed.

Accessing the Enterprise Manager Topology

1. From the **Setup** menu, choose **Management Services and Repository**.
2. On the Management Services and Repository page, click **Launch EM Services Dashboard** in the Overview region.

Figure 14–6 Enterprise Manager Services Dashboard

14.5 Controlling and Configuring Management Agents

Beginning with Enterprise Manager Cloud Control 12c, management of Management Agents can be performed directly from the Enterprise Manager console. This provides a central point where all Management Agents for your managed targets can be compared, configured and controlled.

14.5.1 Management Agent Home Page

The Management Agent home page provides details for a single Management Agent. This page also lets you drill down for more detailed information. You can access an Agent home page by selecting it from the All Targets page.

Figure 14–7 Agent Home Page

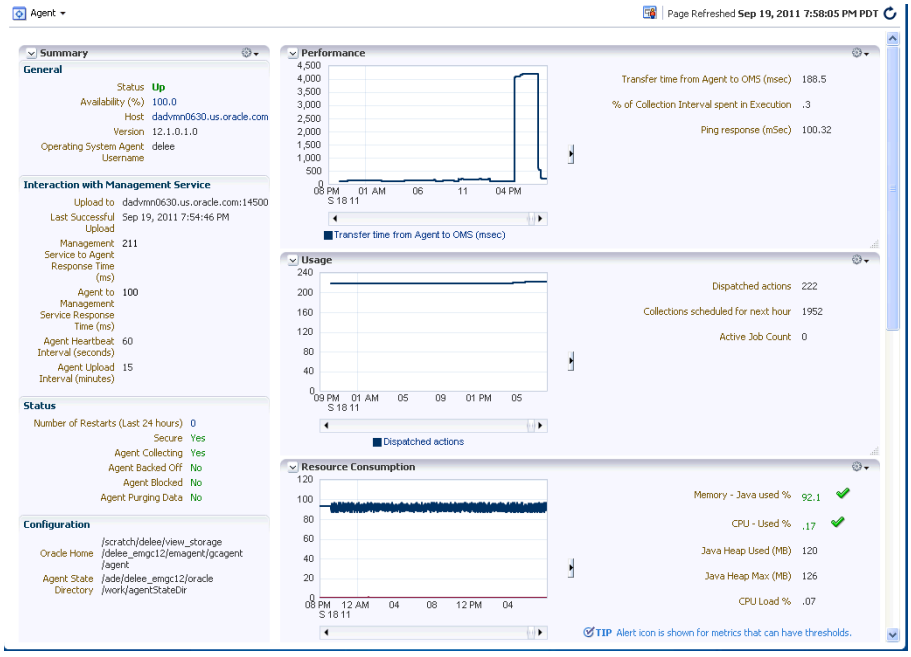

The **Summary** region provides primary details of the Management Agent such as its status and availability. The **Interaction with Management Service** region provides

details on the communication between the OMS and the Management Agent. The **Status** region provides further details on the Management Agent status such as the number of restarts, the action that the Management Agent is performing currently. The **Performance, Usage and Resource Consumption** charts provide further details on the Management Agent in graphical format. The **Incidents** region lists the incidents recorded for the Management Agent. The **Monitoring** section provides details on the targets that are being monitored by the Management Agent, metric extensions and management plug-ins deployed in the Management Agent.

14.5.2 Controlling a Single Agent

Control operations for a single Management Agent can be performed on the Agent home page for that Agent.

1. Navigate to the desired Management Agent home page.
2. From the **Agent** drop-down menu, choose **Control** and then one of the control operations (Start Up/Shut Down, or Restart)

> **Note:** You must have at least operator privileges in order to perform Agent control operations.

Figure 14–8 Control Operations from the Agent Home Page

Upon choosing any of the above control menu options, a pop-up dialog requesting the credentials of the user displays. These operations require the credentials of the OS user who owns the Management Agent. At this point, you can either choose from a previously stored username/password, preferred or named credential. You also have the option of choosing a new set of credentials which could also be saved as the preferred credential or as a named credential for future use.

When storing a named credential or creating a new set of credentials, a **Test** button is shown, thus allowing you to verify if the credential information you just entered is valid. Once you are authenticated, the chosen control operation begins and continues even if the pop-up dialog is closed. Any message of failure/success of the task is displayed in the pop-up dialog.

When choosing the Secure/Resecure/Unsecure options, you must provide the requisite Registration Password.

14.5.3 Configuring Single Management Agents

Configuration operations for a single Management Agent can be performed from the Management Agent home page. To access the Management Agent properties page:

1. Navigate to the desired Management Agent home page.
2. From the **Agent** drop-down menu, select **Properties**.

> **Note:** You must have at least Configure privileges in order to perform Agent configuration operations.

Figure 14–9 Agent Properties Page

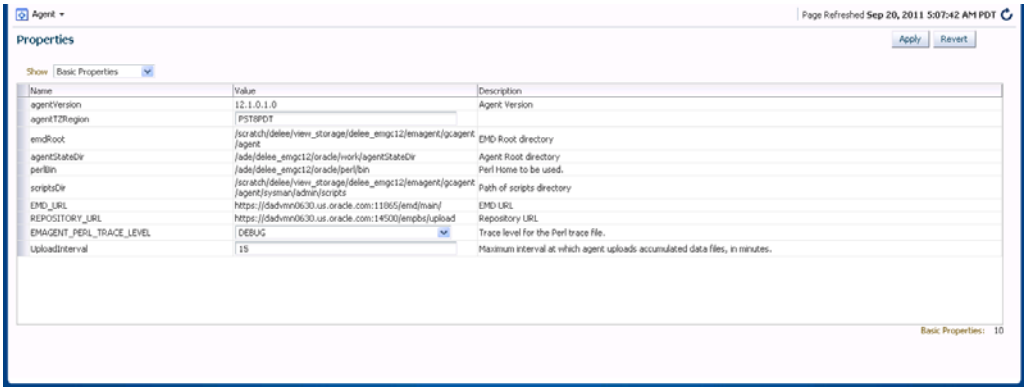

The properties on this page can be filtered to show **All Properties**, **Basic Properties**, or **Advanced Properties**. The **Basic Properties** are a simple name, value combination of a property and its value. **Advanced Properties** are also a combination of name and value but can also be grouped into categories. You must have at least *configure* privileges in order to modify the existing properties and set custom properties.

14.5.4 Controlling Multiple Management Agents

In order to perform control operations on multiple Management Agents, Enterprise Manager makes use of the Job system to automate repetitive tasks. Therefore, you must have Job privileges for controlling multiple Management Agents through a single action. To access

1. From the **Setup** menu, select **Agents**. The Management Agent setup page displays.
2. Select multiple Management Agents from the list.
3. Click one of the control operation buttons (**Start Up**/**Shut Down**/**Restart**/**Secure**/**Resecure**/**Unsecure**).

When you click on any of the control operations, you are taken to the Job creation wizard where you schedule a new job to perform the action on the selected Agents.

Figure 14–10 Multiple Agent Control Operation: Job Creation

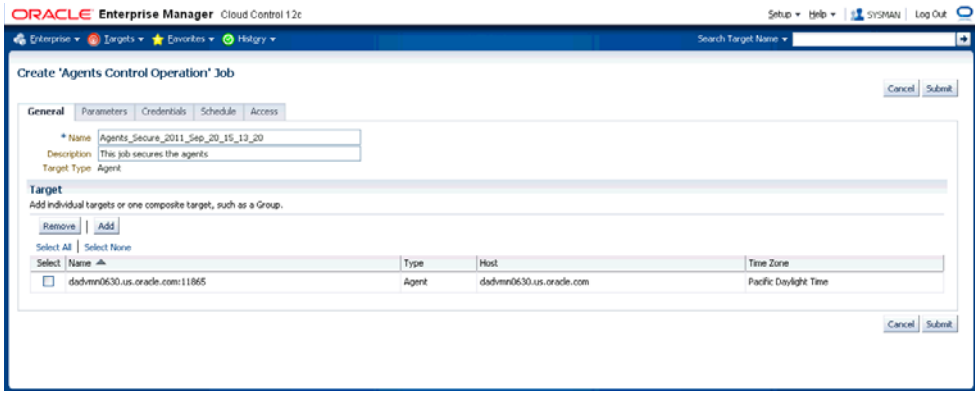

In the Jobs page, you can view the chosen Management Agents in Target section in the General tab. You can add more Management Agents by clicking the **Add** button. You then provide the parameters for the operation in the **Parameters** tab, if needed. The credentials must be specified in the **Credentials** tab where you can either choose from a previously stored username/password, preferred, or named credential. You also have the option of choosing a new set of credentials which could also be saved as the preferred credential or as a named credential for future use.

You are given the option to start the job immediately or schedule the job for a later time. At this point, you can also create a repeating job by specifying the job start time, the frequency, and the end time.

The Access tab displays the Administrator details and the access levels they have to the job. You can then add a new administrator or modify the access level to **View** or **Full**, if you have the requisite privileges.

Figure 14–11 Job Creation: Access Tab

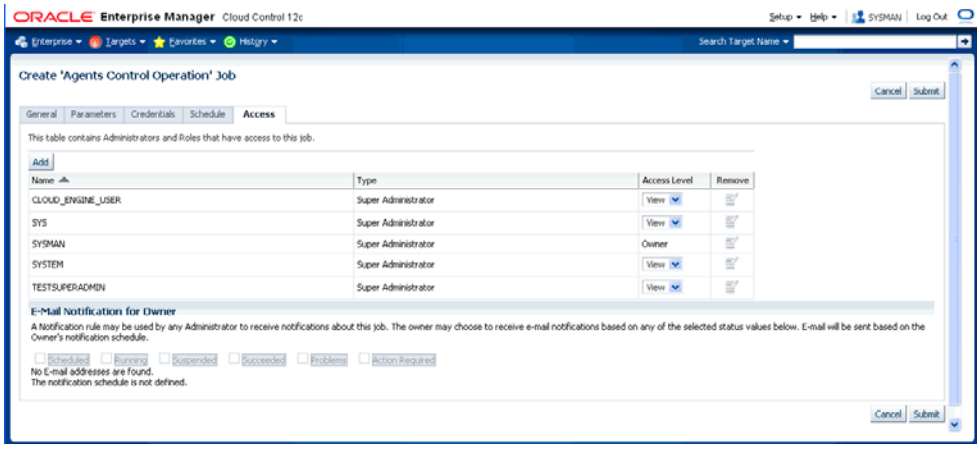

Note: Administrators with insufficient privileges can also schedule jobs for these control operations, but in this situation, the jobs will not complete successfully.

14.5.5 Configuring Multiple Agents

As with multi-Agent control operations, you can also perform Agent configuration on multiple Agents in the same way. This greatly simplifies standardizing Agent configurations across your enterprise. To access Agent properties:

1. From the **Setup** menu, select **Agents**. The Management Agent setup page displays.
2. Select multiple Management Agents from the list.
3. Click **Properties**. As with any multi-Agent operation, configuration is implemented using the Job system.

Figure 14–12 Agent

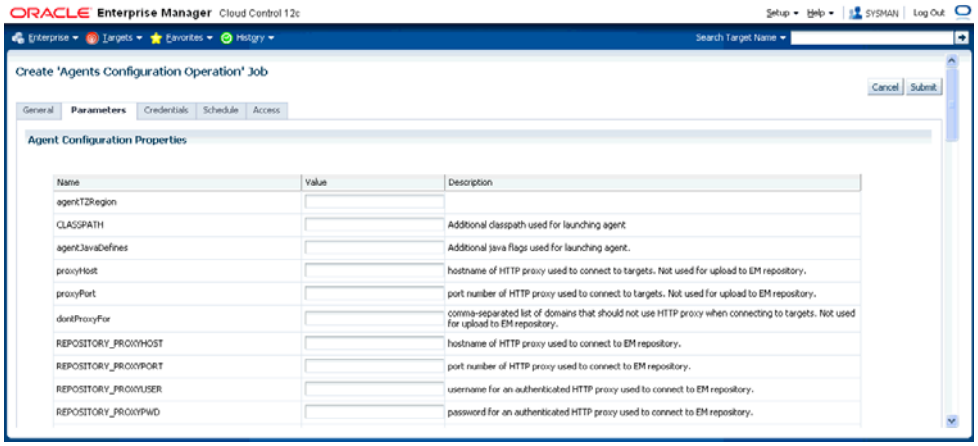

In the Jobs page, you can view the chosen Management Agents in the Target section of the General tab. You can add more Management Agents by clicking the **Add** button if necessary. In the **Parameters** tab, you provide the modified value for a particular set of properties that you want to change. You can also set a custom property for the chosen agents. No credentials are required for modifying Agent properties.

The **Access** tab displays the administrator details and the access levels they have to the job. You can then add a new administrator or modify the access level to View or Full if you have the requisite privileges.

Figure 14–13 Multi-Agent Configuration: Job Access

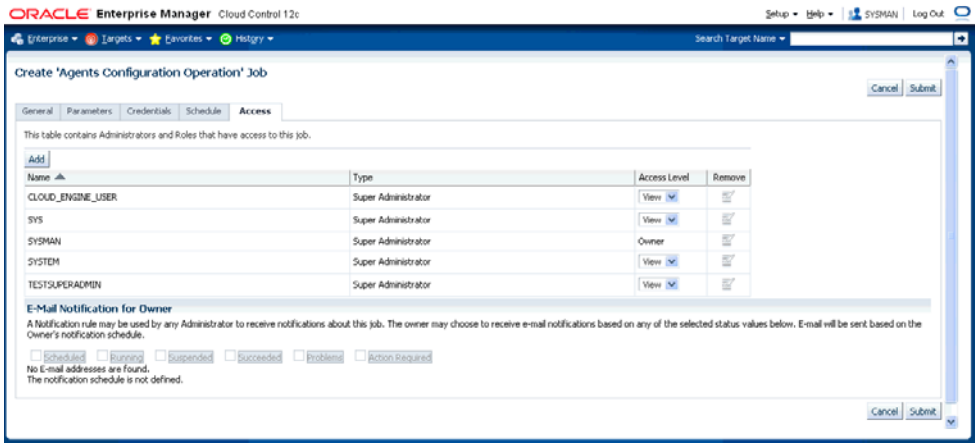

15

Updating Cloud Control

The Self Update feature allows you to expand Enterprise Manager's capabilities by updating Enterprise Manager components whenever new or updated features become available. Updated plug-ins are made available via the Enterprise Manager Store, an external site that is periodically checked by Enterprise Manager Cloud Control to obtain information about updates ready for download.

Cloud Control also provides support for plug-in and connector management. The ability to update plug-ins is particularly important because core Enterprise Manager features - such as Oracle Database management functionality - is now made available via plug-ins.

This chapter contains the following sections:

- Using Self Update
- Setting Up Self Update
- Applying an Update
- Acquiring or Updating Management Agent Software
- Deploying and Updating Plug-ins

15.1 Using Self Update

The Self Update 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 the updates. The Self Update console frees you from having to monitor multiple channels to get informed about new updates that are available from Oracle. The Self Update console automatically informs you whenever new updates are made available by Oracle. Only those updates that are applicable to your site are shown, eliminating the need to wade through unrelated updates.

15.1.1 What Can Be Updated?

Specific updates authored by Oracle that are usually bundled with specific Cloud Control releases can be updated via Self Update. Some examples are Oracle authored Management Plug-ins or Deployment Procedures. In general, Oracle-supplied entities are read-only. You can create a copy and customize the copy as per your needs but you cannot modify the original Oracle-supplied entity.

These entities can also be published on Oracle Web sites such as Oracle Technology Network (OTN) and My Oracle Support (MOS). You can download and import the entity archive into their Cloud Control deployment using specific import features provided by the update-able entity.

Entity Types That Can Be Updated

Examples of update-able entity types are:

- Management Agents
- Management Plug-ins
- Management Connectors
- Database Profiles and Gold Images
- Application Server Profiles and Gold Images
- Provisioning Bundles
- Enterprise Manager Deployment Pre-requisite Checks
- Compliance Content
- Diagnostic Checks

15.2 Setting Up Self Update

Before the Self Update feature can be used, a few prerequisites must be met.

- My Oracle Support credentials have been set up. This is required to enable entities to be downloaded from the My Oracle Support site.
- The Software Library (also known as the local store) has been configured. Updates are downloaded to this local store before being deployed into Cloud Control.

Review the following sections for instructions on setting up Self Update:

- Setting Up Enterprise Manager Self Update Mode
- Assigning Self Update Privileges to Users
- Setting Up the Software Library
- Setting Up the EM CLI Utility (Optional)

15.2.1 Setting Up Enterprise Manager Self Update Mode

In order to setup/modify the Enterprise Manager Self Update feature, you must have Enterprise Manager Super Administrator privileges.

1. Log into on to Enterprise Manager as an administrator with Super Administrator privileges.
2. From the **Setup** menu, select **Extensibility**, then select **Self Update**. The Self Update console appears with the default setup displayed.
3. From the **General** status area, click on the **Connection Mode** status to set either offline or online mode. Enterprise Manager takes you to the Patching Setup page to specify online and offline settings.
4. Once the desired connection mode has been selected, return to the Self Update console.

   From here you can select entity types, schedule updates from the Enterprise Manager Update Store.

15.2.2 Assigning Self Update Privileges to Users

In order for Enterprise Administrators to use the Self Update feature, they must have the requisite privileges. The Enterprise Manager Super Administrator must assign the following Self Update roles to these administrators:

- VIEW_SELF_UPDATE - User can view the Self Update console and can monitor the status of download and apply jobs.
- MANAGE_SELF_UPDATE - User can schedule download and apply jobs. User can also suppress/unsuppress updates. This privilege implicitly contains VIEW_SELF_UPDATE.

By default, the Super Administrator will have MANAGE_SELF_UPDATE privilege granted to him.

To assign Self Update privileges to regular Enterprise Manager administrators:

1. From the **Setup** menu, select **Security**, then select **Administrators**.
2. Select an administrator and click **Edit**.
3. From the Roles page, assign the appropriate Self Update roles.

15.2.3 Setting Up the Software Library

The Software Library is a repository that stores software patches, virtual appliance images, reference gold images, application software and their associated directive scripts. It allows maintaining versions, maturity levels, and states of entities.

In the context of applying updates, it is the "local store" that entities are downloaded to before deployment.

Follow these steps to set up the Software Library.

1. Create a folder in the system where Enterprise Manger is installed. For example, /net/*hostname*/scratch/aime/swlib1.
2. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
3. Click **Actions**, then **Administration**.
4. Click **Add**.
5. In the pop up window, enter a name and location for the folder you want to use as the Software Library. For example, `swlib1` and `/net/hostname/sample/swlib1`. This should be the folder that you created in step 1.
6. Wait for the processing to complete.

15.2.4 Setting Up the EM CLI Utility (Optional)

If you plan to apply software updates in offline mode, or to deploy non-Oracle plug-ins, you will need to use the Enterprise Manager Command Line Utility, or EM CLI, to import entity archives for deployment to Enterprise Manager.

A page is provided in the Cloud Control console with instructions on setting up EM CLI. Access the page by appending `/console/emcli/download` to the URL used to access the Cloud Control console:

`https://emcc_host:emcc_port/em/console/emcli/download`

For example:

```
https://server01.example.com:7801/em/console/emcli/download
```

15.3 Applying an Update

The process for applying updates is essentially as follows:

- Check for the latest updates available from Oracle.
- Download the updates you want to apply to the Software Library.
- Apply the update.

Review the following sections to learn how to apply an update:

- Applying an Update in Online Mode
- Applying an Update in Offline Mode

15.3.1 Applying an Update in Online Mode

Updates must be downloaded to the Software Library (the local store) before they can be applied. You can review the latest available updates from the Self Update console.

Note that Enterprise Manager must have access to the Enterprise Manager Store via the Internet to download available updates. If this access is not possible, you can download entities in offline mode. See Section 15.3.2, "Applying an Update in Offline Mode" for details.

1. From the **Setup** menu, choose **Extensibility**, then choose **Self Update**.
2. Click **Self Update** to get the complete list of available updates.
3. Select the desired entity type and choose **Open** from the **Action** menu. The entity type page appears.
4. Select an update from the list of available updates.
5. Click **Download**. The **Schedule Download** dialog appears.
6. Select when to download the update. Note that multiple downloads can be scheduled simultaneously.

   The following options are available:

   - Immediately
   - Later (specified time)
   - Whether or not to send a notification when the download is complete.
7. Click **Select**. An Enterprise Manager job is created to download the update to the Software Library.

   Enterprise Manager starts downloading the archive from the Oracle Enterprise Manager store. Wait for the download to complete. (When in offline mode the system starts reading from the specified location.)

   When the download is complete, Enterprise Manager displays the Confirmation page, and the downloaded plug-in is shown in the local Oracle Enterprise Manager Store.

> **Note:** The page is not refreshed automatically. Click the refresh icon to view the updated download status.

8. Once an entity has been downloaded to the Software Library, it is ready to be applied to your installation. Select an update from the list whose status is **Downloaded**, then click **Apply**.

   Note that the application process varies according to the entity type:

   - For connectors, diagnostic checks, and compliance content, clicking **Apply** will install the update to Enterprise Manager. No further action is required.
   - For plug-ins, you will be redirected to the plug-in deployment page.
   - For provisioning bundles, you will need to exit the Enterprise Manager console, run Opatch and other commands via a terminal, and then restart the OMS.

15.3.2 Applying an Update in Offline Mode

Under certain circumstances, such as in high security environments, an active Internet connection between Enterprise Manager and the Enterprise Manager Update Store may not be available. In such situations, the Self Update feature can be used in "offline mode".

The update process still requires that a computer exist at your site that has an Internet access, as a connection to the Enterprise Manager Update Store is still required to obtain the updates. Update files from this computer can then be transferred to a computer behind your firewall.

The generic offline mode update procedure is as follows:

1. Ensure Cloud Control is set to offline mode. From the **Setup** menu, select **Provisioning and Patching**, then select **Offline Patching**.
2. Change the setting for Connection to **Offline**.
3. Click **Check for Updates** on the Self Update home page. A message is displayed that contains the URL to be accessed to download a catalog of all updates.
4. From an Internet-enabled computer, download the catalog file using the aforementioned URL.
5. Copy the downloaded file to the Oracle Management Service host or the Management Agent host you will deploy the update to.
6. Run the `emcli import_update` command to import the file into the Oracle Management Service instance or the Management Agent you want to update.

   See Section 15.5.1.2, "Importing an External Archive into Enterprise Manager" for instructions on using the `emcli import_update` command.
7. Review the update from Self Update Home and select and click download. A messages is displayed with URL and instructions.
8. Select and click **Apply** to apply the update.

15.4 Acquiring or Updating Management Agent Software

Management Agent software for the various platforms (operating systems) supported by Enterprise Manager Cloud Control can be downloaded to the Software Library using the Self Update console. Once a Management Agent is persisted to the Software Library, it can be installed on host machines that you want to bring under Cloud Control management using the Management Agent installation wizard.

Steps for obtaining Management Agent software in both online and offliine modes are discussed below.

- Acquiring Management Agent Software in Online Mode
- Acquiring Management Agent Software in Offline Mode

15.4.1 Acquiring Management Agent Software in Online Mode

Using Self Update in online mode requires Enterprise Manager to have access to My Oracle Support via the Internet.

1. From the **Setup** menu, choose **Extensibility**, then choose **Self Update**.
2. Select the entity type *Agent Software* and choose **Open** from the **Action** menu. The entity type page appears to show agent software for different platforms.
3. Select an update from the list of available updates. All entries other than the one which matches the platform of the OMS host should show their status as *Available*.
4. Click **Download**. The Schedule Download dialog opens.
5. Select when to download the update. The following options are available:
   - Immediately
   - Later (specified time)
   - Whether or not to send a notification when the download is complete
6. Click **Select**. An Enterprise Manager job is created to download the Agent software to the Software Library.

   Enterprise Manager starts downloading the archive from the Oracle Enterprise Manager store. Wait for the download to complete. (When in offline mode the system starts reading from the specified location.)

   When the download is complete, Enterprise Manager displays the Confirmation page, and the downloaded plug-in is shown in the local Oracle Enterprise Manager Store.
7. Once the download is complete, select the Management Agent, then click **Apply**. This step will stage the Agent software in the Software Library and make it available to the Add Targets wizard, which you will use to install the Agent on host machines.
8. Click **Agent Software** to launch the Add Targets/Agent Installation wizard.

15.4.2 Acquiring Management Agent Software in Offline Mode

Follow this Self Update process only when Enterprise Manager is in offline mode.

1. Ensure Cloud Control is set to offline mode. From the **Setup** menu, select **Provisioning and Patching**, then select **Offline Patching**.
2. Change the setting for Connection to **Offline**.
3. From the **Setup** menu, select **Extensibility**, the select **Self Update**.
4. Click **Check for updates** on Self Update home page. A message is displayed that contains the URL to be accessed to download a catalog of all updates.

   Note that the archive containing the Management Agent software should also be available from the Oracle Technology Network (OTN) site.

5. From an Internet-enabled computer, download the catalog file using the aforementioned URL.
6. Copy the downloaded file to either of the following:
   - To any host that has a Management Agent and EM CLI installed
   - To the Oracle Management Service (OMS) host
7. Run the `emcli import_update` command to import the archive into the Oracle Management Service instance or the Management Agent you want to update.

   See Section 15.5.1.2, "Importing an External Archive into Enterprise Manager" for instructions on using the `emcli import_update` command.
8. Select the entity type *Agent Software* and choose **Open** from the **Action** menu. The entity type page appears displaying agent software for different platforms.
9. Select an update from the list of available updates. All entries other than the one which matches the platform of the Oracle Management Service host will show their status as Available.
10. Click **Download**. A message is displayed with a URL and instructions.
11. From an Internet-enabled computer, download the file from the URL displayed in step 8. Do one of the following:
    - Copy the file to a Management Agent and follow the instructions displayed in step 8.
    - Copy the file to Oracle Management Service and follow the instructions displayed in step 8.

    At this stage, the update will show up in downloaded state in the Self Update home page.
12. Once the download is complete, select the Managment Agent, then click **Apply**. This step will stage the Management Agent software in the Software Library and make it available to the Add Targets wizard, which you will use to install the Management Agent on host machines.
13. Click **Agent Software** to launch the Add Targets/Agent installation wizard.

15.5 Deploying and Updating Plug-ins

A plug-in is a component (module) that can be plugged into an existing Enterprise Manager Cloud Control installation to extend its management and monitoring capabilities. The plug-in management features provided with Cloud Control simplify lifecycle management by allowing you to install and deploy plug-ins, as well as upgrade to newer versions as they become available.

With all releases of Enterprise Manager, external companies provide plug-ins that monitor specific types of targets, such as non-Oracle databases or applications. In a major architectural change, core Enterprise Manager Cloud Control features for managing and monitoring Oracle technologies - such as Oracle Database, Oracle Fusion Middleware and Oracle Fusion Applications - are also provided via plug-ins that can be downloaded and deployed.

This new “pluggable” framework enables Cloud Control to be updated with management support for the latest Oracle product releases, without having to wait for the next Cloud Control release to provide such functionality. For example, when a new version of Oracle Database is released, you can simply download and deploy the latest

Oracle Database plug-in, which will include management support for the latest release.

See the following sections:

- Importing an External Archive into Enterprise Manager
- Deploying a Plug-in
- Updating a Plug-in

15.5.1 Deploying a Plug-in

The process of enabling a plug-in to monitor targets in Enterprise Manager Cloud Control is essentially as follows:

- The plug-in archive is downloaded to the Software Library.
- The plug-in is deployed to the Oracle Management Service (OMS) instance that manages and monitors targets of the plug-in's type. As part of this process, the plug-in metadata is written to the Management Repository.

  Note that deployment of most plug-ins will require the Oracle Management Service instance to stop, then re-start. This process will occur automatically as part of the plug-in deployment process.
- Targets that the plug-in will monitor are added - or promoted to managed status - in Cloud Control.
- As part of the target promotion process, a Management Agent containing the required plug-in content is assigned to monitor the target. (If a Management Agent already exists on the target host, it will be updated with the plug-in content.)

See the following sections for details:

- Downloading a Plug-in from the Enterprise Manager Store
- Importing an External Archive into Enterprise Manager
- Deploying a Plug-in to Oracle Management Service (OMS)
- Adding Targets for the Plug-in to Monitor
- Important Details Regarding Plug-in Deployment

15.5.1.1 Downloading a Plug-in from the Enterprise Manager Store

Plug-in archives that are provided by Oracle will be made available through the Enterprise Manager Store, just like any other update. The plug-in must be downloaded to the Software Library, again like any other update, before it can be deployed.

See Section 15.3.1, "Applying an Update in Online Mode" for instructions on download plug-in archives.

15.5.1.2 Importing an External Archive into Enterprise Manager

External, non-Oracle plug-ins - that is, plug-ins that have been created by a company other than Oracle - must be imported into the Software Library using EM CLI before the plug-in deployment process can be initiated.

Note that this process can also be used to import other types of entity archives if Self Update is used in offline mode.

The plug-in archive must first be downloaded to an accessible location. Once downloaded, the plug-in can be imported into Enterprise Manager Cloud Control

using the `emcli import_upate` command. See Section 15.2.4, "Setting Up the EM CLI Utility (Optional)" for instructions on setting up EM CLI.

You have two options for importing the plug-in archive, depending on where EM CLI is installed:

- If EM CLI is on the same system as the system as the location you downloaded the plug-in archive (*.opar file) to, run the following command.

  ```
  emcli import_update
   -file="<path to *.opar file>"
   -omslocal
  ```

  The `-omslocal` flag indicates that the plug-in archive is on the same system where you are running this command and the path exists on this system.

- If EM CLI is on a different system than the plug-in archive, run the following command:

  ```
  emcli import_update
          -file="<path to *.opar file you created>"
          -host="host1.example.com"
          -credential_name="host1_creds"
          -credential_owner="admin1"
  ```

  The command syntax is as follows:

  - `-file`: The absolute path to the *.opar file on the system where you created the archive.
  - `-host`: The target name for a host target where the file is available.
  - `-credential_name`: The name of the credentials on the remote system you are connecting to.
  - `-credential_owner`: The owner of the credentials on the host system you are connecting to.

- As an alternative to the previous step, you can also run the following command:

  ```
  emcli import_update
        -file="<path to *.opar file you created>"
        -host="hostname"
        -credential_set_name="setname"
  ```

  `-credential_set_name`: The set name of the preferred credential stored in the Management Repository for the host target. It can be one of the following:

  - `HostCredsNormal`: The default unprivileged credential set.
  - `HostCredsPriv`: The privileged credential set.

Once the archive has been imported, the plug-in (or other entity downloaded in offline mode) can be deployed.

15.5.1.3 Deploying a Plug-in to Oracle Management Service (OMS)

A plug-in must be deployed on Oracle Management Service (OMS) before it can be used to monitor targets. Follow the steps below to deploy the plug-in on Enterprise Manager Cloud Control.

> **Note:** A plug-in deployment failure could put the Management Repository in an inconsistent state. Therefore it is strongly recommended that you back up the Management Repository before deploying the plug-in.

1. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**. Enterprise Manager displays the list of plug-ins that have been downloaded and can be deployed on the Plug-ins page.
2. On the Plug-ins page, select the specific plug-in you want to deploy. Note that the plug-in archive must have already been downloaded to the Software Library.
3. Click **Deploy On>Management Servers**.

   Be sure that dependent plug-ins are deployed and that all existing Management Agents are compatible with the version of the specified plug-in. Enterprise Manager prompts for credentials if the agent is not available.

   > **Note:** Plug-ins must be deployed on Oracle Management Service prior to being deployed on Management Agents.

4. Specify the required details on the Deploy Plug-in dialog box. Note that you will need the Management Repository SYS user password to complete the deployment process.

   In the Version of Plug-in to Deploy section, select or choose the **Plug-in** version from the Plug-in drop-down. The **Target Type** information is displayed in the table. Enter the **Repository sys Password**, then click **Continue**.
5. Proceed through the steps in the Deploy Plug-in dialog box.
6. Click **Deploy** to deploy the selected plug-in on all Enterprise Manager servers.

   Deployment time varies by plug-in, depending on the volume of data populated in the Management Repository. A page is displayed that allows you to monitor the deployment status.

   Note that deployment of most plug-ins will require the Oracle Management Service instance to stop, then re-start. This process will occur automatically as part of the plug-in deployment process.

   You can also monitor the deployment status by going to the Enterprise Manager Cloud Control console, then going to the plug-ins page as described in step 1, selecting the plug-in and select the **Recent Deployment Activities** tab at the bottom of the page for the selected plug-ins. This bottom section also lets you see details of your plug-in, which includes the plug-in ID, version, vendor, and so on.

   If any of the steps during plug-in deployment fails, the log is available in $ORACLE_HOME/cfgtoollogs/pluginca/*. Append these in while logging a support request for failure while deploying the plug-in. You can also use them to debug the problem.

15.5.1.4 Adding Targets for the Plug-in to Monitor

In the current Cloud Control release, deployment of a plug-in to a Management Agent that will monitor targets is no longer required. Instead, the plug-in for a specific target type is automatically deployed with the Management Agent that will monitor targets of that type.

This is a significant change from previous releases, in which plug-ins had to first be manually deployed to a Management Agent, and a target instance then had to be manually added to the Management Agent.

You can add targets that the plug-in will monitor through Enterprise Manager Cloud Control by selecting **Add Targets** from the **Setup** menu. The process for adding targets - known in Cloud Control terminology as *target promotion* - will vary depending on the option you choose.

See Chapter 1, "Discovering Targets" for details on adding targets for the plug-in to monitor.

15.5.1.5 Important Details Regarding Plug-in Deployment

- You can import multiple versions of the same plug-in. The version to deploy can be selected from a list if you are using Cloud Control to deploy the plug-in, or can be specified on the command line if using EM CLI.
- Only one version can be deployed on the Oracle Management Service (OMS) at any given time. If a later version has been deployed previously, it cannot be downgraded to an earlier version.
- Updating a plug-in to a new version does not remove the content of the older plug-in.
- The Management Agent can have the same or earlier version of the plug-in that is deployed on the OMS host. However a version later than the version on the OMS host is not allowed on the Management Agent host.
- If the plug-in version available on the OMS is not certified on the Management Agent platform (operating system), the version of the plug-in version that *is* supported will be deployed to the Management Agent.
- The plug-in on the OMS host and the plug-in on the Management Agent host can be updated independently of each other.
- Available updates are visible on the Plug-ins page. They can be downloaded from the Enterprise Manager store or imported using EM CLI as described in Section 15.5.1.2, "Importing an External Archive into Enterprise Manager".

15.5.2 Updating a Plug-in

When a new version of a plug-ins is released, you have the option of updating the Oracle Management Service instances and Management Agents that use the plug-in with the new version.

Updated versions of Oracle plug-ins are made available through the Oracle Enterprise Manager Store, where then can be downloaded for deployment. The Oracle Enterprise Manager Store is a central store maintained by Oracle, where plug-in metadata and archives are published. Plug-ins must be downloaded to a location (the Local Store) before being deployed.

Updated versions of external non-Oracle plug-ins are not available through the Oracle Enterprise Manager Store, and must first be imported into Oracle Management Service before being deployed. See Section 15.5.1.2, "Importing an External Archive into Enterprise Manager" for instructions.

Note that before a plug-in can be upgraded on a Management Agent, it must first be upgraded on the associated Oracle Management Service instance. In addition, you do not need to remove the existing version of a plug-in from OMS or Management Agents before upgrading to the latest version.

For more information, see the following sections:

- Downloading the Latest Plug-in Archive from the Oracle Enterprise Manager Store
- Updating a Plug-in on Oracle Management Service
- Updating a Plug-in on a Management Agent
- Un-deploying a Plug-in

15.5.2.1 Downloading the Latest Plug-in Archive from the Oracle Enterprise Manager Store

Plug-ins must be downloaded to the Software Library (the local store) before being deployed.

To download the latest plug-in archive from the Oracle Enterprise Manager Store to the Software Library, follow the instructions in Section 15.5.1.1, "Downloading a Plug-in from the Enterprise Manager Store".

15.5.2.2 Updating a Plug-in on Oracle Management Service

Once the plug-in is available in the Local Store, you can deploy it to Oracle Management Service. See Section 15.5.1.3, "Deploying a Plug-in to Oracle Management Service (OMS)" for instructions.

The plug-in must be updated on the OMS instance managing relevant targets before it is updated on active Management Agents.

Note that you do not need to remove the existing version of a plug-in from the OMS instance before upgrading to the latest version.

15.5.2.3 Updating a Plug-in on a Management Agent

Once the updated version of the plug-ins has been deployed on OMS, all Management Agents that use the plug-in can be updated with the latest version.

Note that you do not need to remove the existing version of a plug-in from the Management Agents before upgrading to the latest version.

1. From the **Setup** menu, choose **Extensibility**, then select **Plug-ins**.
2. Select the row for the plug-in you want to update to in the table.
3. Click **Deploy On>Management Agent**.

> **Note:** Plug-ins must be deployed on a Managed Server prior to deploying on Management Agents.

4. On the Deploy Plug-in dialog box, use the Management Agents to Deploy section to **Add** or **Remove** the agents to which you want to deploy the plug-in.

   When you click **Add**, Enterprise Manager displays the Search and Select dialog box where you can select the agents to add. Click **OK** to return to the Deploy Plug-in dialog box. Only agents running the operating systems supported by the selected plug-in may be selected.

   To remove an agent, highlight it in the Management Agents to Deploy table and click **Remove**.

5. Specify other required details on the Deploy Plug-in on Agents dialog box.

Select the **Plug-in** version from the Plug-in drop-down. Click **Next**.

6. Click **Deploy** to deploy the chosen plug-in on the selected Management Agents.
7. Enterprise Manager displays a page that monitors the deployment status and begins the deployment process with the Install option. Deployment occurs in parallel on all selected agents.

15.5.2.4 Un-deploying a Plug-in

To un-deploy a plug-in from Oracle Management Service or a Management Agent, follow the steps below. Un-deploying a plug-in completely removes monitoring of affected targets, and also removes all plug-in metadata from the Management Repository.

> **Notes:** A plug-in must be deployed from all Management Agents before being un-deployed from an Oracle Management Service instance. If not, un-deployment will fail.
>
> In addition, default plug-ins provided by Oracle, such as the Database Management plug-in, cannot be un-deployed.

1. From the **Setup** menu, choose **Extensibility**, then select **Plug-ins**.
2. Select the row for the plug-in you want to remove to in the table.
3. Click **Undeploy From**, then either **Management Servers** or **Management Agen**t. You can then select the OMS or Management Agent you want to remove the plug-in from.
4. Confirm the plug-in removal. Enterprise Manager notifies the connected and relevant Enterprise Manager users and begins the de-configuration process.

16

Patching Enterprise Manager

This chapter provides an overview of patching, and describes how you can patch the Enterprise Manager core components, mainly Oracle Management Service (OMS), Oracle Management Agent (Management Agent), and Oracle Management Repository (Management Repository).

- Overview
- Patching OMS and Management Repository
- Patching Enterprise Manager Agents

16.1 Overview

A patch is an entity that contains one more bugs fixes. In order to transform a software product with a defect to a software product without a defect, you must apply patches, this process is called Patching. The patching cycle involves downloading patches, applying patches, and verifying the applied patch to ensure that the bug fixes present in the patch reflect appropriately.

Patching involves migrating from one version of the software product to another, within a particular release, unlike upgrading which involves moving from one release of a product to another newer release of the software product.

Oracle periodically releases the following types of patches to fix the bugs encountered in the core Enterprise Manager components:

- **Interim Patches** are released to fix a bug, or a collection of bugs.
- **Interim Patches (for Security bug fixes)** are released to provide customer specific security fixes.
- **Diagnostic Patches** mainly help diagnose and verify a fix, or a collection of bug fixes.
- **Bundle Patch Updates** are cumulative collection of fixes for a specific product or component.
- **Patch Set Updates** (PSU), are cumulative patch bundles that contain well-tested and proven bug fixes for critical issues. PSUs have limited new content, and do not include any changes that require re-certification.
- **Security Patch Updates** are cumulative collection of security bug fixes.

16.2 Patching OMS and Management Repository

Patching OMS and Management Repository follows the Manual patching approach that requires you to follow step-by-step instructions to patch a target. This mechanism of patching expects you to meet certain prerequisites, manually validate the patch for applicability and conflicts, and patch only one target at a time.

This section contains the following topics:

- OMS Patches
- Repository Patches
- Applying OMS and Repository Patches

16.2.1 OMS Patches

OMS patches typically fix one or more OMS errors encountered. These patches can be downloaded from *My Oracle Support* portal. For information about the critical OMS patches released by Oracle that apply to your environment, see the Patch Recommendation region available on the **Patches and Updates** tab in *My Oracle Support*. Alternately, if you know the patch number of the OMS patch that you need to apply, go to the Patch Search region available on the **Patches and Updates** tab, enter the patch number, and click **Search**. The OMS patches are applied using the `OPatch` utility, or the `emctl` utility. For information on applying the patch see "Applying OMS and Repository Patches"

In Enterprise Manager 12*c*, OMS patches are available as OMS Core Patches and OMS Plugin Patches, required to patch the core component and plugins respectively. Ensure that you navigate to the correct directory location under `<middleware_home>` to patch OMS core or OMS plugin:

```
<middleware_home>
    |_____oms
    |_____plugins
          |_____oracle.sysman.db.oms.plugin_12.1.0.1.0
          |_____oracle.sysman.emas.oms.plugin_12.1.0.1.0
          |_____oracle.sysman.mos.oms.plugin_12.1.0.1.0
```

For example, to patch OMS core, you must navigate to `<middleware_home>/oms` and for plugins, navigate to `<middleware_home>/plugins`.

16.2.2 Repository Patches

Enterprise Manager Repository patches typically update PL/SQL procedures, or other SQL content. They are patched using the `emctl` patching tool.

> **Note:** The Repository Patches are applied on the OMS targets. For location details of a core, and plugin patch, see "OMS Patches"

For information on applying the patch see "Applying OMS and Repository Patches"

16.2.3 Applying OMS and Repository Patches

To apply OMS or repository patches, follow these steps:

1. Log into *My Oracle Support* (`https://support.oracle.com`) console with the necessary credentials.

> **Note:** Check the Patch Recommendation region to view the patches recommended for your environment. You can also provide the recommended patch number in the patch search region to download the recommended patch.

2. On the *My Oracle Support* home page, click **Patches and Updates.**
3. Enter the patch number in the Patch Search region, then click **Search.**
4. Select the patch, and from the context menu, click **Download**.
5. After downloading the zip file, follow the instructions available in the `Readme.html` or `Readme.txt` to patch the target.

> **Note:** Depending on whether it is a core patch or a plugin patch you must log into the host machine, navigate to the directory location, and unzip the patch file.

16.3 Patching Enterprise Manager Agents

Oracle offers two approaches to apply the Agent patches: automated approach, and the manual approach. Oracle strongly recommends you to use the automated approach because it not only saves time and effort in mass-deploying patches but also reduces human intervention, thereby minimizing the errors involved while patching.

This section contains the following topics:

- Management Agent Patches
- Automated Agent Patching
- Manual Agent Patching

16.3.1 Management Agent Patches

Management Agent Patches are released to fix one or more errors encountered in the agent targets. In addition to the Management Agent running on OMS, you can patch all the agent targets that report to the OMS running on your host machine.

A GUI based utility called Patches and Updates is used to patch the Agent targets. For information about accessing the tool, see "Accessing Patches and Updates".

In Enterprise Manager 12*c*, there are separate Agent patches for core components and for plugins. Ensure that you navigate to the correct directory location under `<installation_base_directory>` to patch Agent core or Agent plugin:

```
<installation_base_directory>
    |_____core
          |_____12.1.0.1.0
    |_____plugins
    |_____plugins.txt
    |_____plugins.txt.status
    |_____agent_inst
    |_____sbin
    |_____agentimage.properties
```

For example, to patch Agent core component, you must navigate to `<installation_base_directory>/core/` and for plugins, navigate to `<installation_base_directory>/plugins`.

16.3.1.1 Patches and Updates Versus My Oracle Support

The Table 16–1 captures the advantages of using Enterprise Manager Cloud Console over *My Oracle Support* for Automated Patching:

Table 16–1 Advantages of using Enterprise Manager Cloud Console to My Oracle Support

Enterprise Manager Cloud Control	My Oracle Support
Enterprise Manager Cloud Control enables you to leverage existing Cloud Control Agents to perform configuration data collection without having to install individual configuration managers on each managed target.	You must install Oracle Configuration Manager 10.3.2 (or higher) or Enterprise Manager Cloud Control for My Oracle Support configuration collection (Enterprise Manager harvester) on each of the managed target for the latest patch recommendations.
Enterprise Manager Cloud console is inherently built with the intelligence to support the entire lifecycle of patching, it allows you to select the recommended patch for your environment, add to a plan, and deploy the patch on the selected targets after validating the patches for applicability and conflicts in your environment.	MOS only supports adding the recommended patches to a plan, and validating the patch for conflicts. You cannot deploy the patch from MOS.
In Enterprise Manager, once the patches are applied, they will not appear in the recommended patches list for the selected target.	In MOS this is not an automated process, depending on the data collection by the configuration manager present on the patched targets, the patch recommendation region is automated.

16.3.2 Automated Agent Patching

Automated Patching is a quick-and-easy, reliable, and a GUI-based patching mechanism that is facilitated using Patch Plans, a new concept introduced through the Patches and Updates functionality within the Enterprise Manager Cloud Control console (Cloud Control console).

Automated patching can be performed in the Online mode, and in the Offline mode. In the Online Mode, you can connect to *My Oracle Support* Web site to download the patches. However, if you are patching in the offline mode, then you must ensure that the patches to be applied are already available on the Software Library.

This section includes the following:

- Accessing Patches and Updates
- Viewing Patch Recommendations
- Searching Patches
- Applying Management Agent Patches
- Verifying the Applied Agent Patches
- Validating Agent Patch Errors
- Deinstalling the Applied Agent Patches

16.3.2.1 Accessing Patches and Updates

To access the Patches and Updates page, in Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching,** then select **Patches and Updates**.

The following page appears:

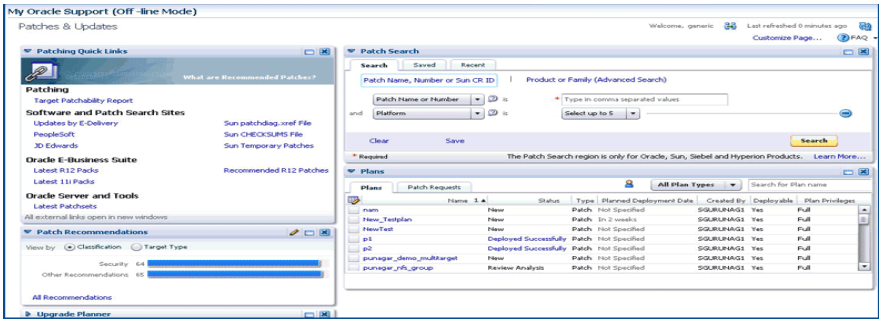

Some of the advantages of using the Patches and Updates page (Automated Patching approach) are:

- Patching operations are more organized, done through a single window, and is always initiated only from the OMS.
- Allow you to schedule jobs that will run periodically, and connect to *My Oracle Support*, check for the latest patches, and automatically download them. This relieves you of the manual effort involved in searching the latest patches and patch sets, and downloading them whenever they are available.
- One patch plan can be used to add multiple patches to multiple sets of homogeneous targets. For example, both core and plugin agent patches can be added to the same plan.

16.3.2.2 Viewing Patch Recommendations

The Patch Recommendation region is one of the regions in the patching domain that proactively communicates all the recommendations that are applicable to your environment. This region minimizes human effort in terms of searching for Oracle recommended patches, which may or may not apply to your environment.

Note: Keep the following points in mind:

- Recommendations are not available for custom plugins. Oracle only supports the default plugins that ship with the product, and releases timely updates for them.
- You must use the configuration manager release 10.3.2 or higher for Patch Recommendations to be enabled.

Figure 16–1 captures the Patch Recommendations region as it appears in the Patches and Updates page.

Figure 16–1 Patch Recommendations

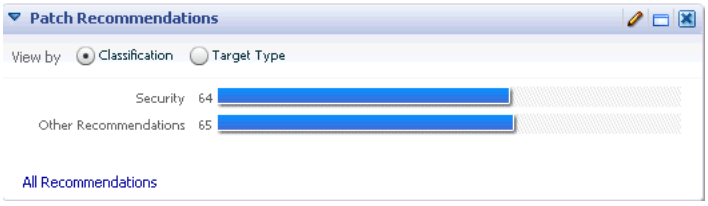

Patches are primarily classified as Security patches, and Other Recommended patches. For example, if you select **Security** from the graph, all the security related patches are displayed on the Patch Recommendation page. Alternately, you can also view the patches by their target types. Click the bar graph to drill down to a list of recommended patches, view details about those patches, download the patches, or add them to a patch plan. The bar graph summarizes the number of issues found (for example, if there is one issue, then there is one recommendation for one target).

Patch Recommendation region allows you to:

- Compare the patches installed in your configuration with what Oracle recommends, and identifies any missing patches.
- Identify and prioritize targets missing Critical Patch Updates.
- Identify missing recommended patches issued by Oracle.

16.3.2.3 Searching Patches

The Patch Search is a region in the patching domain that allows you to search for Oracle, Sun, Siebel, and Hyperion products. The primary purpose of searching a patch is to limit results to the exact patch name or number.

From Figure 16–2 you can see that the Patch Search region contains three tabs: **Search**, **Saved**, and **Recent**. The Search tab lets you apply the desired filters to drill down to the exact results. You can choose to save the search, and view it later. All the saved searches appear in the Saved region. A history of all the searches is registered in the Recent tab, and you can check the logs if you want to access them.

Figure 16–2 Patch Search

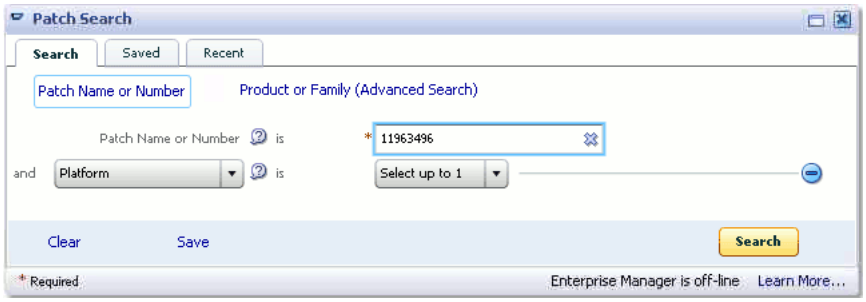

Use any of the following approaches to search for a patch:

- Basic Search
- Advanced Search

Basic Search

To perform a **Basic Search**, follow these steps:

1. From the **Enterprise** menu, select **Patching and Provisioning**, then select **Patches and Updates**.
2. On Patches and Updates page, enter the **Patch Name or Number, or Sun CR ID,** if there are multiple values, then they must be separated by commas. You can select the platform name from the list in the Patch Search Region, and then click **Search**.
3. The Patch Search Results page displays the results based on the search criterion provided.

   From the Patch Search results page, you can do the following:

   – You can highlight one or more patches in the search results and, from the inline tool bar, add the patches to a patch plan (if you use the collector), download the patches, or copy patch details to the system clipboard.

   – If a single patch is highlighted, you can view the patch readme.

Advanced Search

To perform an **Advanced Search**, follow these steps:

1. From the **Enterprise** menu, select **Patching and Provisioning**, then select **Patches and Updates**.
2. On the Patches and Updates page, click **Product or Family (Advanced Search)**.
3. Enter a **Product** name, and set the **Release** number to narrow down your search.

   Additionally, if you are accessing Patches and Updates page in the Offline mode, then click **(+)** icon to add more filters like **Type**, **Platform, or Language**. In the Online mode, in addition to all the filters available in the Offline mode, you can use advanced search categories like: **Classification, Description, Patch Target, or Updated** to drill down to the desired results.

   Note: Advanced Search allows you to search for recommended patches for your environment through the **Classification** search category available when you are accessing Patches and Updates page in the online mode.

4. Click **Search**.

   After updating the appropriate search filters, you can save the search combination by clicking **Save**.

5. The Patch Search Results page displays the results based on the search criterion provided.

   From the Patch Search results page, you can do the following:

   – You can highlight one or more patches in the search results and, from the inline tool bar, add the patches to a patch plan (if you use the collector), download the patches, or copy patch details to the system clipboard.

   – If a single patch is highlighted, you can view the patch readme.

16.3.2.4 Applying Management Agent Patches

To apply a Management Agent patch using patch plans, follow these steps:

Note: Applying the core agent patch, or the plugin agent patch follows the same patching process (as listed in this section.)

Ensure that the patches selected are applied on homogeneous set of targets, which means that the patches selected should have the same platform and version, as the targets being patched. For example, 12.1.0.1.0 Linux x86 patches can only be applied on 12.1.0.1.0 LinuxX86 targets, any mismatch will result in a patching error.

---

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Patches and Updates**.
2. On the Patches and Updates page, select Management Agent patches from one of the following regions:
   - In the Patch Recommendation region, click the graph to drill down to the list of recommended patches for your environment.

     For more information on Patch Recommendation, see "Viewing Patch Recommendations"
   - In the Patch Search region, enter the patch number of the Agent patch or perform an Advanced Search to search, and select the desired patch.

     For more information on Basic and Advanced search, see "Searching Patches"
3. Select one or more patches, and from the context menu, click **Add to a Plan**.

   Click **Add to New** to create a new plan, if not you can update an existing plan by selecting **Add to Existing** option.
4. If you select **Add to New**, then one of the following dialog box appears:
   - In the Create a New Plan dialog box, enter a unique name for your plan, and click **Create Plan**.

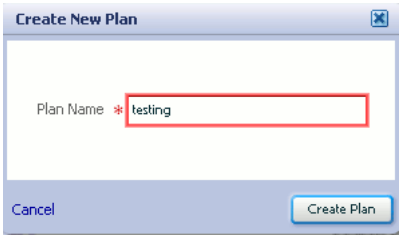

     The patch selected along with the associated target gets added to the plan.
   - In the Add Patch to Plan dialog box as shown in Figure 16–3, enter a unique name for the plan, and select the targets. To search and select targets, follow one of these approaches:
     - If you know the target name, then enter the name of the target in the search field.
     - Click the search icon to view all the targets reporting to the OMS running on your host machine, group them by type **Agent**, and select the desired Agent targets.
     - Create a group of Agent targets, and provide the group name in the search field to add all the targets in that particular group to the plan. To create a group of targets, from **Setup** menu, select **Add Target**, and then click

group. On the Create Group page, add all the Agent targets to the group, and create the group with a unique name.

After selecting the targets, click **Create Plan**. The selected patches and associated targets are added to the plan after validating for conflicts.

Figure 16–3 Add Patch To Plan

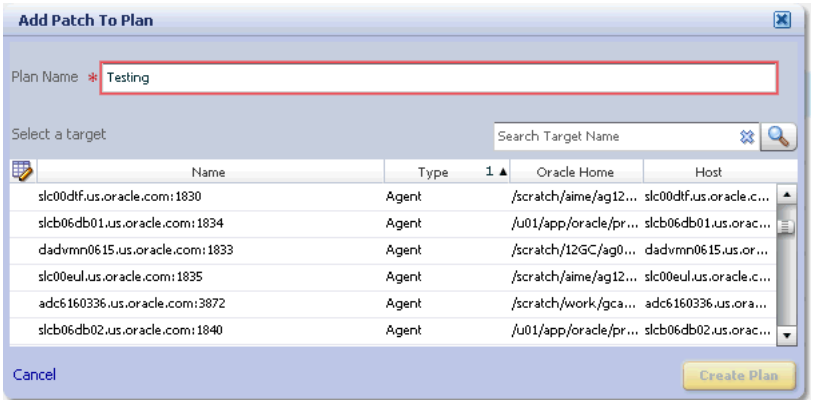

Note: If you want to add to an existing plan, click **Add to Existing** and, then click **More.** From the Add Selected Patch to Plan dialog box, select the plan name and click **Select Targets**. From the Add Selected Patch to Plan <plan_name> dialog box, select the targets and click **Add to Plan**.

5. If the selected patches are applied on homogeneous targets, then the plan gets successfully created with a link to **View Plan**. Click the link to view the plan details.

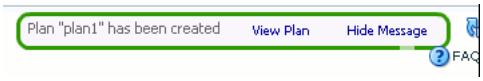

If any of the agent targets added to the patch plan are NFS-Agents, then you may see a warning message **Issues with Adding Patch** as shown in Figure 16–4. As a solution to this problem, a list of all the targets impacted appear, click **Add All To Plan** to add all the affected targets to the patch plan.

Figure 16–4 Issues with Adding Patch: Impacts Other Targets

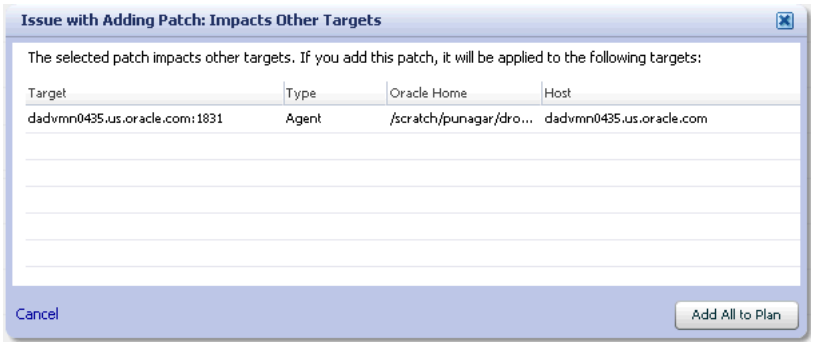

However, if there is a mismatch between the platform and version of the patch selected, and that of the target, you may see one of the following warnings:

Note: Oracle recommends that you fix the warning before proceeding as it may result in an error during the plan validation. However, if you still want to proceed, you can select the patches, and click **Ignore Warnings and Add** to proceed

---

a. **Target is Down**: This warning message appears when the target added is not up and running. All the other homogeneous targets get added to the plan, and the plan is created without this target. Click **View Plan** to see the details.

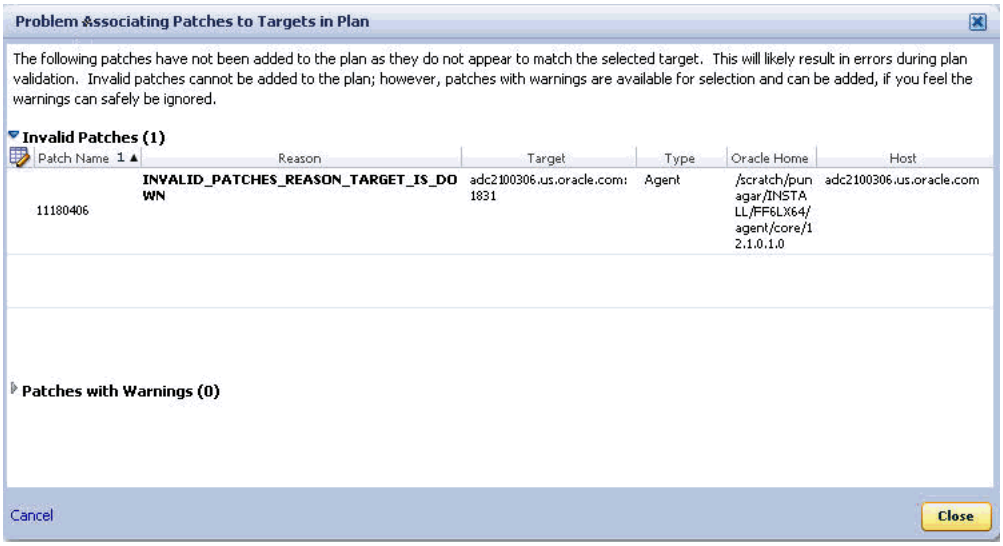

b. **Null Platform**: This error occurs when the target selected appears with a null platform, the validation fails as there is a mismatch in the platform of the patch and that of the target. This could happen when a target is down, in this case the plan is not created until the error is fixed.

6. On the Patches & Updates page, in the Plans region, click the plan name you want to view, and from the context menu, click **View**.

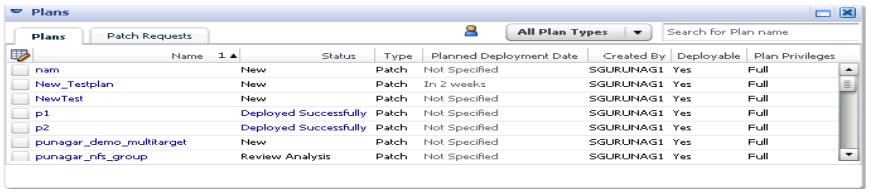

To filter the plans table, select **All Plan Types** or **Patch** depending on your preference. To search for a plan, enter a plan name or partial plan name in the search box, then click the search button.

7. In the Create Plan Wizard, on the Plan Information page, in the Overview section, validate the Patch Plan name. You can choose to edit it if you want.

   (Optional) Select a date and time when you want to deploy the Patch Plan, and enter a short description to describe the Patch Plan.

8. Click **Next**.

9. On the Patches page, review the patches added to the Patch Plan.

   To associate additional targets to a patch that is already in your Patch Plan, click **Add Patch**. In the Edit Search dialog box, enter the patch number and click **Search**. Select the patch, and click **Add to This Plan**. From Add Patch To Plan dialog box, select the targets, and click **Add to This Plan.**

10. Click **Next**.

11. In the Deployment Options page, retain the default location (`%emd_emstagedir%`) available on the target machine or edit the **Stage Location** to provide a new location for staging the Agent patches.

    In the Credentials section, select Oracle Home Preferred Credentials if you have already set them earlier. You can otherwise click **Override Oracle Home Preferred credentials** and set the Normal Oracle Home Credentials and Privileged Oracle Home Credentials to access the Oracle home of the target.

12. Click **Next**.

13. On the Validation page, click **Analyze** to validate the patch before deploying it. A Validation job is submitted, that performs an exhaustive list of checks like: check for conflicts, check for the latest OPatch version, check if the version and platform of the targets and the patch match (homogeneity rule), and so on in the background.

    To track the progress of the job, from **Enterprise** menu, select **Job**, and then click **Activity**. On the Job Activity Page, in the Advanced Search region, enter the name of the job, and then click **Go**. Select the job, and drill down to the steps by click **Expand All**. If the status is **Succeeded**, then the job is valid. If not, then review the issues, and try to resolve it according to the corresponding problem description available on the page. After resolving the issue, click **Re-Analyze**.

    Upon validation, if there are conflicts between the two patches, then you might be recommended to request for replacement patches. In this case, click **Request Patch**.

    If there is a Merge Patch already available, you can directly opt to replace the conflicting patches with the Merge Patch. In this case, click **Replace Patch**.

    See Also: For more information about the common errors during the Validation phase, see "Validating Agent Patch Errors"

14. Click **Next**.
15. On the Review & Deploy page, review the details you have provided for the patch plan, then click **Deploy**.

    A job is submitted, to track the progress of the job, from **Enterprise** menu, select **Job**, and then click **Activity**. On the Job Activity Page, in the Advanced Search region, enter the name of the job, and then click **Go**.

> **Note:** For a demonstration about how to apply patches on Enterprise Manager 12*c* Agents using Cloud Control Console, see *My Oracle Support* note 1359221.1.

16.3.2.5 Verifying the Applied Agent Patches

To verify the applied patches using Enterprise Manager, perform the following steps:

1. In Cloud Control, click **Targets**, then select **All Targets**
2. On the All Targets page, enter the **target name** in the Search Target Name field to search for the target you patched.

   For example, enter **adc2101818** in the search field, and click the search icon.
3. Click the target name to select Oracle home of the target that was patched.

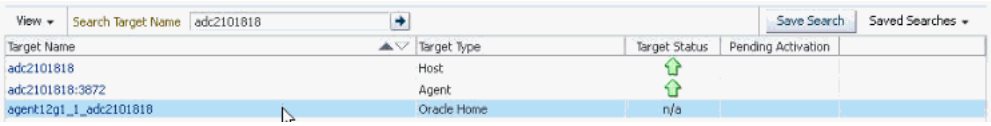

   A summary page with details about the target is displayed.
4. On the Summary page, in the Patch Advisories region, select **Patches Applied** tab to verify all the patches that have been successfully applied on the target.

16.3.2.6 Validating Agent Patch Errors

Here are some of the errors that may appear during the Validation phase of patch plans:

- **Problems Associating the Patches To the Plan**

  This error occurs when patches do not match the targets selected.

  For example if the patch is released for Linux x86 platform, and you are trying to apply the patch on a Linux x64 target, then the plan fails.

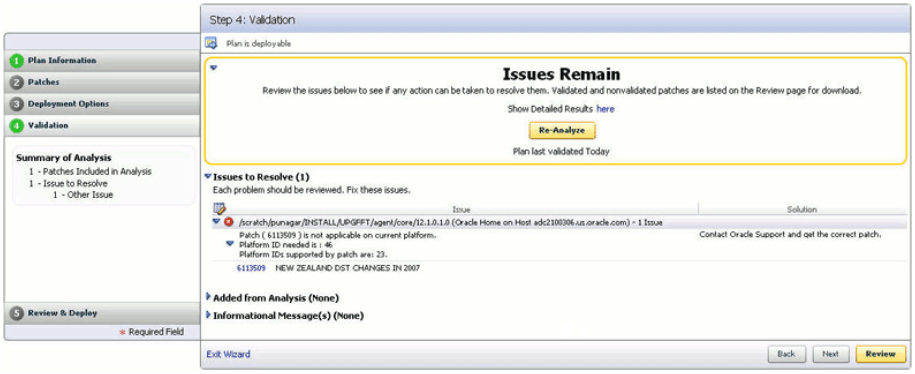

- **Oracle Home Credentials Not Set**

  This error occurs when the oracle Home credentials like Privileged Oracle Home credentials or the Normal Oracle Home credentials are not set.

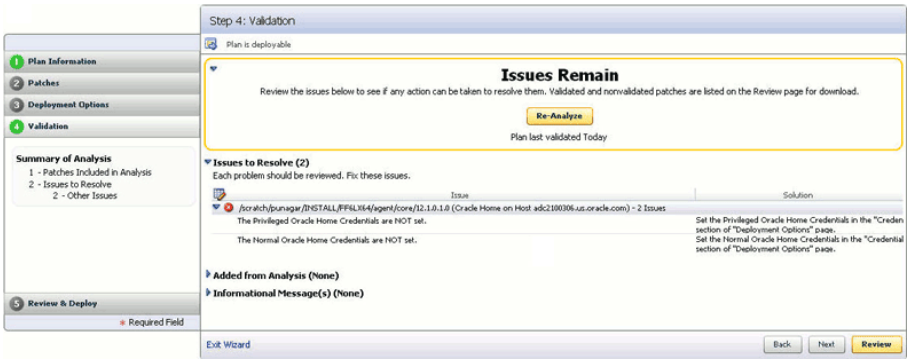

- **Conflict Check Analysis Failure**

  This error occurs when there is a conflict in the patches added.

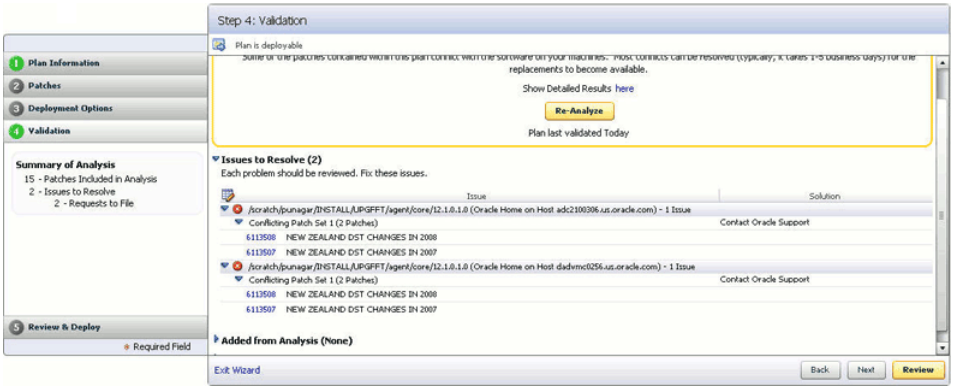

16.3.2.7 Deinstalling the Applied Agent Patches

To roll back applied Agent patches, follow the deinstallation steps available in the `Readme.html` or `Readme.txt` for the patch.

16.3.3 Manual Agent Patching

Manual patching is a patching mechanism that requires you to follow step-by-step instructions to patch a Management Agent manually. This mechanism of patching expects you to meet certain prerequisites, manually validate the patch for applicability and conflicts, and patch only one Management Agent at a time.

> **Note:** Oracle recommends you to use the automated patching mechanism because it not only saves time and effort in mass-deploying patches but also reduces human intervention, thereby minimizing the errors involved while patching.

To patch manually, you must perform the following steps:

1. Log into *My Oracle Support* (`https://support.oracle.com`) console with the necessary credentials.

> **Note:** Check the Patch Recommendation region to view the patches recommended for your environment. You can also provide the recommended patch number in the patch search region to download the recommended patch.

2. On the *My Oracle Support* home page, click **Patches and Updates.**
3. Enter the patch number in the Patch Search region, then click **Search.**
4. Select the patch, and from the context menu, click **Download**.
5. After downloading the zip file, follow the instructions available in the `Readme.html` or `Readme.txt` to install the patch.

17

Starting and Stopping Enterprise Manager Compone nts

This chapter explains how to use the Enterprise Manager command line utility (emctl) to start and stop the Management Service, the Management Agent, and Cloud Control.

This chapter also explains the various emctl commands for Management Service and Management Agent and how to use log information to troubleshoot emctl.

Following are the sections in this chapter:

- Controlling the Oracle Management Agent
- Controlling the Oracle Management Service
- Guidelines for Starting Multiple Enterprise Manager Components on a Single Host
- Starting and Stopping Oracle Enterprise Manager 12c Cloud Control
- Additional Management Agent Commands
- emctl Commands
- Using emctl.log File

17.1 Controlling the Oracle Management Agent

The following sections describe how to use the Enterprise Manager command line utility (`emctl`) to control the Oracle Management Agent:

- Starting, Stopping, and Checking the Status of the Management Agent on UNIX
- Starting and Stopping the Management Agent on Windows
- Checking the Status of the Management Agent on Windows
- Troubleshooting Management Agent Startup Errors

17.1.1 Starting, Stopping, and Checking the Status of the Management Agent on UNIX

When you start the agent on UNIX systems, it starts the parent watchdog process and the child Java process for the agent. The watchdog monitors the agent Java process and attempts to start it if it fails abnormally.

To start, stop, or check the status of the Management Agent on UNIX systems:

1. Change directory to the `AGENT_INSTANCE_HOME/bin` directory.
2. Use the appropriate command described in Table 17–1.

   For example, to stop the Management Agent, enter the following commands:

```
$PROMPT> cd AGENT_INSTANCE_HOME/bin
$PROMPT> emctl stop agent
```

Table 17–1 Starting, Stopping, and Checking the Status of the Management Agent

Command	Purpose
emctl start agent	Starts the Management Agent
emctl stop agent	Stops the Management Agent
emctl status agent	If the Management Agent is running, this command displays status information about the Management Agent, including the Agent Home, the process ID, and the time and date of the last successful upload to the Management Repository (Example 17–1).

Example 17–1 Checking the Status of the Management Agent

```
$ emctl status agent
Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
---------------------------------------------------------------
Agent Version     : 12.1.0.1.0
OMS Version       : 12.1.0.1.0
Protocol Version  : 12.1.0.1.0
Agent Home        : /ade/jexample_username/oracle/work/agentStateDir
Agent Binaries    : /scratch/user/view_storage/example_
username/emagent/gcagent/agent
Agent Process ID  : 22299
Parent Process ID : 22226
Agent URL         : https://example.us.oracle.com:11852/emd/main/
Repository URL    : https://example.us.oracle.com:14487/empbs/upload
Started at        : 2011-08-09 06:19:12
Started by user   : user
Last Reload       : (none)
Last successful upload                       : (none)
Last attempted upload                        : (none)
Total Megabytes of XML files uploaded so far : 0
Number of XML files pending upload           : 787
Size of XML files pending upload(MB)         : 2.21
Available disk space on upload filesystem    : 47.91%
Collection Status                            : Collections enabled
Last attempted heartbeat to OMS              : 2011-08-09 06:20:51
Last successful heartbeat to OMS             : (none)

---------------------------------------------------------------
```

On IBM AIX environment with a large memory configuration where the Management Agent is monitoring a large number of targets, the Agent may not start. To prevent this issue, prior to starting the Management Agent, add the following parameters to the common environment file:

```
LDR_CNTRL="MAXDATA=0x80000000"@NOKRTL
AIX_THREADSCOPE=S
```

The LDR_CNTRL variable sets the data segment size and disables loading of run time libraries in kernel space. The AIX_THREADSCOPE parameter changes AIX Threadscope context from the default Processwide 'P' to Systemwide 'S'. This causes less mutex contention.

17.1.2 Starting and Stopping the Management Agent on Windows

When you install the Oracle Management Agent on a Windows system, the installation procedure creates one new service in the Services control panel.

The procedure for accessing the Services control panel varies, depending upon the version of Microsoft Windows you are using. For example, on Windows 2000, locate the Services Control panel by selecting **Settings** and then **Administrative Tools** from the **Start** menu.

> **Note:** The `emctl` utility described in Section 17.2.1 is available in the `bin` subdirectory of the Oracle home where you installed the Management Agent; however, Oracle recommends that you use the Services control panel to start and stop the Management Agent on Windows systems.

Table 17–2 describes the Windows service that you use to control the Management Agent.

Table 17–2 Service Installed and Configured When You Install the Management Agent on Windows

Component	Service Name Format	Description

Oracle Management Agent	`Oracle<agent_home>Agent`
	For example:
	`OracleOraHome1Agent`
		Use this to start and stop the Management Agent.

> **Note:** If you are having trouble starting or stopping the Management Agent on a Windows NT system, try stopping the Management Agent using the following emctl command:
>
> ```
> $PROMPT> <AGENT_HOME>\bin\emctl istop agent
> ```
>
> After stopping the Management Agent using the emctl istop agent command, start the Management Agent using the Services control panel.
>
> This problem and solution applies only to the Windows NT platform, not to other Windows platforms, such as Windows 2000 or Windows XP systems.

17.1.3 Checking the Status of the Management Agent on Windows

To check the status of the Management Agent on Windows systems:

1. Change directory to the following location in the `AGENT_INSTANCE_HOME` directory:

   AGENT_INSTANCE_HOME`\bin`

2. Enter the following emctl command to check status of the Management Agent:

   ```
   $PROMPT> emctl status agent
   ```

   If the Management Agent is running, this command displays status information about the Management Agent, including the Agent Home, the process ID, and the

time and date of the last successful upload to the Management Repository (Example 17–1).

17.1.4 Troubleshooting Management Agent Startup Errors

If the agent fails to start, see the `emctl.log` and `emagent.nohup` log files for details. Following are common issues and troubleshooting suggestions:

17.1.4.1 Management Agent starts up but is not ready

The Management Agent goes through the following process when it starts up:

1. Starting up (the Management Agent has just received the request to start up and is going to start the initialization sequence)
2. Initializing (the Management Agent is iterating over each of its components and is initializing them)
3. Ready (All components have been initialized and the Management Agent is ready to accept requests)

The command to start the Management Agent (`emctl start agent`) has a default timeout of 120 seconds. At the end of that timeout, it will return control to the caller and will indicate what the last state of the Management Agent was when it returns control. Depending on the number of targets being monitored by the Management Agent, step 2 listed above could take a long time and it is possible that when the command exits, the state of the agent is "Initializing" and the command reports that the "agent is running but is not ready".

You can increase the timeout by setting an environment variable "EMAGENT_TIME_FOR_START_STOP". The value should indicate the number of seconds to wait before returning control to the caller.

17.1.4.2 Management Agent fails to start because of time zone mismatch between agent and OMS

The Management Agent uses the time zone set in emd.properties file. During the install process of the Management Agent, the agent and the host target are registered with the OMS along with the time zone. If the Management Agent's time zone is modified at any point after the installation, the OMS will signal the Management Agent to shut down as soon as it detects this mismatch.

To reset the Management Agent's time zone, run the following command:

```
emctl resettz agent
```

For more information about setting the time zone for the agent, see Section 17.5.5, "Changing the Management Agent Time Zone".

17.1.4.3 Agent fails to start due to possible port conflict

If the Management Agent cannot start and emctl reports that there is a possible port conflict, check the Management Agent's port (based on emd.properties:EMD_URL) and see if there is another application, such as another agent, running on the machine that is already bound to the port.

To resolve this issue, stop the application currently bound to the Management Agent's port.

17.1.4.4 Agent secure/unsecure fails

Securing or unsecuring of the Management Agent can fail if the password to secure the agent against the OMS is incorrect or if the OMS is locked or down. You can find the reason for the failure in the `<agent state directory>/sysman/log/secure.log` file.

17.2 Controlling the Oracle Management Service

When you start the Management Service, the following services are started:

1. OPMN process. This is the watchdog for the Apache process. The OPMN process starts the Apache process if it crashes.
2. Apache processes to start the HTTP server
3. Node Manager Java process. This is the watchdog for the Managed Server and Admin Server processes. It restarts the Managed Server and Admin Server processes if they crash.
4. Admin Server Java process (if the command to start OMS is executed on the first OMS machine). This is the WebLogic Server instance that maintains configuration data for configured Enterprise Manager domain.
5. Managed Server Java process. This is the Managed WebLogic Server on which Enterprise Manager application is deployed.
6. (On Windows only) Node Manager service process. This is the Windows service for starting and stopping the Node Manager (equivalent to the Node Manager process on Linux).
7. (On Windows only) OMS service process. This is the Windows service for starting and stopping the OMS.

The following sections describe how to control the Oracle Management Service:

- Controlling the Management Service on UNIX
- Controlling the Management Service on Windows

17.2.1 Controlling the Management Service on UNIX

To start, stop, or check the status of the Management Service with the Enterprise Manager command-line utility, follow these steps:

1. Change directory to the `ORACLE_HOME/bin` directory in the Management Service home.
2. Use the appropriate command described in Table 17–3.

   For example, to stop the Management Service, enter the following commands:

```
$PROMPT> cd bin
$PROMPT> ./emctl stop oms
```

Table 17–3 Starting, Stopping, and Checking the Status of the Management Service

Command	Purpose
emctl start oms	Starts the Fusion Middleware components required to run the Management Service. Specifically, this command starts HTTP Server, the Node Manager, OPMN process, and the managed server on which the Management Service is deployed. In addition if this command is run on the host that has the Administration Server, the Administration Server is started too.

Table 17–3 (Cont.) Starting, Stopping, and Checking the Status of the Management

Command	Purpose
emctl stop oms	Stops the OMS managed server and HTTP server but leaves Node Manager and Administration Server running. **Note:** The `emctl stop oms` command does not stop Fusion Middleware. Use `emctl stop oms -all` to stop all processes including Administration Server, HTTP Server, Node Manager, and management server.
emctl status oms	Displays a message indicating whether or not the Management Service is running. Run `emctl status oms -details` to view information about the configuration of the management service such as ports being used and the URLs for console and upload.

17.2.2 Controlling the Management Service on Windows

When you install the Oracle Management Service on a Windows system, the installation procedure creates three new services in the Services control panel.

The procedure for accessing the Services control panel varies, depending upon the version of Microsoft Windows you are using. For example, on Windows 2000, locate the Services control panel by selecting **Settings,** then **Administrative Tools** from the **Start** menu.

Note: The `emctl` utility described in Section 17.2.1 is available in the `bin` subdirectory of the Oracle home where you installed the Management Service; however, Oracle recommends that you use the Services control panel to start and stop the Management Service on Windows systems.

Table 17–4 describes the Windows services that you use to control the Oracle Management Service.

Table 17–4 Summary of Services Installed and Configured When Installing the Oracle Management Service on Windows

Component	Service Name Format	Description
WebLogic Server	`OracleWeblogicNodeManager_EMGC_OMS1_1`	Use this service to start and stop the node manager of the WebLogic Server that was installed and configured to deploy the Management Service J2EE application.
Oracle Management Server	`OracleManagementServer_EMGC_OMS1_1`	Use this service to start and stop all components that were installed and configured as part of the Management Service J2EE application.

17.2.3 Troubleshooting Oracle Management Service Startup Errors

Following are the log files you can check if the OMS fails to start:

Management Service Fails to Start

Check the logs located as indicated in Table 17–5. The INSTANCE_HOME mentioned in the table is the OMS instance home and n is the index of the OMS server.

Table 17–5 OMS Log Files Location

OMS Log File	Log File Location
emctl log file	$INSTANCE_HOME/sysman/log/emctl.log file
Managed Server log files	$INSTANCE_HOME/user_projects/domains<DOMAIN_NAME>/servers/EMGC_OMS<n>/logs/EMGC_OMS<n>.log
	$INSTANCE_HOME/user_projects/domains<DOMAIN_NAME>/servers/EMGC_OMS<n>/logs/EMGC_OMS<n>.out
OMS log files	$INSTANCE_HOME/sysman/log/emoms_pbs.log
	$INSTANCE_HOME/sysman/log/emoms_pbs.trc
	$INSTANCE_HOME/sysman/log/emoms.trc
	$INSTANCE_HOME/sysman/log/emoms.log

WebTier Service Fails to Start

Check logs under `<WebTier Instance Home>/diagnostics` folder in case WebTier start fails.

17.3 Guidelines for Starting Multiple Enterprise Manager Components on a Single Host

Oracle Enterprise Manager components are used to manage a variety of Oracle software products. In most cases, in a production environment, you will want to distribute your database and WebLogic Server instances among multiple hosts to improve performance and availability of your software resources. However, in cases where you must install multiple WebLogic Servers or databases on the same host, consider the following guidelines.

When you start Fusion Middleware Control, the Management Agent, or Database Control, Enterprise Manager immediately begins gathering important monitoring data about the host and its managed targets. Keep this in mind when you develop a process for starting the components on the host.

Specifically, consider staggering the startup process so that each Enterprise Manager process has a chance to start before the next process begins its startup procedure. When you start up all the components (for example, after a restart of the system), use a process such as the following:

1. Use the `emctl start` command to start Oracle Management Service.
2. Wait 15 seconds.
3. Use the `emctl start agent` command to start the Management Agent for the host.

Using a staggered startup procedure such as the preceding example will ensure that the processes are not in contention for resources during the CPU-intensive startup phase for each component.

17.4 Starting and Stopping Oracle Enterprise Manager 12c Cloud Control

As described in the previous sections, you use separate commands to control the Oracle Management Service and Management Agents.

The following sections describe how to stop and start all the Cloud Control components that are installed by the Oracle Enterprise Manager 12c Cloud Control Console installation procedure.

You can use this procedure to start all the framework components after a system reboot or to shutdown all the components before bringing the system down for system maintenance.

17.4.1 Starting Cloud Control and All Its Components

The following procedure summarizes the steps required to start all the components of the Cloud Control. For example, use this procedure if you have restarted the host computer and all the components of the Cloud Control have been installed on that host.

To start all the Cloud Control components on a host, use the following procedure:

1. If your Oracle Management Repository resides on the host, change directory to the Oracle Home for the database where you installed the Management Repository and start the database and the Net Listener for the database:
   a. Set the ORACLE_HOME environment variable to the Management Repository database home directory.
   b. Set the ORACLE_SID environment variable to the Management Repository database SID (default is asdb).
   c. Start the Net Listener:

      ```
      $PROMPT> $ORACLE_HOME/bin/lsnrctl start
      ```

   d. Start the Management Repository database instance:

      ```
      ORACLE_HOME/bin/sqlplus /nolog
      SQL> connect SYS as SYSDBA
      SQL> startup
      SQL> quit
      ```

      See Also: *Oracle Database Administrator's Guide* for information about starting and stopping an Oracle Database.

2. Start the Oracle Management Service:

   ```
   $PROMPT> OMS_HOME/bin/emctl start oms
   ```

   See Also: "Controlling the Oracle Management Service" on page 17-5

3. Change directory to the home directory for the Oracle Management Agent and start the Management Agent:

   ```
   $PROMPT> AGENT_HOME/bin/emctl start agent
   ```

   See Also: "Controlling the Oracle Management Agent" on page 17-1

Note: Be sure to run the `emctl start agent` command in the Oracle Management Agent home directory and not in the Management Service home directory.

17.4.2 Stopping Cloud Control and All Its Components

The following procedure summarizes the steps required to stop all the components of the Cloud Control. For example, use this procedure if you have installed all the components of the Cloud Control on the same host you want to shut down or restart the host computer.

To stop all the Cloud Control components on a host, use the following procedure:

1. Stop the Oracle Management Service:

```
$PROMPT> $ORACLE_HOME/bin/emctl stop oms -all
```

See Also: "Controlling the Oracle Management Service" on page 17-5

2. Change directory to the home directory for the Oracle Management Agent and stop the Management Agent:

```
$PROMPT> AGENT_HOME/bin/emctl stop agent
```

See Also: "Controlling the Oracle Management Agent" on page 17-1

Note: Be sure to run the `emctl stop agent` command in the Oracle Management Agent home directory and not in the Oracle Management Service home directory.

3. If your Oracle Management Repository resides on the same host, change directory to the Oracle Home for the database where you installed the Management Repository and stop the database and the Net Listener for the database:

   a. Set the ORACLE_HOME environment variable to the Management Repository database home directory.

   b. Set the ORACLE_SID environment variable to the Management Repository database SID (default is asdb).

   c. Stop the database instance:

```
$PROMPT> ORACLE_HOME/bin/sqlplus /nolog
SQL> connect SYS as SYSDBA
SQL> shutdown
SQL> quit
```

See Also: *Oracle Database Administrator's Guide* for information about starting and stopping an Oracle Database.

   d. Stop the Net Listener:

```
$PROMPT> $ORACLE_HOME/bin/lsnrctl stop
```

17.5 Additional Management Agent Commands

The following sections describe additional `emctl` commands you can use to control the Management Agent:

- Uploading and Reloading Data to the Management Repository
- Specifying New Target Monitoring Credentials
- Listing the Targets on a Managed Host
- Controlling Blackouts

17.5.1 Uploading and Reloading Data to the Management Repository

Under normal circumstances, the Management Agent uploads information about your managed targets to the Management Service at regular intervals.

To use these commands, change directory to the `AGENT_HOME/bin` directory (UNIX) or the `AGENT_HOME\bin` directory (Windows) and enter the appropriate command as described in Table 17–6.

Table 17–6 Manually Reloading and Uploading Management Data

Command	Description
`emctl upload (agent)`	Use this command to force an immediate upload of the current management data from the managed host to the Management Service. Use this command instead of waiting until the next scheduled upload of the data.
`emctl reload` (agent)	This command can be used to apply the changes after you have manually modified the `emd.properties` file. For example, to change the upload interval, emd.properties can be modified, and `emctl reload` can then be run. **Note:** Oracle does not support manual editing of the `targets.xml` files unless the procedure is explicitly documented or you are instructed to do so by Oracle Support.

17.5.2 Specifying New Target Monitoring Credentials

To monitor the performance of your database targets, Enterprise Manager connects to your database using a database user name and password. This user name and password combination is referred to as the database monitoring credentials.

> **Note:** The instructions in this section are specific to the monitoring credentials for a database target, but you can use this procedure for any other target type that requires monitoring credentials. For example, you can use this procedure to specify new monitoring credentials for your Oracle Management Service and Management Repository.

When you first add an Oracle9i Database target, or when it is added for you during the installation of the Management Agent, Enterprise Manager uses the DBSNMP database user account and the default password for the DBSNMP account as the monitoring credentials.

When you install Oracle Database 11*g*, you specify the DBSNMP monitoring password during the database installation procedure.

As a result, if the password for the DBSNMP database user account is changed, you must modify the properties of the database target so that Enterprise Manager can continue to connect to the database and gather configuration and performance data.

Similarly, immediately after you add a new Oracle Database 11*g* target to the Cloud Control, you may need to configure the target so it recognizes the DBSNMP password that you defined during the database installation. Otherwise, the Database Home page may display no monitoring data and the status of the database may indicate that there is a metric collection error.

Note: You can modify the Enterprise Manager monitoring credentials by using the Oracle Enterprise Manager 12*c* Cloud Control Console.

17.5.3 Listing the Targets on a Managed Host

There are times when you need to provide the name and type of a particular target you are managing. For example, you must know the target name and type when you are setting the monitoring credentials for a target.

To list the name and type of each target currently being monitored by a particular Management Agent:

1. Change directory to the `AGENT_HOME/bin` directory (UNIX) or the `AGENT_HOME\bin` directory (Windows).
2. Enter the following command to specify new monitoring credentials:

   ```
   $PROMPT>./emctl config agent listtargets
   ```

   Example 17–2 shows the typical output of the command.

Example 17–2 Listing the Targets on a Managed Host

```
ade:[ exampname_1208_qc_ag ] [example_username@example emagent]$ emctl config
agent listtargets
Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
[example.us.oracle.com:11852, oracle_emd]
[example.us.oracle.com, host]
[chronos_test, oracle_webcache]
[chronos_apache_test, oracle_apache]
[mytestBeacon, oracle_beacon]
[CSAcollector, oracle_csa_collector]
[database, oracle_database]
[database2, oracle_database]
[database3, oracle_database]
[listener, oracle_listener]
[listener2, oracle_listener]
[listener3, oracle_listener]
[Management Services and Repository, oracle_emrep]
ade:[ example_username_1208_qc_ag ] [example_username@example emagent]$
```

17.5.4 Controlling Blackouts

Blackouts allow Enterprise Manager users to suspend management data collection activity on one or more managed targets. For example, administrators use blackouts to prevent data collection during scheduled maintenance or emergency operations.

You can control blackouts from the Oracle Enterprise Manager 12c Cloud Control Console or from the Enterprise Manager command line utility (`emctl`). However, if you are controlling target blackouts from the command line, you should not attempt to control the same blackouts from the Cloud Control Console. Similarly, if you are controlling target blackouts from the Cloud Control Console, do not attempt to control those blackouts from the command line.

> **See Also:** "Creating, Editing, and Viewing Blackouts" in the Enterprise Manager online help for information about controlling blackouts from the Cloud Control Console

From the command line, you can perform the following blackout functions:

- Starting Immediate Blackouts
- Stopping Immediate Blackouts
- Checking the Status of Immediate Blackouts

> **Note:** When you start a blackout from the command line, any Enterprise Manager jobs scheduled to run against the blacked out targets will still run. If you use the Cloud Control Console to control blackouts, you can optionally prevent jobs from running against blacked out targets.

To use the Enterprise Manager command-line utility to control blackouts:

1. Change directory to the `AGENT_HOME/bin` directory (UNIX) or the `AGENT_INSTANCE_HOME\bin` directory (Windows).
2. Enter the appropriate command as described in Table 17–7.

> **Note:** When you start a blackout, you must identify the target or targets affected by the blackout. To obtain the correct target name and target type for a target, see Section 17.5.3, "Listing the Targets on a Managed Host".

Table 17–7 Summary of Blackout Commands

Blackout Action	Command
Set an immediate blackout on a particular target or list of targets	`emctl start blackout <Blackoutname> [<Target_name>[:<Target_Type>]]....[-d <Duration>]` Be sure to use a unique name for the blackout so you can refer to it later when you want to stop or check the status of the blackout. The `-d` option is used to specify the duration of the blackout. Duration is specified in `[days] hh:mm where:` ■ days indicates number of days, which is optional ■ hh indicates number of hours ■ mm indicates number of minutes If you do not specify a target or list of targets, Enterprise Manager will blackout the local host target. All monitored targets on the host are not blacked out unless a list is specified or you use the `-nodelevel` argument. If two targets of different target types share the same name, you must identify the target with its target type.
Stop an immediate blackout	`emctl stop blackout <Blackoutname>`
Set an immediate blackout for all targets on a host	`emctl start blackout <Blackoutname> [-nodeLevel] [-d <Duration>]` The `-nodeLevel` option is used to specify a blackout for all the targets on the host; in other words, all the targets that the Management Agent is monitoring, including the Management Agent host itself. The `-nodeLevel` option must follow the blackout name. If you specify any targets after the `-nodeLevel` option, the list is ignored.
Check the status of a blackout	emctl status blackout `[<Target_name>[:<Target_Type>]]....`

Use the following examples to learn more about controlling blackouts from the Enterprise Manager command line:

- To start a blackout called "bk1" for databases "db1" and "db2," and for Oracle Listener "ldb2," enter the following command:

  ```
  $PROMPT> emctl start blackout bk1 db1 db2 ldb2:oracle_listener -d 5 02:30
  ```

  The blackout starts immediately and will last for 5 days 2 hours and 30 minutes.

- To check the status of all the blackouts on a managed host:

  ```
  $PROMPT> emctl status blackout
  ```

- To stop blackout "bk2" immediately:

  ```
  $PROMPT> emctl stop blackout bk2
  ```

- To start an immediate blackout called "bk3" for all targets on the host:

  ```
  $PROMPT> emctl start blackout bk3 -nodeLevel
  ```

- To start an immediate blackout called "bk3" for database "db1" for 30 minutes:

  ```
  $PROMPT> emctl start blackout bk3 db1 -d 30
  ```

- To start an immediate blackout called "bk3" for database "db2" for five hours:

```
$PROMPT> emctl start blackout bk db2 -d 5:00
```

17.5.5 Changing the Management Agent Time Zone

The Management Agent may fail to start after the upgrade if it realizes that it is no longer in the same time zone that it was originally configured with.

You can reset the time zone used by the Management Agent using the following command:

```
emctl resetTZ agent
```

This command will correct the Management Agent side time zone and specify an additional command to be run against the Management Repository to correct the value there.

> **IMPORTANT: Before you change the Management Agent time zone, first check to see if there are any blackouts that are currently running or scheduled to run on any target managed by that Management Agent.**

To check for blackouts:

1. From the Cloud Control home page, click **Targets** and then **All Targets**. In the All Targets page, locate the Management Agent in the list of targets. Click on the Management Agent's name. This brings you to the Management Agent's home page.
2. The list of targets monitored by the Management Agent are listed in the Monitored Targets section.
3. For each of target in the list:
   a. Click the target name. This brings you to the target's home page.
   b. From the <Target> menu, select **Monitoring** and then click **Blackouts**. This allows you to check any currently running blackouts or blackouts that are scheduled in the future for this target.

If such blackouts exist, then:

1. From the Cloud Control Console, stop all currently running blackouts on all targets monitored by that Management Agent.
2. Stop all scheduled blackouts on all targets monitored by that Management Agent.

Once you have stopped all currently running and scheduled blackouts, you can run the `emctl resetTZ agent` command to change the Management Agent's time zone.

Once you have changed the Management Agent's time zone, create new blackouts on the targets as needed.

17.5.6 Reevaluating Metric Collections

Use the following command to perform an immediate reevaluation of a metric collection:

```
emctl control agent runCollection <targetName>:<targetType> <colletionItemName>
```

where <collectionItemName> is the name of the Collection Item that collects the metric.

Performing this command causes the reevaluated value of the metric to be uploaded into the Management Repository, and possibly trigger alerts if the metric crosses its threshold.

Related metrics are typically collected together; collectively a set of metrics collected together is called a Metric Collection. Each Metric Collection has its own name. If you want to reevaluate a metric, you first need to determine the name of the Metric Collection to which it belongs, then the CollectionItem for that Metric Collection.

When you run the previous command to reevaluate the metric, all other metrics that are part of the same Metric Collection and Collection Item will also be reevaluated.

Perform the following steps to determine the Metric Collection name and Collection Item name for a metric:

1. Go to $INSTALL_BASE/ngagent/plugins directory, where $INSTALL_BASE is the root of the installation. The Oracle Home of the Management Agent exists in this directory.
2. Locate the XML file for the target type. For example, if you are interested in the host metric 'Filesystem Space Available(%)' metric, look for the host.xml file.
3. In the xml file, look for the metric in which you are interested. The metric that you are familiar with is actually the display name of the metric. The metric name would be preceded by a tag that started with:

   ```
   <Label NLSID=
   ```

   For example, in the host.xml file, the metric 'Filesystem Space Available(%)" would have an entry that looks like this:

   ```
   <Label NLSID="host_filesys_pctAvailable">Filesystem Space Available (%)
   </Label>
   ```

4. Once you have located the metric in the xml file, you will notice that its entry is part of a bigger entry that starts with:

   ```
   <Metric NAME=
   ```

   Take note of the value defined for "Metric NAME". This is the Metric Collection name. For example, for the 'Filesystem Space Available(%)' metric, the entry would look like this:

   ```
   <Metric NAME="Filesystems"
   ```

   So for the 'Filesystem Space Available(%)' metric, the Metric Collection name is 'Filesystems'.
5. The Collection Item name for this Metric Collection needs to be determined next. Go to the $INSTALL_BASE/ngagent/plugins/default_collection directory, where $INSTALL_BASE is the Oracle Home of the Management Agent.
6. In this directory, look for the collection file for the target type. In our example, this would be host.xml.
7. In cases where a Metric Collection is collected by itself, there would be a single Collection Item of the same name in the collection file. To determine if this is the case for your Metric Collection, look for an entry in the collection file that starts with:

   ```
   <CollectionItem NAME=
   ```

where the value assigned to the CollectionItem NAME matches the Metric NAME in step (4).

For the 'Filesystem Space Available(%)' metric, the entry in the collection file would look like:

```
<CollectionItem NAME = "Filesystems"
```

8. If you find such an entry, then the value assigned to "CollectionItem NAME" is the collection item name that you can use in the emctl command.
9. Otherwise, this means the Metric Collection is collected with other Metric Collections under a single Collection Item. To find the Collection Item for your Metric Collection, first search for your Metric Collection. It should be preceded by the tag:

   `<MetricColl NAME=`

   Once you have located it, look in the file above it for: `<CollectionItem NAME=`

   The value associated with the CollectionItem NAME is the name of the collection item that you should use in the emctl command.

   For example if the you want to reevaluate the host metric "Open Ports", using the previous steps, you would do the following:

   a. Go to the `$INSTALL_BASE/ngagent/plugins` directory where $INSTALL_BASE is the Oracle Home of the Management Agent. Look for the host.xml file and in that file locate: `<Metric NAME="openPorts".`
   b. Then go to the `$INSTALL_BASE/ngagent/plugins/default_collection` directory. Look for the host.xml file and in that file look for `<CollectionItem NAME="openPorts".`

      Failing this, look for `<MetricColl NAME="openPorts".`
   c. Look above this entry in the file to find the `<CollectionItem NAME=` string and find `<CollectionItem NAME="oracle_security".`

   The CollectionItem NAME oracle_security is what you would use in the emctl command to reevaluate the Open Ports metric.

17.6 emctl Commands

This section lists the emctl commands for the Enterprise Manager Management Agent and Oracle Management Service.

Table 17–8 explains the emctl commands for OMS.

Table 17–8 emctl Commands for OMS

emctl Command	Description

emctl [getversion] oms	Gets the version of the Management Service. Sample output is as follows:
	`./emctl getversion oms`
	`Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`Enterprise Manager 12c OMS Version 12.1.0.1.0`

emctl [start] oms	Starts the Management Service.

Table 17–8 (Cont.) emctl Commands for OMS

emctl Command	Description
emctl stop oms -all	Stops the Management Service including Administration Server, HTTP Server, Node Manager, and management server.
emctl stop oms -all -force and emctl stop oms -force	-force can be used with both emctl stop oms -all and emctl stop oms. If the emctl stop oms commands do not shutdown the relevant processes, using -force option will forcefully stop the relevant processes.
emctl status oms	Lists the status of the Management Service
emctl status oms -details	Lists Management Service details such as port numbers, lock status, domain information, and so on.
emctl config oms -list_ repos_details	Lists the Management Service repository details.
emctl config oms -store_ repos_details [-repos_host <host> -repos_port <port> -repos_sid <sid> | -repos_ conndesc <connect descriptor>] -repos_user <username> [-repos_pwd <pwd>]	Configures the settings used by Management Service to connect to the Management Repository.
emctl config oms -change_ repos_pwd [-old_pwd <old_pwd>] [-new_pwd <new_pwd>] [-use_sys_ pwd [-sys_pwd <sys_ pwd>]]	Configures the password used by Management Service to connect to the Management schema in the Management Repository.
emctl config oms -change_ view_user_pwd [-sysman_ pwd <sysman_pwd>] [-user_pwd <user_pwd>] [-auto_generate]	Configures the password used by Management Service for MGMT_VIEW user that is used for report generation.
emctl upload	Uploads xml files that are pending to upload to the OMS under the upload directory.

Table 17–9 explains emctl commands for Management Agents.

Table 17–9 emctl Commands for Management Agent

emctl Command	Description
emctl start | stop agent	Starts or stops Management Agent.
emctl status agent	Lists the status of Management Agent.

Table 17–9 (Cont.) emctl Commands for Management Agent

emctl Command	Description

emctl status agent -secure	Lists the secure status of the agent and the port on which the agent is running in secure mode and also the OMS security status of the agent it points to. This command also gives the OMS secure port. Below is an example output:
	`bash-3.00$ emctl status agent -secure`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`Checking the security status of the Agent at location set in /ade/example_username_cpap4_ag/oracle/sysman/config/emd.properties... Done.`
	`Agent is secure at HTTPS Port 1838.`
	`Checking the security status of the OMS at http://example.us.oracle.com:7654/em/upload/... Done.`
	`OMS is secure on HTTPS Port 4473`
	`bash-3.00$`

emctl status agent scheduler	Lists all Running, Ready, and Scheduled Collection threads.

emctl status agent jobs	Lists the status of the jobs that are running at present on the agent. The following is an example output:
	`bash-3.00$ emctl status agent jobs`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`----------------------------------------------------`
	`step id typ pid stat command line`
	`------- --- --- ---- ------------`
	`----------------------------------------------------`
	`Agent is Running and Ready`

Table 17–9 (Cont.) emctl Commands for Management Agent

emctl Command	Description
emctl status agent target <target name>,<target type>,<metric>	Lists the detailed status of the specified targets in the order of target name, target type. The following is an example of an oracle_database target. You can also provide a particular metric name in the emctl command to get the status of a particular metric of a target.

```
bash-3.00$ emctl status agent target database,oracle_
database
Oracle Enterprise Manager 12c Release 1 Cloud Control
12.1.0.1.0.
Copyright (c) 1996, 2011 Oracle Corporation.  All
rights reserved.
---------------------------------------------------
Target Name : database
Target Type : oracle_database
Current severity state
----------------------
Metric      Column name    Key   State  Timestamp
---------------------------------------------------
DeferredTrans errortrans_count  n/a CLEAR 2011-07-09
02:38:07
DeferredTrans deftrans_count    n/a   CLEAR 2011-07-09
02:38:07
ha_recovery missing_media_files n/a   CLEAR 2011-07-09
02:28:57
ha_recovery corrupt_data_blocks n/a  CLEAR 2011-07-09
02:28:57
ha_recovery datafiles_need_recovery n/a CLEAR
2011-07-09 02:28:57
Response     Status        n/a     CLEAR  2011-07-09
02:38:04
Response userLogon         n/a CLEAR  2011-07-09
02:38:04
Response    State          n/a  CLEAR  2011-07-09
02:38:04
OCMInstrumentation  NeedToInstrument n/a CLEAR
2011-07-09 02:31:55
health_check  Status  n/a  CLEAR   2011-07-09 02:40:00
health_check  Unmounted  n/a  CLEAR  2011-07-09
02:40:00
health_check   Mounted  n/a CLEAR  2011-07-09 02:40:00
health_check Unavailable  n/a  CLEAR  2011-07-09
02:40:00
health_check Maintenance n/a  CLEAR  2011-07-09
02:40:00
sql_response       time  n/a  CLEAR  2011-07-09
02:38:50
sga_pool_wastage java_free_pct n/a  CLEAR  2011-07-09
02:28:
58
UserAudit username DBSNMP_example  CLEAR 2011-07-09
02:32:48
-----------------------------------------------------
Agent is Running and Ready
```

Table 17–9 (Cont.) emctl Commands for Management Agent

emctl Command	Description

emctl status agent mcache <target name>,<target type>,<metric>	Lists the names of the metrics for which the values are present in the metric cache. See the following example for a simple host target:
	`bash-3.00$ emctl status agent mcache example.us.oracle.com,host`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`---`
	`Metric cache contains value for following metrics at 2011-07-09 02:54:47`
	`CPUUsage`
	`DiskActivity`
	`FileMonitoring`
	`LPAR Performance on AIX`
	`Load`
	`Network`
	`PagingActivity`
	`---`
	`Agent is Running and Ready`
	
	The metrics listed above are the ones whose values are present in the metric cache.

emctl reload agent dynamicproperties [<Target_name>:<Target_Type>]...	Recomputes the dynamic properties of a target and generates the dynamic properties for the target.
	Sample output for oracle_database is as follows:
	`bash-3.00$ emctl reload agent dynamicproperties database:oracle_database`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`---`
	`EMD recompute dynprops completed successfully`

emctl pingOMS [agent]	Pings the OMS to check if the agent is able to connect to the OMS. Agent will wait for the reverse ping from the OMS so that agent can say the pingOMS is successful.
emctl config agent getTZ	Gets the current timezone set in the environment.
emctl config agent getSupportedTZ	Prints the supported timezone based on the setting in the environment.
emctl config console <fileloc> [<EM loc>]	Allows you to configure the console based on the configuration entries that you have mentioned in the file <fileloc>. <EM loc> is optional and can be used to operate on a different Oracle Home.
emctl config [agent] listtargets [<EM loc>]	Lists all targets present in targets.xml. <EM loc> is optional and can be used to operate on a different Oracle Home.

Table 17–9 (Cont.) emctl Commands for Management Agent

emctl Command	Description

emctl control agent runCollection <target_name>:<target_type> <metric_name>	Allows to manually run the collections for a particular metric of a target. Sample output is as follows:
	`emctl control agent runCollection example.us.oracle.com:host CPUUsage`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`---`
	`EMD runCollection completed successfully`

emctl resetTZ agent	Resets the timezone of the agent. Stop the agent first and then run this command to change the current timezone to a different timezone. Then start the agent.

emctl getversion agent	Prints the version of the agent. Sample output is as follows:
	`./emctl getversion agent`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control Agent 12.1.0.1.0`

emctl dumpstate agent <component> . . .	Generates the dumps for the agent. This command allow you to analyze the memory/cpu issues of the agent. Sample output is as follows:
	`./emctl dumpstate agent`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`Dumpstate succeeded`

emctl gensudoprops	Generates the sudo properties of the agent.
emctl clearsudoprops	Clears the sudo properties.
emctl clearstate	Clears the state directory contents. The files that are located under $ORACLE_HOME/sysman/emd/state will be deleted if this command is run. The state files are the files which are ready for the agent to convert them into corresponding xml files.

emctl getemhome	Prints the agent home directory. The sample output is as follows:
	`bash-3.00$ emctl getemhome`
	`Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0.`
	`Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.`
	`EMHOME=/scratch/aime/gcagent/ngagent/agent_inst`

emctl start blackout <Blackoutname> [-nodeLevel] [<Target_name>[:<Target_Type>]].... [-d <Duration>]	Starts blackout on a target. <Target_name:Target_type> defaults to local node target if not specified. If -nodeLevel is specified after <Blackoutname>,the blackout will be applied to all targets and any target list that follows will be ignored. Duration is specified in [days] hh:mm

Table 17–9 (Cont.) emctl Commands for Management Agent

emctl Command	Description
emctl stop blackout <Blackoutname>	Stops the blackout that was started on a particular target. Only those blackouts that are started by the emctl tool can be stopped using emctl. This command cannot stop the blackouts that are started using the Console or emcli.
emctl status blackout [<Target_name>[:<Target_Type>]]....	Provides the status of the blackout of the target. The status includes the type of blackout, whether one time, repeating, or a scheduled blackout. This command also specifies whether the blackout has started or stopped.
emctl secure agent [registration password]	Secures the agent against an OMS. The registration password must be provided.
emctl unsecure agent	Unsecures the agent. This will make the agent unsecure and the agent's port will be changed to http port.
emctl verifykey	Verifies the communication between the OMS and agent by sending pingOMS.
emctl deploy agent [-s <install-password>] [-o <omshostname:consoleSrvPort>] [-S] <deploy-dir> <deploy-hostname>:<port> <source-hostname>	'agent' creates and deploys only the agent. [-s <password>]: Install password for securing agent. [-S]: Password will be provided in STDIN. [-o <omshostname:consoleSrvPort>]: The OMS Hostname and console servlet port. Choose the unsecured port. <deploy-dir> : Directory to create the shared (state-only) installation port. <deploy-hostname:port> : Host name and port of the shared (state-only) installation. Choose unused port. <source-hostname>: The host name of the source install. Typically the machine where EM is installed. This is searched and replaced in targets.xml by the host name provided in argument <deploy-hostname:port>. <sid>: The instance of the remote database. Only specified when deploying "dbconsole".

17.7 Using emctl.log File

The `emctl.log` file is a file that captures the results of all emctl commands you run. For Management Agent, the log file resides in the `$AGENT_INSTANCE_HOME/sysman/log` directory of the Management Agent, and for Management Service, the log file resides in the `$OMS_INSTANCE_HOME/sysman/log` directory. The file is updated every time you run an emctl command. If your emctl command fails for some reason, access this log file to diagnose the issue.

For example, run the following command from the Oracle home directory of the Management Agent to check its status:

```
<agent_home>emctl status agent
```

After running the command, navigate to the log directory to view the following information in the `emctl.log` file:

```
1114306 :: Wed Jun 10 02:29:36 2011::AgentLifeCycle.pm: Processing status agent
1114306 :: Wed Jun 10 02:29:36 2011::AgentStatus.pm:Processing status agent
1114306 :: Wed Jun 10 02:29:37 2011::AgentStatus.pm:emdctl status returned 3
```

Here, the first column, that is, 1114306, is the PID that was used to check the status. The second column shows the date and time when the command was run. The third

column mentions the Perl script that was run for the command. The last column describes the result of the command, where it shows the progress made by the command and the exit code returned for the command. In this case, the exit code is 3, which means that the Management Agent is up and running.

Similarly, for the Management Service, you can run the following command from the Oracle home directory of the Management Service to check its status:

```
<agent_home>emctl status oms
```

In another example, run the following command from the Oracle home directory of the Management Agent to upload data:

```
<agent_home>emctl upload agent
```

After running the command, navigate to the log directory to view the following information in the `emctl.log` file:

```
1286220 :: Tue Jun  9 07:13:09 2011::AgentStatus.pm:Processing upload
1286220 :: Tue Jun  9 07:13:10 2011::AgentStatus.pm:emdctl status agent returned 3
1286220 :: Tue Jun  9 07:13:41 2011::AgentStatus.pm: emdctl upload returned with
exit code 6
```

Here, the entries are similar to the entries in the first example, but the exit code returned is 6, which means the upload operation is failing for some reason.

The exit codes returned depend on the emctl command executed. In general, exit code of zero means success and any exit code other than zero means failure. For details about the cause of failure, view the error message.

For more information about Management Agent and Oracle Management Service log files, see Chapter 18.

18

Locating and Configuring Enterprise Manager Log Files

When you install the Oracle Management Agent (Management Agent) or the Oracle Management Service (OMS), Enterprise Manager automatically configures the system to save certain informational, warning, and error information to a set of log files.

Log files can help you troubleshoot potential problems with an Enterprise Manager installation. They provide detailed information about the actions performed by Enterprise Manager and whether or not any warnings or errors occurred.

This chapter not only helps you locate and review the contents of Enterprise Manager log files, but also includes instructions for configuring the log files to provide more detailed information to help in troubleshooting or to provide less detailed information to save disk space.

This chapter contains the following sections:

- Managing Log Files
- Locating Management Agent Log and Trace Files
- Locating and Configuring Oracle Management Service Log and Trace Files

18.1 Managing Log Files

Many Enterprise Manager components generate log files containing messages that record errors, notifications, warnings, and traces.

Table 18–1 describes the columns in the Log Message table. For any given component, the optional column may not be populated in the message.

Table 18–1 Message Columns

Column Name	Description
Time	The date and time when the message was generated. This reflects the local time zone.
Message Type	The type of message. Possible values are: Incident Error Warning, Notification, and Trace. In addition, the value Unknown may be used when the type is not known.
Message ID	The ID that uniquely identifies the message within the component. The ID consists of a prefix that represents the component, followed by a dash, then a 5-digit number. For example: `OHS-51009`
Message	The text of the error message.

Table 18–1 (Cont.) Message Columns

Column Name	Description
Target (Expanded)	Expanded target name.
Target	Target name
Target Type	Target type
Execution Context	The Execution Context ID (ECID), which is a global unique identifier of the execution of a particular request in which the originating component participates. You can use the ECID to correlate error messages from different components.
	The Relationship ID, which distinguishes the work done in one thread on one process, from work done by any other threads on this and other processes, on behalf of the same request.
Component	The component that originated the message.
Module	The identifier of the module that originated the message.
Incident ID	The identifier of the incident to which this message corresponds.
Instance	The name of the Oracle instance to which the component that originated the message belongs.
Message Group	The name of the group to which this message belongs.
Message Level	The message level, represented by an integer value that qualifies the message type. Possible values are from 1 (highest severity) through 32 (lowest severity).
Hosting Client	The identifier for the client or security group to which this message relates.
Organization	The organization ID for the originating component. The ID is `oracle` for all Oracle components.
Host	The name of the host where the message originated.
Host IP Address	The network address of the host where the message originated.
User	The name of the user whose execution context generated the message.
Process ID	The ID for the process or execution unit that generated the message.
Thread ID	The ID of the thread that generated the message.
Upstream Component	The component that the originating component is working with on the client (upstream) side.
Downstream Component	The component that the originating component is working with on the server (downstream) side.
Detail Location	A URL linking to additional information regarding the message.
Supplemental Detail	Supplemental information about the event, including more detailed information than the message text.
Target Log Files	Link to the log files page for this target.
Log File	Log file that this message contains.

Using Log Viewer, you can do the following:

- Viewing Log Files and Their Messages
- Searching Log Files
- Downloading Log Files

18.1.1 Viewing Log Files and Their Messages

You can view Enterprise Manager-related logs in a domain, an Oracle WebLogic Server, or a component.

For example, to view the log files and their messages:

1. From the **Enterprise** menu, select **Monitoring**, then select **Logs**.

   or

   From the **Targets** menu, select **Middleware**, then select an Enterprise Manager domain. From the **Farm** menu, select **Logs**, then select **View Log Messages**.

   The Log Messages page is displayed.

2. If the targets for which you want to view logs are added, expand **Selected Targets** and in the row for a particular component or application, click the **Target Log Files** icon.

   The Log Files page is displayed. On this page, you can see a list of log files related to the Managed Server.

3. Select a file and click **View Log File.**

   The View Log File page is displayed. On this page, you can view the list of messages.

4. To view the details of a message, select the message.

   By default, the messages are sorted by time, in ascending order. You can sort the messages by the any of the columns, such as Message Type, by clicking the column name.

5. To view messages that are related by time or ECID, click **View Related Messages** and select **by Time** or **by ECID (Execution Context ID).**

   The Related Messages page is displayed.

18.1.2 Searching Log Files

You can search for diagnostic messages using the Log Messages page. By default, this page shows a summary of the logged issues for the last hour.

You can modify the search criteria to identify messages of relevance. You can view the search results in different modes, allowing ease of navigation through large amounts of data.

The following sections describe how to search log files:

- Searching Log Files: Basic Searches
- Searching Log Files: Advanced Searches

18.1.2.1 Searching Log Files: Basic Searches

You can search for all of the messages for all of the entities in a domain, an Oracle WebLogic Server, a component, or an application.

For example, to search for messages for a domain:

1. From the **Targets** menu, select **Middleware**, then select a domain. From the **Farm** menu, select **Logs**, then select **View Log Messages**.

   or

   From the **Enterprise** menu, select **Monitoring**, then select **Logs**.

The Log Messages page displays a Search section and a table that shows a summary of the messages for the last hour.

2. In the Date Range section, you can select either:
   - **Most Recent:** If you select this option, select a time, such as 3 hours. The default is 1 hour.
   - **Time Interval:** If you select this option, select the calendar icon for **Start Date.** Select a date and time. Then, select the calendar icon for **End Date.** Select a date and time.
3. In the Message Types section, select one or more of the message types.
4. You can specify more search criteria, as described in Searching Log Files: Advanced Searches.
5. Click **Search.**
6. To help identify messages of relevance, in the table, for **Show,** select one of the following modes:
   - **Messages:** Shows the matching messages.

     To see the details of a particular message, click the message. The details are displayed below the table of messages.

     To view related messages, select a message, then click **View Related Messages** and select **by Time** or **by ECID (Execution Context ID).**
   - **Group by Message Type:** Summarizes the matching messages by grouping them based on message type at the target level.

     To see the messages, click the count in one of the message type columns. The Messages by Message Type page is displayed. To see the details of a particular message, click the message. The details are displayed below the table of messages.
   - **Group by Message ID:** Summarizes the matching messages by grouping them based on message ID, message type, and module IDs at the target level.

     To see the associated messages, click the count in the **Occurrences** column. The Messages by Message ID page is displayed. To see the details of a particular message, click the message. The details are displayed below the table of messages.

18.1.2.2 Searching Log Files: Advanced Searches

You can refine your search criteria using the following controls in the Log Messages page:

- **Message:** You can select an operator, such as **contains** and then enter a value to be matched.
- **Add Fields:** Click this to specify additional criteria, such as Host, which lets you narrow the search to particular hosts. Then click **Add.**

  For each field you add, select an operator, such as **contains** and then enter a value to be matched.
- **Selected Targets:** Expand this to see the targets that are participating in the search. To add targets, click **Add** and provide information in the dialog box. To remove targets, select the target and click **Remove.**

18.1.3 Downloading Log Files

You can download the log messages to a file. You can download either the matching messages from a search or the messages in a particular log file.

To download the matching messages from a search to a file:

1. From the **Enterprise** menu, select **Monitoring**, then select **Logs**.

   or

   From the **Targets** menu, select **Middleware**, then select a domain. From the **Farm** menu, select **Logs**, then select **View Log Messages**.

   The Log Messages page is displayed.

2. Search for particular types of messages as described in Searching Log Files: Basic Searches.

3. Select a file type by clicking **Export Messages to File** and select one of the following:

   - **As Oracle Diagnostic Log Text (.txt)**
   - **As Oracle Diagnostic Log Text (.xml)**
   - **As Comma-Separated List (.csv)**

   An Opening dialog box is displayed.

4. Select either **Open With** or **Save to Disk.** Click **OK.**

To export specific types of messages or messages with a particular Message ID to a file:

1. From the **Enterprise** menu, select **Monitoring**, then select **Logs**.

   or

   From the **Targets** menu, select **Middleware**, then select a domain. From the **Farm** menu, select **Logs**, then select **View Log Messages**.

   The Log Messages page is displayed.

2. Search for particular types of messages as described in Searching Log Files: Basic Searches.

3. For **Show,** select **Group by Message Type** or **Group by Message ID.**

4. To download the messages into a file, if you selected Group by Message Type, select the link in one of the columns that lists the number of messages, such as the Errors column. If you selected Group by Message ID, select one of the links in the Occurrences column.

   The Messages by Message Type page or Message by Message ID is displayed.

5. Select a file type by clicking the arrow near **Export Messages to File.**

   You can select one of the following:

   - **As Oracle Diagnostic Log Text (.txt)**
   - **As Oracle Diagnostic Log Text (.xml)**
   - **As Comma-Separated List (.csv)**

   An Opening dialog box is displayed.

6. Select either **Open With** or **Save to Disk.** Click **OK.**

To download the log files for a specific component:

1. From the **Enterprise** menu, select **Monitoring**, then select **Logs**.

   or

   From the **Targets** menu, select **Middleware**, click a domain. From the **Farm** menu, select **Logs**, then select **View Log Messages**.

   The Log Messages page is displayed.

2. Click **Target Log Files.**

   The Log Files page is displayed. On this page, you can see a list of log files related to the component or application.

3. Select a log file and click **Download.**

4. An Opening dialog box is displayed.

5. Select either **Open With** or **Save to Disk.** Click **OK.**

18.2 Locating Management Agent Log and Trace Files

The following sections provide information on the log and trace files for the Oracle Management Agent:

- About the Management Agent Log and Trace Files
- Locating the Management Agent Log and Trace Files

18.2.1 About the Management Agent Log and Trace Files

Oracle Management Agent log and trace files store important information that support personnel can later use to troubleshoot problems. The agent main log is located in `$EMSTATE/sysman/log`. The log is segmented by default to 11 segments, 5MB each. The segments are named gcagent.log and gcagent.log.# where # is a number in the range of 1-10. These settings are controlled by properties in emd.properties as explained in the following sections. The latest segment is always gcagent.log and the oldest is the gcagent.log.X where X is the highest number.

The Management Agent uses the following log files:

- Oracle Management Agent log file (`gcagent.log`)

  This log file contains trace, debug, information, error, or warning messages from the agent.

- Oracle Management Agent trace log file (`gcagent_sdk.trc`)

  This log file contains logging information about fetchlets and receivelets.

- Oracle Management Agent errors log file (`gcagent_errors.log`)

  This error log file contains information about errors. The errors in this file are duplicate of the errors in `gcagent.log`.

- Oracle Management Agent log file (`gcagent_mdu.log`)

  This log tracks the metadata updates to the agent.

- Enterprise Manager Control log file (`emctl.log`)

  The information is saved to `emctl.log` file, when you run the Enterprise Manager Control commands. For more information about `emctl.log` file, see chapter *Starting and Stopping Enterprise Manager Components*.

Note: All the agent logs mentioned above (existing in `$EMSTATE/sysman/log`) are transient. Agent logs are segmented and have a limited overall size and hence need not be deleted or managed.

18.2.1.1 Structure of Agent Log Files

The log contain individual log messages with the following format:

```
YYYY-MM-DD HH:MM:SS,### [<tid>:<thread code or code:name>] <level> -<the message>
```

Where:

- YYYY-MM-DD HH:MM:SS,### is a timestamp (in 24 hours format and ### is the fraction in msec).
- <tid> is the thread id (as a decimal number)
- <thread name or code> is the thread full name or an abbreviated hexadecimal code (see the following example).
- <level> is the logging level that can be one of (in ascending order of importance): DEBUG, INFO, WARN, ERROR, FATAL.
- <the message> is the free text message that is being logged. The message can contain new lines and spawn multiple lines.

For example:

```
2011-06-07 15:00:00,016 [1:3305B9:main] DEBUG - ADR_BASE='/ade/example_
user/oracle/example/agentStateDir' 2011-06-07 15:00:01,883 [1:3305B9] INFO - Agent
is starting up
```

18.2.2 Locating the Management Agent Log and Trace Files

The log and trace files for the Agent are written in the Agent runtime directory. You can find the runtime directory by using this command:

```
$ emctl getemhome
```

The log and trace files will be located at `<EMHOME>/sysman/log`.

18.2.3 Setting Oracle Management Agent Log Levels

Every log message is logged using a specific log level. The log levels are ordered in priority order: DEBUG, INFO, WARN, ERROR, and FATAL. The log setting determines the minimum level that will be included in the log. For example, if the log level is set to INFO (the default), only log messages of level INFO and above (INFO, WARN, ERROR and FATAL) are going to be included in the log.

Agent logging is based on log4j so the logging configuration is also log4j-based. The logging configuration properties are located in emd.properties and they all start with the prefix "Logger".

The "rootCategory" property controls the default log level. It is by default set to INFO. For example:

```
Logger.log4j.rootCategory=INFO,Rolling,Errors,Test
```

The agent uses an embedded HTTP server (jetty) to service client requests made on HTTP. There is a dedicated class which enables jetty logging as shown below:

```
Logger.log4j.category.oracle.sysman.gcagent.comm.agent.http.HTTPListener=INFO
```

This can be configured to DEBUG to try to troubleshoot certain communication problems, but the number of entries logged for DEBUG setting cannot be controlled. It is recommended that you configure the DEBUG option only under the direction of Oracle Support as in doing so may cause other valuable logging to be sacrificed.

18.2.3.1 Setting gcagent.log

The gcagent.log is configured using properties that starts with "Logger.log4j.appender.Rolling". The following is a sample gcagent.log:

```
Logger.log4j.appender.Rolling=org.apache.log4j.RollingFileAppender
Logger.log4j.appender.Rolling.File=/ade/user_
name/oracle/agentStateDir/sysman/log/gcagent.log
Logger.log4j.appender.Rolling.Append=true
Logger.log4j.appender.Rolling.MaxFileSize=5000000
Logger.log4j.appender.Rolling.MaxBackupIndex=10
Logger.log4j.appender.Rolling.layout=oracle.sysman.gcagent.util.logging.GCPattern
```

Where:

- "File" property controls the location and name of the main log. It is recommended that you do not change the setting for it.
- "Append" determines if the log file should be appended (true) or overwritten (false) on agent startup.
- " MaxFileSize" determines the size of each of the log segments.
- "MaxBackupIndex" determines the number of "backup" segments for the log (the total number of segments are this number plus one).

18.2.3.2 Setting gcagent_error.log

The gcagent_errors.log is a subset of the gcagent.log. The logging configuration for gcagent_errors.log is specified in `emd.properties`. Following are the settings for gcagent_errors.log:

```
Logger.log4j.appender.Errors=org.apache.log4j.RollingFileAppender
Logger.log4j.appender.Errors.File=/ade/user_
name/oracle/work/agentStateDir/sysman/log/gcagent_errors.log
Logger.log4j.appender.Errors.Append=true
Logger.log4j.appender.Errors.Threshold=ERROR
Logger.log4j.appender.Errors.layout=oracle.sysman.gcagent.util.logging.GCPattern
Logger.log4j.appender.Errors.MaxFileSize=50000000
Logger.log4j.appender.Errors.MaxBackupIndex=3
```

18.2.3.3 Setting the Log Level for Individual Classes and Packages

The logging level for individual class and/or packages can also be set. The following are examples that are currently configured by default:

```
# Set the class loaders to level INFO
Logger.log4j.category.oracle.sysman.gcagent.metadata.impl.ChainedClassLoader=INFO

Logger.log4j.category.oracle.sysman.gcagent.metadata.impl.ReverseDelegationClassLo
ader=INFO
Logger.log4j.category.oracle.sysman.gcagent.metadata.impl.PluginLibraryClassLoader
=INFO
Logger.log4j.category.oracle.sysman.gcagent.metadata.impl.PluginClassLoader=INFO
```

The above configuration changed the default level of logging for the four classes to be INFO. When the default level of logging is INFO it does not make any difference but if the default log level is set to DEBUG (when debugging the code) it will prevent those four classes from logging at DEBUG level (as they are normally too verbose).

The reverse is also true, for example if the following configuration is added (not set by default):

```
Logger.log4j.category.oracle.sysman.gcagent.metadata.impl.collection=DEBUG
```

It will cause all classes in the "oracle.sysman.gcagent.metadata.impl.collection" package to log at DEBUG level even if the default log level is INFO.

18.2.3.4 Setting gcagent_mdu.log

A set of entries are created in the gcagent_mdu.log file for each client command that modifies target instances, target instance collections, or blackouts. Entries are as follows:

```
2011-08-18 22:56:40,467 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] - SAVE
TARGET(S)
<Target IDENTIFIER="TARGET_GUID=6A3A159D0BB320C50B7926E0671A1A98"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_NAME="EM Management
Beacon" NAME="EM Management Beacon" TYPE="oracle_beacon"/>
<Target IDENTIFIER="TARGET_GUID=51F9BBC6F5B833058F4278B51E496000"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_NAME="mytestBeacon"
NAME="mytestBeacon" TYPE="oracle_beacon"><Property VALUE="***"
NAME="proxyHost"/><Property VALUE="***" NAME="proxyPort"/><Property VALUE="***"
NAME="dontProxyFor"/></Target>
<Target IDENTIFIER="TARGET_GUID=7C4336B536C9F241DBCAC4D1D082AD22"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_NAME="CSAcollector"
NAME="CSAcollector" TYPE="oracle_csa_collector"><Property VALUE="***"
NAME="recvFileDir"/></Target>
<Target IDENTIFIER="TARGET_GUID=207B57A3FE300C86F81FE7D409F5DD1C"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_NAME="Oemrep_Database"
NAME="Oemrep_Database" TYPE="oracle_database"><Property VALUE="***"
NAME="MachineName"/><Property VALUE="***" NAME="Port"/><Property VALUE="***"
NAME="SID"/><Property VALUE="***" NAME="OracleHome"/><Property ENCRYPTED="FALSE"
VALUE="***" NAME="UserName"/><Property ENCRYPTED="FALSE" VALUE="***"
NAME="Role"/><Property ENCRYPTED="FALSE" VALUE="***" NAME="password"/></Target>
<Target IDENTIFIER="TARGET_GUID=0C48C5AE0FAFB42ED91F897FF398FC84"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_NAME="Management Services
and Repository" NAME="Management Services and Repository" TYPE="oracle_
emrep"><Property VALUE="***" NAME="ConnectDescriptor"/><Property ENCRYPTED="FALSE"
VALUE="***" NAME="UserName"/><Property ENCRYPTED="FALSE" VALUE="***"
NAME="password"/><AssocTargetInstance ASSOC_TARGET_TYPE="oracle_oms" ASSOC_TARGET_
NAME="adc2190447.us.oracle.com:41034_Management_Service" ASSOCIATION_NAME="app_
composite_contains"/><AssocTargetInstance ASSOC_TARGET_TYPE="oracle_oms" ASSOC_
TARGET_NAME="adc2190447.us.oracle.com:41034_Management_Service" ASSOCIATION_
NAME="internal_contains"/><CompositeMembership><Member ASSOCIATION=""
NAME="adc2190447.us.oracle.com:41034_Management_Service_CONSOLE" TYPE="oracle_oms_
console"/><Member ASSOCIATION="" NAME="adc2190447.us.oracle.com:41034_Management_
Service_PBS" TYPE="oracle_oms_pbs"/><Member ASSOCIATION=""
NAME="adc2190447.us.oracle.com:41034_Management_Service" TYPE="oracle_
oms"/></CompositeMembership></Target>
<Target IDENTIFIER="TARGET_GUID=DF64B4A7C0F2EEBA7894EA3AD4CAF61E"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_
NAME="adc2190447.us.oracle.com:41034_Management_Service"
NAME="adc2190447.us.oracle.com:41034_Management_Service" TYPE="oracle_
oms"><Property VALUE="***" NAME="InstanceHome"/><Property VALUE="***"
NAME="OracleHome"/><AssocTargetInstance ASSOC_TARGET_TYPE="oracle_oms_console"
```

```
ASSOC_TARGET_NAME="adc2190447.us.oracle.com:41034_Management_Service_CONSOLE"
ASSOCIATION_NAME="app_composite_contains"/><AssocTargetInstance ASSOC_TARGET_
TYPE="oracle_oms_pbs" ASSOC_TARGET_NAME="adc2190447.us.oracle.com:41034_
Management_Service_PBS" ASSOCIATION_NAME="app_composite_
contains"/><AssocTargetInstance ASSOC_TARGET_TYPE="oracle_oms_console" ASSOC_
TARGET_NAME="adc2190447.us.oracle.com:41034_Management_Service_CONSOLE"
ASSOCIATION_NAME="internal_contains"/><AssocTargetInstance ASSOC_TARGET_
TYPE="oracle_oms_pbs" ASSOC_TARGET_NAME="adc2190447.us.oracle.com:41034_
Management_Service_PBS" ASSOCIATION_NAME="internal_
contains"/><CompositeMembership><MemberOf ASSOCIATION="" NAME="Management Services
and Repository" TYPE="oracle_emrep"/><Member ASSOCIATION=""
NAME="adc2190447.us.oracle.com:41034_Management_Service_CONSOLE" TYPE="oracle_oms_
console"/><Member ASSOCIATION="" NAME="adc2190447.us.oracle.com:41034_Management_
Service_PBS" TYPE="oracle_oms_pbs"/></CompositeMembership></Target>
<Target IDENTIFIER="TARGET_GUID=4D290260F13596502EFD8F3E22752404"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_
NAME="adc2190447.us.oracle.com:41034_Management_Service_CONSOLE"
NAME="adc2190447.us.oracle.com:41034_Management_Service_CONSOLE" TYPE="oracle_oms_
console"><Property VALUE="***" NAME="InstanceHome"/><Property VALUE="***"
NAME="OracleHome"/><CompositeMembership><MemberOf ASSOCIATION="" NAME="Management
Services and Repository" TYPE="oracle_emrep"/><MemberOf ASSOCIATION=""
NAME="adc2190447.us.oracle.com:41034_Management_Service" TYPE="oracle_
oms"/></CompositeMembership></Target>
<Target IDENTIFIER="TARGET_GUID=D0A23AE06A9E678221B075A216364541"
STATUS="MONITORED" TIMEZONE_REGION="" ON_HOST="" DISPLAY_
NAME="adc2190447.us.oracle.com:41034_Management_Service_PBS"
NAME="adc2190447.us.oracle.com:41034_Management_Service_PBS" TYPE="oracle_oms_
pbs"><Property VALUE="***" NAME="InstanceHome"/><Property VALUE="***"
NAME="OracleHome"/><CompositeMembership><MemberOf ASSOCIATION="" NAME="Management
Services and Repository" TYPE="oracle_emrep"/><MemberOf ASSOCIATION=""
NAME="adc2190447.us.oracle.com:41034_Management_Service" TYPE="oracle_
oms"/></CompositeMembership></Target>
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
2011-08-18 22:57:10,084 [HTTP Listener-34 - /emd/main/ (DispatchRequests)] -
SUCCESS
```

For the batch of saved targets in the above example, the original request came in at 22:56:40 and the list of targets saved are found in the line(s) following the SAVE TARGET(S) message. In this case, there were 8 targets. The result of saving the targets is available in the next 8 lines (for the same thread) and in this case all were saved successfully by 22:57:10.

The pattern is the same for saved collection items (or collections) and blackouts.

The logging configuration for the gcagent_mdu log is specified in `emd.properties` but you must not be modify this log. For example, these entries are logged at INFO level, which means that if you decided to save space and change this to WARN only

by editing the mdu log entries in the `emd.properties` file, you will have effectively disabled this log.

Following are the settings for gcagent_mdu log:

```
Logger.log4j.appender.Mdu=org.apache.log4j.RollingFileAppender
Logger.log4j.appender.Mdu.File=/ade/user_
name/oracle/work/agentStateDir/sysman/log/gcagent_mdu.log
Logger.log4j.appender.Mdu.Append=true
Logger.log4j.appender.Mdu.Threshold=INFO
Logger.log4j.appender.Mdu.layout=org.apache.log4j.PatternLayout
Logger.log4j.appender.Mdu.layout.ConversionPattern=%d [%t] - %m%n
Logger.log4j.appender.Mdu.MaxFileSize=50000000
Logger.log4j.appender.Mdu.MaxBackupIndex=3
```

18.2.3.5 Setting the TRACE Level

The following `_enableTrace` property when set to "true" will enable the TRACE logging level that shows as DEBUG messages.

```
Logger._enableTrace=true
```

The default log level for the agent log must be set to DEBUG for the tracing level to work.

18.3 Locating and Configuring Oracle Management Service Log and Trace Files

The following sections describe how to locate and configure the OMS log files:

- Locating Oracle Management Service Log and Trace Files
- Controlling the Size and Number of Oracle Management Service Log and Trace Files
- Controlling the Contents of the Oracle Management Service Trace File
- Controlling the Oracle WebLogic Server and Oracle HTTP Server Log Files

18.3.1 About the Oracle Management Service Log and Trace Files

OMS log and trace files store important information that Oracle Support can later use to troubleshoot problems. OMS uses the following six types of log files:

- Oracle Management Service log file (`emoms.log`)

  The Management Service saves information to the log file when it performs an action (such a starting or stopping) or when it generates an error. This is a log file for console application.

- Oracle Management Service trace file (`emoms.trc`)

  OMS trace file provides an advanced method of troubleshooting that can provide support personnel with even more information about what actions the OMS was performing when a particular problem occurred. This is a trace file for Console application.

- Oracle Management Service log file (`emoms_pbs.log`)

  The Management Service saves information to this log file for background modules such as the loader, job system, event system, notification system, and so on. This file contains messages logged at ERROR or WARN levels.

- Oracle Management Service trace file (`emoms_pbs.trc`)

  This trace file provides additional logging for the background modules such as the loader, job system, event system, notification system, and so on when DEBUG or INFO level logging is enabled for these modules. This file can provide Support personnel with even more information about actions these modules were performing when a particular problem occurred.

- Enterprise Manager Control log file (`emctl.log`)

  The information is saved to `emctl.log` file, when you run the Enterprise Manager Control commands. For more information about emctl.log file, see chapter *Starting and Stopping Enterprise Manager Components*.

18.3.2 Locating Oracle Management Service Log and Trace Files

OMS log and trace files are stored in the following location:

```
<EM_INSTANCE_BASE>/em/<OMS_NAME>/sysman/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.

For example, if the Oracle Middleware Home is `/u01/app/Oracle/Middleware`, then the instance base directory is `/u01/app/Oracle/gc_inst`, and the log and trace files are available in `/u01/app/Oracle/gc_inst/em/EMGC_OMS1/sysman/log/` directory path.

18.3.3 Controlling the Size and Number of Oracle Management Service Log and Trace Files

OMS log and trace files increases in size over time as information is written to the files. However, the files are designed to reach a maximum size. When the files reach the predefined maximum size, the OMS renames (or rolls) the logging information to a new file name and starts a new log or trace file. This process keeps the log and trace files from growing too large.

As a result, you will often see multiple log and trace files in the OMS log directory. The following example shows one archived log file and the current log file in the `/u01/app/Oracle/gc_inst/em/EMGC_OMS1/sysman/log/` directory:

```
emoms.log
emoms.log.1
```

To control the maximum size of the OMS log and OMS trace files, as well as the number of rollover files, run the following command, and specify details as described in Table 18–2:

```
emctl set property -name <property> -value <property value> -module logging
```

Note: In Oracle Enterprise Manager Cloud Control 12*c*, you do not have to restart OMS for the changes to take effect.

Table 18–2 Oracle Management Service Log File Properties in the emomslogging.properties File

Property	Purpose	Example
log4j.appender.emlogAppender.MaxFileSize	When OMS log file reaches this size, then OMS copies the logging data to a new rollover file and creates a new `emoms.log` log file. The size of the log is specified in units of bytes. This property is also applicable for emoms_pbs.log.	log4j.appender.emlogAppender.MaxFileSize=20000000
log4j.appender.emlogAppender.MaxBackupIndex	This optional property indicates how many times OMS will rollover the log file to a new file name before deleting logging data. This property is also applicable for emoms_pbs.log. **Note:** Because the log file does not contain as much data as the trace file, it is usually not necessary to create more than one rollover file.	log4j.appender.emlogAppender.MaxBackupIndex=1
log4j.appender.emtrcAppender.MaxFileSize	When the OMS trace file reaches this size, then OMS copies the logging data to a new rollover file and creates a new `emoms.trc` log file. This property is also applicable for emoms_pbs.trc.	log4j.appender.emtrcAppender.MaxFileSize=5000000
log4j.appender.emtrcAppender.MaxBackupIndex	This property indicates how many times the OMS will rollover the trace file to a new file name before deleting tracing data. This property is also applicable for emoms_pbs.trc.	log4j.appender.emtrcAppender.MaxBackupIndex=10

18.3.4 Controlling the Contents of the Oracle Management Service Trace File

By default, the OMS will save all critical and warning messages to the `emoms.trc` file. However, you can adjust the amount of logging information that the OMS generates.

To change the amount of logging information generated by the OMS, run the following command:

```
emctl set property -name "log4j.rootCategory" -value "<LEVEL>, emlogAppender,
emtrcAppender" -module logging
```

Note: If you change the `root` logging level for the `emoms.trc` file, then a lot of messages are written to the trace file filling up the space quickly, and potentially slowing down the system. Run the following command to enable debug selectively for specific modules that need to be assessed:

```
emctl set property -name <logging module> -value DEBUG -module
logging
```

Where, `<logging module>` represents the logging module from a specific subsystem.

For example, `oracle.sysman.emdrep.dbjava.loader.`

The logging level can be changed for specific modules by running the following command:

```
emctl set property -name "<CATEGORY>" -value "<LEVEL>" -module
logging
```

where LEVEL can be DEBUG, INFO, WARN, or ERROR, and CATEGORY is specific to the module for which level has to be changed. To change the logging module, contact Oracle Support.

Note: The location of emoms.trc, emoms.log, emoms_pbs.trc, and emoms_pbs.log files can be changed to a different location from the default location. However, it is not advisable to do so.

18.3.5 Controlling the Oracle WebLogic Server and Oracle HTTP Server Log Files

Oracle Management Service is a Java EE application deployed on an Oracle WebLogic Server. Different components of the Oracle WebLogic Server generate their own log files. These files contain important information that can be used later by support personnel to troubleshoot problems.

Table 18–3 lists the location of the log files for some components.

Table 18–3 Component Log File Location

Component	Location

Oracle HTTP Server (OHS)	`<EM_INSTANCE_BASE>/<webtier_instance_name>/diagnostics/logs/OHS/<ohs_name>`
	For example,
	`/u01/app/Oracle/gc_inst/WebTierIH1/diagnostics/logs/OHS/ohs1`

OPMN	`<EM_INSTANCE_BASE>/<webtier_instance_name>/diagnostics/logs/OPMN/<opmn_name>`
	For example,
	`/u01/app/Oracle/gc_inst/WebTierIH1/diagnostics/logs/OPMN/opmn`

Table 18–3 (Cont.) Component Log File Location

Component	Location

Oracle WebLogic	The log data from WebLogic will be at:
	`<EM_INSTANCE_BASE>/user_projects/domains/<domain_name>/servers/<SERVER_NAME>/logs/<SERVER_NAME>.log`
	This log can be restricted, rotated by size, time, and other conditions from the WebLogic Console. The default settings are:
	■ In production mode, they are rotated at a default of 5MB.
	■ The log level is WARNING.
	■ The number files are restricted to 10.
	For example,
	`/u01/app/Oracle/gc_inst/user_projects/domains/GCDomain/servers/EMGC_OMS1/logs/EMGC_OMS1.log`
	The messages written to sysout and syserr will be available in the .out files. They cannot be rotated by size or time. They are rotated only when the server starts. They are located at:
	`<EM_INSTANCE_BASE>/user_projects/domains/<domain_name>/servers/<SERVER_NAME>/logs/<SERVER_NAME>.out`
	For example,
	`/u01/app/Oracle/gc_inst/user_projects/domains/GCDomain/servers/EMGC_OMS1/logs/EMGC_OMS1.out`
	The node manager logs are at `<INST_HOME>/NodeManager/emnodemanager` and the admin server logs are at `<INST_HOME>/user_projects/domains/GCDomain/servers/EMGC_ADMINSERVER/logs`.

By default, the Enterprise Manager Cloud Control configures Oracle HTTP Server logs to roll over periodically to a new file, so that each file does not grow too large in size. You must also ensure that you delete the old rollover files periodically to free up the disk space. You can use an operating system scheduler, like cron on UNIX, to periodically delete the rollover files.

Note: Following are log files that you will need to maintain and manually purge:

- `<gc_inst>/user_projects/domains/<domain_name>/servers/EMGC_ADMINSERVER/logs/<domain_name>.log*`
- All files under `<gc_inst>/WebTierIH1/diagnostics/logs/OHS/ohs1/`. For example:

```
em_upload_http_access_log.*
access_log.*
em_upload_https_access_log.*
ohs1-*.log
console~OHS~1.log*
mod_wl_ohs.log*
```

- Log files under the admin server and emgc_oms server:

```
<gc_inst>\users_projects\domains\<domain_name>\servers\EMGC_
ADMINSERVER\logs\*.out*
<gc_inst>\users_projects\domains\<domain_name>\servers\EMGC_
OMS?\logs\*.out*
```

For instructions on controlling the size and rotation of these log files, refer to chapter "Managing Log Files and Diagnostic Data" in *Oracle Fusion Middleware Administrator's Guide*.

For information about configuring Enterprise Manager to view Fusion Applications PL/SQL and C diagnostic log files, see chapter "Managing Oracle Fusion Applications Log Files and Diagnostic Tests" in *Oracle Fusion Applications Administrator's Guide*.

19

Maintaining and Troubleshooting the Management Repository

This chapter describes maintenance and troubleshooting techniques for maintaining a well-performing Management Repository.

Specifically, this chapter contains the following sections:

- Management Repository Deployment Guidelines
- Management Repository Data Retention Policies
- Changing the SYSMAN Password
- Dropping and Recreating the Management Repository
- Troubleshooting Management Repository Creation Errors
- Cross Platform Enterprise Manager Repository Migration

19.1 Management Repository Deployment Guidelines

To be sure that your management data is secure, reliable, and always available, consider the following settings and configuration guidelines when you are deploying the Management Repository:

- Install a RAID-capable Logical Volume Manager (LVM) or hardware RAID on the system where the Management Repository resides. At a minimum the operating system must support disk mirroring and stripping. Configure all the Management Repository data files with some redundant configuration.
- Use Real Application Clusters to provide the highest levels of availability for the Management Repository.
- If you use Enterprise Manager to alert administrators of errors or availability issues in a production environment, be sure that the Cloud Control components are configured with the same level of availability. At a minimum, consider using Oracle Data Guard to mirror the Management Repository database. Configure Data Guard for zero data loss. Choose between Maximum Availability or Maximum Protection based on your environment and needs.

  See Also: *Oracle Database High Availability Architecture and Best Practices*

  Oracle Data Guard Concepts and Administration

- Oracle strongly recommends that archive logging be turned on and that a comprehensive backup strategy be in place prior to an Enterprise Manager

implementation going live in a production environment. The backup strategy should include both incremental and full backups as required.

> **See Also:** *Oracle Enterprise Manager Cloud Control Installation and Basic Configuration* for information about the database initialization parameters required for the Management Repository

19.2 Management Repository Data Retention Policies

When the various components of Enterprise Manager are configured and running efficiently, the Oracle Management Service gathers large amounts of raw data from the Management Agents running on your managed hosts and loads that data into the Management Repository. This data is the raw information that is later aggregated, organized, and presented to you in the Cloud Control console.

After the Oracle Management Service loads information into the Management Repository, Enterprise Manager aggregates and purges the data over time.

The following sections describe:

- The default aggregation and purging policies used to maintain data in the Management Repository.
- How you can modify the length of time the data is retained before it is aggregated and then purged from the Management Repository.

19.2.1 Management Repository Default Aggregation and Purging Policies

Enterprise Manager aggregates collected metric data by hour and by day to enhance query performance and help minimize the size of the Management Repository. Before the data is aggregated, each data point is stored in a raw metric data table. Once a day, the previous day's raw metric data is rolled up, or aggregated, into a one-hour and a one-day table. These hourly and daily records will have hourly and daily metric data averages, minimums, maximums and standard deviations respectively.

After Enterprise Manager aggregates the data, the data is then considered eligible for purging. A certain period of time must pass for data to actually be purged. This period of time is called the retention time.

The raw data, with the highest insert volume, has the shortest default retention time, which is set to 7 days. As a result, 7 days after it is aggregated into a one-hour record, a raw data point is eligible for purging.

> **Note:** This data retention policy varies for JVMD and ADP data.

Hourly aggregate metric data records are purged after 31 days. The highest level of aggregation, one day, is kept for 12 months (roughly 365 days).

The default data retention policies are summarized in Table 19–1.

Table 19–1 Default Repository Purging Policies

Aggregate Level	Retention Time
Raw metric data	7 days
Hourly aggregated metric data	31 days
Daily aggregated metric data	12 months (~365 days)

If you have configured and enabled Application Performance Management, Enterprise Manager also gathers, saves, aggregates, and purges response time data. The response time data is purged using policies similar to those used for metric data. The Application Performance Management purging policies are shown in Table 19–2.

Table 19–2 Default Repository Purging Policies for Application Performance Management Data

Aggregate Level	Retention Time
Raw response time data	24 hours
One-hour aggregated response time data	7 days
One-hour distribution response time data	24 hours
One-day aggregated response time data	31 days
One-day distribution aggregated response time data	31 days

19.2.2 Management Repository Default Aggregation and Purging Policies for Other Management Data

Besides the metric data and Application Performance Monitoring data, other types of Enterprise Manager data accumulates over time in the Management Repository.

For example, the last availability record for a target will also remain in the Management Repository indefinitely, so the last known state of a target is preserved.

19.2.3 Modifying the Default Aggregation and Purging Policies

The Enterprise Manager default aggregation and purging policies were designed to provide the most available data for analysis while still providing the best performance and least disk-space requirements for the Management Repository. As a result, you should not modify these policies to improve performance or increase your available disk space.

However, if you plan to extract or review the raw or aggregated data using data analysis tools other than Enterprise Manager, you may want to increase the amount of raw or aggregated data available in the Management Repository. You can accomplish this by increasing the retention times for the raw or aggregated data.

A PL/SQL API has been provided to modify the default retention time for the core metric data tables in the Enterprise Manager repository. Table 19–3 shows the default number of partitions retained for each of the three tables and the size of the partitions for each table. The API will allow you to change the number of partitions retained only.

Table 19–3 Core EM Metric Data Tables and Default Data Retention in the Management Repository

Table Name	Partitions Retained	Partition Size
EM_METRIC _VALUES	7	DAY
EM_METRIC_VALUES_HOURLY	32	DAY
EM_METRIC_VALUES_DAILY	12	MONTH

The PL/SQL API is:

```
gc_interval_partition_mgr.set_retention(<repository schema name>,
                                        <table name>,
                                        <number of partitions to retain>);
```

An example of using the PL/SQL API to change the number of partitions retained in EM_METRIC_VALUES (raw data) from the default of 7 to 10 follows:

```
BEGIN
  gc_interval_partition_mgr.set_retention('SYSMAN', 'EM_METRIC_VALUES', 10);
END;
/
```

19.2.4 Modifying Data Retention Policies When Targets Are Deleted

By default, when you delete a target from the Grid Control console, Enterprise Manager automatically deletes all target data from the Management Repository.

However, deleting raw and aggregated metric data for database and other data-rich targets is a resource consuming operation. Targets can have hundreds of thousands of rows of data and the act of deleting this data can degrade performance of Enterprise Manager for the duration of the deletion, especially when several targets are deleted at once.

To avoid this resource-consuming operation, you can prevent Enterprise Manager from performing this task each time you delete a target. When you prevent Enterprise Manager from performing this task, the metric data for deleted targets is not purged as part of target deletion task; instead, it is purged as part of the regular purge mechanism, which is more efficient.

In addition, Oracle strongly recommends that you do not add new targets with the same name and type as the deleted targets within 24 hours of target deletion. Adding a new target with the same name and type will result in the Grid Control console showing data belonging to the deleted target for the first 24 hours.

To disable raw metric data deletion:

1. Use SQL*Plus to connect to the Management Repository as the Management Repository user.

   The default Management Repository user is SYSMAN. For example:

   ```
   SQL> connect sysman/sysman_password;
   ```

2. To disable metric deletion, run the following SQL command.

   ```
   SQL> EXEC MGMT_ADMIN.DISABLE_METRIC_DELETION();
   SQL> COMMIT;
   ```

To enable metric deletion at a later point, run the following SQL command:

1. Use SQL*Plus to connect to the Management Repository as the Management Repository user.

   The default Management Repository user is SYSMAN. For example:

   ```
   SQL> connect sysman/oldpassword;
   ```

2. To enable metric deletion, run the following SQL command.

   ```
   SQL> EXEC MGMT_ADMIN.ENABLE_METRIC_DELETION();
   SQL> COMMIT;
   ```

19.2.5 How to Modify the Retention Period of Job History

Enterprise Manager Cloud Control has a default purge policy which removes all finished job details which are older than 30 days. This section provides details for modifying this default purge policy.

The actual purging of completed job history is implemented via a DBMS_SCHEDULER job that runs once a day in the repository database. When the job runs, it looks for finished jobs that are 'n' number of days older than the current time (value of sysdate in the repository database) and deletes these jobs. The value of 'n' is, by default, set to 30 days.

The default purge policy cannot be modified via the Enterprise Manager console, but it can be changed using SQL*Plus.

To modify this purge policy, follow these steps:

1. Log in to the repository database as the SYSMAN user, via SQL*Plus.
2. Check the current values for the purge policies using the following command:

```
SQL> select * from mgmt_job_purge_policies;

POLICY_NAME                    TIME_FRAME
------------------------------ ----------
SYSPURGE_POLICY                        30
REFRESHFROMMETALINKPURGEPOLICY          7
FIXINVENTORYPURGEPOLICY                 7
OPATCHPATCHUPDATE_PAPURGEPOLICY         7
```

   The purge policy responsible for the job deletion is called SYSPURGE_POLICY. As seen above, the default value is set to 30 days.

3. To change the time period, you must drop and recreate the policy with a different time frame:

```
SQL> execute MGMT_JOBS.drop_purge_policy('SYSPURGE_POLICY');

PL/SQL procedure successfully completed.

SQL> execute MGMT_JOBS.register_purge_policy('SYSPURGE_
POLICY', 60, null);

PL/SQL procedure successfully completed.

SQL> COMMIT;

Commit complete.

SQL> select * from mgmt_job_purge_policies;

POLICY_NAME                    TIME_FRAME
------------------------------ ----------
SYSPURGE_POLICY                        60
....
```

The above commands increase the retention period to 60 days. The time frame can also be reduced below 30 days, depending on the requirement.

You can check when the purge job will be executed next. The actual time that the purge runs is set to 5 AM repository time and can be verified using these steps:

1. Login to the Repository database using the SYSMAN account.
2. Execute the following command:

```
SQL> select job_name,
            to_char(last_start_date, 'DD-MON-YY HH24:MI:SS') last_run,
            to_char(next_run_date,   'DD-MON-YY HH24:MI:SS') next_run
from all_scheduler_jobs
where job_name ='EM_JOB_PURGE_POLICIES';

JOB_NAME              LAST_RUN           NEXT_RUN
--------------------- ------------------ ------------------
EM_JOB_PURGE_POLICIES                    07-SEP-11 05:00:00
```

The schedule can also be verified from the Enterprise Manager console by following these steps:

a. From the **Setup** menu, select **Management Service**, the select **Repository**.

b. Click the **Repository Operations** tab.

c. Find the Next Scheduled Run and Last Scheduled Run information for Job Purge in the list.

Please note that the time of the next scheduled execution of the Job Purge does not represent the cutoff time for the retention period; the cutoff time is determined by the purge policy at the time the Job Purge runs.

19.2.6 DBMS_SCHEDULER Troubleshooting

Enterprise Manager uses the database scheduler (dbms_scheduler) to run various processes in the repository. When the dbms_scheduler is stopped or has insufficient resources to operate, the Enterprise Manager processes do not run or are delayed. The following is a list of common causes that may prohibit the dbms_scheduler from running normally.

Job Queue Processes

The dbms_scheduler uses a separate job-queue process for each job it runs. The maximum number of these processes is controlled by the database parameter, *job_queue_processes*. If all processes are in use, no new jobs will be started.

The following query returns the number of currently running jobs.

```
SELECT count(*)
FROM dba_scheduler_running_jobs;
```

If the count is close to the setting of *job_queue_processes*, it could mean that Enterprise Manager dbms_scheduler jobs cannot be started (on time). Determine if any of the running dbms_scheduler jobs are stuck and consider increasing the setting for job_queue_processes.

Job Slave Processes

The dbms_scheduler also depends on the setting of the dbms_scheduler property MAX_JOB_SLAVE_PROCESSES. If the number of running dbms_scheduler jobs exceeds this setting, no new jobs will be started. This attribute can be checked using this query.

```
SELECT value
FROM dba_scheduler_global_attribute
WHERE attribute_name='MAX_JOB_SLAVE_PROCESSES';
```

If the count equals the number of running dbms_scheduler jobs, then determine if any of the running dbms_scheduler jobs are stuck and consult the dbms_scheduler documentation about how to adjust this attribute.

DBMS_SCHEDULER Program Disabled

The dbms_scheduler has an attribute that can be set to disable this feature in the database. When set, the Enterprise Manager dbms_scheduler jobs will not run. To check if this attribute has been set (inadvertently), run this query.

```
SELECT *
FROM dba_scheduler_global_attribute
WHERE attribute_name = 'SCHEDULER_DISABLED';
```

When a row is returned, the dbms_scheduler is disabled. Execute `dbms_scheduler.set_scheduler_attribute('SCHEDULER_DISABLED', 'FALSE');`

Consult the dbms_scheduler documentation about how to remove this attribute.

Too Many Database Sessions

Each dbms_scheduler job requires two database sessions. When no more sessions are available, Enterprise Manager dbms_scheduler jobs will not run. The following two queries give the maximum number of allowed sessions and the current number of active sessions:

```
SELECT value
FROM v$parameter
WHERE name='sessions';

SELECT count(*)
FROM v$session;
```

When the current number of sessions approaches the maximum, then you should determine if any of the sessions are stuck and consult the Oracle Database documentation about how to increase the maximum number of sessions.

Also the high water mark of the number of sessions may indicate that this issue has played a role in the past:

```
select *
from v$resource_limit
where resource_name = 'sessions' ;
```

If the MAX_UTILIZATION column indicates a value that is close the maximum number of sessions, it could explain why some of the Enterprise Manager dbms_ scheduler jobs may not have run (on time) in the past.

Insufficient Memory

The database may not be able to spawn a new job queue process when there is insufficient memory available. The following message in the database alert file, *Unable to spawn jobq slave processes*, in combination with, *(free memory = 0.00M)*, would be indicative of this problem. Please consult the Oracle Database documentation about how to diagnose this memory problem further.

19.3 Changing the SYSMAN Password

The SYSMAN account is the default super user account used to set up and administer Enterprise Manager. It is also the database account that owns the objects stored in the

Oracle Management Repository. From this account, you can set up additional administrator accounts and set up Enterprise Manager for use in your organization.

The SYSMAN account is created automatically in the Management Repository database during the Enterprise Manager installation. You also provide a password for the SYSMAN account during the installation.

If you later need to change the SYSMAN database account password, use the following procedure:

1. Shut down all the Oracle Management Service instances that are associated with the Management Repository.
2. Stop the Management Agent that is monitoring the target OMS and Repository.

   Failure to do this will result in the Management Agent attempting to connect to the target with a wrong password once it is changed with SQL*Plus. This may also result in the SYSMAN account being locked which can subsequently prevent logins to the Cloud Control console to change the password of the target OMS and Repository.
3. Change the password of the SYSMAN database account using the following SQL*Plus commands:

   ```
   SQL>connect sysman/oldpassword;

   SQL>alter user sysman identified by newpassword;
   ```
4. To change the password of the SYSMAN user, enter the following command. Oracle strongly recommends you use this emctl command to change the password which updates both the *sysman* and the *sysman_mds* passwords. If you use other supported methods to change the sysman password but do not change the sysman_mds password, you may experience issues. This is the only command you need to run to change the repository user password:

   ```
   emctl config oms -change_repos_pwd [-old_pwd <old_pwd>]
   [-new_pwd <new_pwd>] [-use_sys_pwd [-sys_pwd <sys_pwd>]]
   ```

   You must run this command on each Management Service in your environment.

Parameter	Description
`-old_pwd`	This is the current SYSMAN password.
`-new_pwd`	This is the new password.
`-use_sys_pwd`	This parameter is optional and is used to connect to the database as a SYS user.
`-sys_pwd`	This is the password for the SYS user.

5. In the Cloud Control console, click the **Targets** tab, then click **All Targets** on the sub tab.
6. Select the **Management Services and Repository** target, then click **Configure**. Enterprise Manager displays the Monitoring Configurations page.
7. Enter the new password in the Repository password field, then click **OK**.

In Release 12c, there are two more schemas (SYSMAN_OPSS, SYSMAN_APM) that use the same password. The *emctl config oms -change_repos_pwd* command updates all these schemas with the new password.

There are two modes in which this emctl command can be launched:

- When the current SYSMAN password is known,. In this case the command is launched with both old and new passwords. This command updates the passwords for all the schemas to the new password in the repository and also updates the datasource and all relevant locations where these passwords are stored.
- When the current SYSMAN password is not known or if the SYSMAN acct is locked. In these cases you need to provide the SYS credentials and launch the emctl command with the *-use_sys_pwd* option. This command uses the SYS credentials to update the passwords for all the schemas to the new password in the repository and also updates the datasource and all relevant locations where these passwords are stored.

There are no other commands required to update the schema passwords.

19.4 Dropping and Recreating the Management Repository

This section provides information about dropping the Management Repository from your existing database and recreating the Management Repository after you install Enterprise Manager.

19.4.1 Dropping the Management Repository

To recreate the Management Repository, you first remove the Enterprise Manager schema from your Management Repository database. You accomplish this task using the `-action drop` argument to the `RepManager` script, which is described in the following procedure.

To remove the Management Repository from your database:

1. Locate the `RepManager` script in the following directory of the Middleware Home where you have installed and deployed the Management Service:

   ```
   ORACLE_HOME/sysman/admin/emdrep/bin
   ```

2. At the command prompt, enter the following command:

   ```
   $PROMPT> RepManager repository_host repository_port repository_SID
   -sys_password password_for_sys_account -action dropall
   ```

   In this syntax example:

   - *`repository_host`* is the machine name where the Management Repository database is located
   - *`repository_port`* is the Management Repository database listener port address, usually 1521
   - *`repository_SID`* is the Management Repository database system identifier
   - *`password_for_sys_account`* is the password of the SYS user for the database. For example, `change_on_install`
   - `-action drop` indicates that you want to drop the Management Repository

Alternatively, you can use a connect descriptor to identify the database on the RepManager command line. The connect descriptor identifies the host, port, and name of the database using a standard Oracle database syntax.

For example, you can use the connect descriptor as follows to create the Management Repository:

```
$PROMPT> ./RepManager -connect "(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)
(HOST=host1)(PORT=1521)) (CONNECT_DATE=(SERVICE_NAME=servicename)))"
-sys_password efkl34lmn -action dropall
```

> **See Also:** "Establishing a Connection and Testing the Network" in the *Oracle Database Net Services Administrator's Guide* for more information about connecting to a database using connect descriptors.

19.4.2 Recreating the Management Repository

The preferred method for creating the Management Repository is to create the Management Repository during the Enterprise Manager installation procedure, which is performed using Oracle Universal Installer.

> **See Also:** *Oracle Enterprise Manager Cloud Control Installation and Basic Configuration* for information about installing Enterprise Manager.

However, if you need to recreate the Management Repository in an existing database, you can use the `RepManager` script, which is installed when you install the Management Service. Refer to the following sections for more information:

- Using the RepManager Script to Create the Management Repository
- Using a Connect Descriptor to Identify the Management Repository Database

19.4.2.1 Using the RepManager Script to Create the Management Repository

To create a Management Repository in an existing database:

1. Review the hardware and software requirements for the Management Repository as described in *Oracle Enterprise Manager Cloud Control Installation and Basic Configuration*. and review the section "Management Repository Deployment Guidelines" on page 19-1.
2. Locate the `RepManager` script in the following directory of the Oracle Management Service home directory:

   ```
   ORACLE_HOME/sysman/admin/emdrep/bin
   ```

3. At the command prompt, enter the following command:

   ```
   $PROMPT> ./RepManager repository_host repository_port repository_SID
   -sys_password password_for_sys_account -action create
   ```

In this syntax example:

- *`repository_host`* is the machine name where the Management Repository database is located
- *`repository_port`* is the Management Repository database listener port address, usually 1521 or 1526
- *`repository_SID`* is the Management Repository database system identifier
- *`password_for_sys_account`* is the password of the SYS user for the database. For example, `change_on_install`

Enterprise Manager creates the Management Repository in the database you specified in the command line.

19.4.2.2 Using a Connect Descriptor to Identify the Management Repository Database

Alternatively, you can use a connect descriptor to identify the database on the `RepManager` command line. The connect descriptor identifies the host, port, and name of the database using a standard Oracle database syntax.

For example, you can use the connect descriptor as follows to create the Management Repository:

```
$PROMPT> ./RepManager -connect "(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)
(HOST=host1)(PORT=1521)) (CONNECT_DATA=(SERVICE_NAME=servicename)))"
-sys_password efkl34lmn -action create
```

> **See Also:** "Establishing a Connection and Testing the Network" in the *Oracle Database Net Services Administrator's Guide* for more information about connecting to a database using a connect descriptor

The ability to use a connect string allows you to provide an address list as part of the connection string. The following example shows how you can provide an address list consisting of two listeners as part of the `RepManager` command line. If a listener on one host becomes unavailable, the second listener can still accept incoming requests:

```
$PROMPT> ./RepManager -connect "(DESCRIPTION=
(ADDRESS_LIST=
(ADDRESS=(PROTOCOL=TCP)(HOST=host1)(PORT=1521)
(ADDRESS=(PROTOCOL=TCP)(HOST=host2)(PORT=1521)
(CONNECT_DATE=(SERVICE_NAME=servicename)))"
-sys_password efkl34lmn -action create
```

> **See Also:** *Oracle Database High Availability Architecture and Best Practices*
>
> *Oracle Enterprise Manager Cloud Control Installation and Basic Configuration*

19.5 Troubleshooting Management Repository Creation Errors

Oracle Universal Installer creates the Management Repository using a configuration step at the end of the installation process. If the repository configuration tool fails, note the exact error messages displayed in the configuration tools window, wait until the other configuration tools have finished, exit from Universal Installer, and then use the following sections to troubleshoot the problem.

19.5.1 Package Body Does Not Exist Error While Creating the Management Repository

If the creation of your Management Repository is interrupted, you may receive the following when you attempt to create or drop the Management Repository at a later time:

```
SQL> ERROR:
ORA-00604: error occurred at recursive SQL level 1
ORA-04068: existing state of packages has been discarded
ORA-04067: not executed, package body "SYSMAN.MGMT_USER" does not exist
ORA-06508: PL/SQL: could not find program unit being called
ORA-06512: at "SYSMAN.SETEMUSERCONTEXT", line 5
ORA-06512: at "SYSMAN.CLEAR_EMCONTEXT_ON_LOGOFF", line 4
ORA-06512: at line 4
```

To fix this problem, see "General Troubleshooting Techniques for Creating the Management Repository" on page 19-12.

19.5.2 Server Connection Hung Error While Creating the Management Repository

If you receive an error such as the following when you try to connect to the Management Repository database, you are likely using an unsupported version of the Oracle Database:

```
Server Connection Hung
```

To remedy the problem, upgrade your database to the supported version as described in *Oracle Enterprise Manager Cloud Control Installation and Basic Configuration*.

19.5.3 General Troubleshooting Techniques for Creating the Management Repository

If you encounter an error while creating the Management Repository, drop the repository by running the `-drop` argument to the `RepManager` script.

See Also: Section 19.4.1, "Dropping the Management Repository"

If the `RepManager` script drops the repository successfully, try creating the Management Repository again.

If you encounter errors while dropping the Management Repository, do the following:

1. Connect to the database as SYSDBA using SQL*Plus.
2. Check to see if the SYSMAN database user exists in the Management Repository database.

   For example, use the following command to see if the SYSMAN user exists:

   ```
   prompt> SELECT username FROM DBA_USERS WHERE username='SYSMAN';
   ```

3. If the SYSMAN user exists, drop the user by entering the following SQL*Plus command:

   ```
   prompt> DROP USER SYSMAN CASCADE;
   ```

4. Check to see if the following triggers exist:

   ```
   SYSMAN.EMD_USER_LOGOFF
   SYSMAN.EMD_USER_LOGON
   ```

   For example, use the following command to see if the EMD_USER_LOGOFF trigger exists in the database:

   ```
   prompt> SELECT trigger_name FROM ALL_TRIGGERS
           WHERE trigger_name='EMD_USER_LOGOFF';
   ```

5. If the triggers exist, drop them from the database using the following commands:

   ```
   prompt> DROP TRIGGER SYSMAN.EMD_USER_LOGOFF;
   prompt> DROP TRIGGER SYSMAN.EMD_USER_LOGON;
   ```

19.6 Cross Platform Enterprise Manager Repository Migration

There are user requirements for migrating an Enterprise Manager repository across servers - same and cross platforms.

The Enterprise Manager repository migration process is not exactly the same as database migration. In case of Enterprise Manager Repository migration you must take care of Enterprise Manager specific data, options, and pre-requisites for the repository move. You should make sure data integrity is maintained from both the Enterprise Manager and Oracle database perspective.

This brings up need for defining the process that can be followed by end users for successful and reliable migration of repository in minimum time and with maximum efficiency.

The overall strategy for migration depends on:

- The source and target database version
- The amount of data/size of repository
- Actual data to migrate [selective/full migration]

Cross platform transportable tablespace along with data pump (for metadata) is the fastest and best approach for moving large Enterprise Manager Cloud Control repository from one platform to another. Other option that can be considered for migration is to use Data Pump for both data and metadata moves but this would require more time than the cross platform transportable tablespace approach for the same amount of data. The advantage to using the data pump approach is that it provides granular control over options and the overall process, as in the case of selective data being migrated and not the whole of source data. If the source and target is not on version 11g then export/import is the only way to get the data migrated cross platform.

More details on cross platform transportable tablespace, data pump, and export/import options can be found at the *Oracle Technology Network* (OTN) or in the *Oracle Database Administrator's Guide*.

19.6.1 Common Prerequisites

The following lists the common prerequisites for a repository migration:

- Source and target database must use the same character set and should be at same version.
- Source and target database should meet all the pre-requisites mentioned for Enterprise Manager Repository software requirements mentioned in Enterprise Manager install guide.
- If source and target database are NOT on 11g - only Export/Import can be used for cross platform migration.
- If Source and target database are on 11g - either of three options Cross platform transportable tablespaces migration, Data Pump or Export/Import can be used for cross platform repository migration.
- You cannot transport a tablespace to a target database in which a tablespace with the same name already exists. However, you can rename either the tablespace to be transported or the destination tablespace before the transport operation.
- To plug a transportable tablespace set into an Oracle Database on a different platform, both databases must have compatibility set to at least 10.0.
- Most of the platforms (but not all) are supported for cross-platform tablespace transport. You can query the V$TRANSPORTABLE_PLATFORM view to see the platforms that are supported, and to determine their platform IDs and their endian format (byte ordering).

- Source and Destination host should have Enterprise Manager Management Agent running and configured to the instance which is to be migrated.
- If target database has Enterprise Manager repository installed, it should be first dropped using RepManager before target database related steps are carried out.

19.6.2 Methodologies

The following sections discuss the methodologies of a repository migration.

19.6.2.1 Cross Platform Transportable Tablespaces

Oracle's transportable tablespace feature allows users to quickly move a user tablespace across Oracle databases. It is the most efficient way to move bulk data between databases. Prior to Oracle Database Release 11g, if you want to transport a tablespace, both source and target databases need to be on the same platform. Oracle Database Release 11g adds the cross platform support for transportable tablespaces. With the cross platform transportable tablespace, you can transport tablespaces across platforms.

Cross platform transportable tablespaces allows a database to be migrated from one platform to another (use with Data Pump or Import/Export).

19.6.2.1.1 Preparation for Transportable Tablespaces

Use these steps to prepare for transportable tablespaces:

1. Prepare set of user tablespaces and Check for containment violation.

   ```
   execute DBMS_TTS.TRANSPORT_SET_CHECK('MGMT_TABLESPACE,MGMT_
   ECM_DEPOT_TS', TRUE);

   select * FROM transport_set_violations;
   ```

2. Shutdown OMS instances and prepare for migration.

   Shutdown OMS, set job queue_processes to 0 and run:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   remove_dbms_jobs.sql
   ```

3. Make the tablespaces to be transported read only.

   ```
   alter tablespace MGMT_TABLESPACE read only;

   alter tablespace MGMT_ECM_DEPOT_TS read only;
   ```

19.6.2.1.2 Extract metadata

Extract Metadata for transportable tablespaces using Data Pump Utility:

1. Create the data pump directory.

   ```
   create directory data_pump_dir as
   '/scratch/user/EM102/ttsdata';
   ```

2. Extract the metadata using data pump (or export).

   ```
   expdp DUMPFILE=ttsem102.dmp TRANSPORT_TABLESPACES=MGMT_
   TABLESPACE,MGMT_ECM_DEPOT_TS TRANSPORT_FULL_CHECK=Y
   ```

3. Extract other objects (packages, procedures, functions, temporary tables. and so on -- Not contained in user tablespaces).

```
expdp SCHEMAS=SYSMAN CONTENT=METADATA_ONLY
EXCLUDE=INDEX,CONSTRAINT DUMPFILE=data_pump_dir:postexp.dmp
LOGFILE=data_pump_dir:postexp.log JOB_NAME=expmet
```

19.6.2.1.3 Endian check and conversion

Run Endian check and convert the datafiles if endian is different between source and destination:

1. For Endian check, run this on both source and destination database:

   ```
   SELECT endian_format

   FROM v$transportable_platform tp, v$database d

   WHERE tp.platform_name = d.platform_name;
   ```

   If the source platform and the target platform are of different endianness, then an additional step must be done on either the source or target platform to convert the tablespace being transported to the target format. If they are of the same endianness, then no conversion is necessary and tablespaces can be transported as if they were on the same platform.

   Example:

   ```
   Source Endian
   Linux IA (32-bit) - Little

   Destination Endian
   Solaris[tm] OE (32-bit) - Big
   ```

2. Ship datafiles, metadata dump to target and Convert datafiles using RMAN:

   Ship the datafiles and the metadata dump to target and On target convert all datafiles to destination endian:

   ```
   CONVERT DATAFILE
   '/d14/em10g/oradata/em102/mgmt.dbf',
   '/d14/em10g/oradata/em102/mgmt_ecm_depot1.dbf'
   FROM PLATFORM 'Linux IA (32-bit)';
   ```

   Conversion via RMAN can be done either on source or target (For more details refer RMAN doc). Parallelism can be used to speed up the process if the user tablespaces contains multiple datafiles.

19.6.2.1.4 Import metadata and plugin tablespaces

Use the following steps to import metadata and plugin tablespaces:

1. Run RepManager to drop the target repository (if the target database has the Enterprise Manager repository installed):

   ```
   RepManager repository_host repository_port repository_SID
   -sys_password password_for_sys_account -action drop
   ```

2. Run the pre-import steps to create the sysman user and grant privileges on the target database:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_repos_user.sql

   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   pre_import.sql
   ```

3. Invoke the Data Pump utility to plug the set of tablespaces into the target database:

```
impdp DUMPFILE=ttsem102.dmp DIRECTORY=data_pump_dir

TRANSPORT_
DATAFILES=/d14/em10g/oradata/em102/mgmt.dbf,/d14/em10g/oradat
a/em102/mgmt_ecm_depot1.dbf
```

4. Import other objects (packages, procedures, functions, and so on):

```
impdp CONTENT=METADATA_ONLY EXCLUDE=INDEX,CONSTRAINT
DUMPFILE=data_pump_dir:postexp.dmp LOGFILE=data_pump_
dir:postexp.log
```

19.6.2.1.5 Post Plug In Steps

Follow these post plug-in steps:

1. Run post plug-in steps to recompile any invalids, create public synonyms, create other users, enable VPD policy, repin packages:

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
create_synonyms.sql

@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
post_import.sql
```

   Check for invalid objects -- compare source and destination schemas for any discrepancy in the counts and invalids.

2. Bring user tablespaces back to read write mode:

```
alter tablespace MGMT_TABLESPACE read write;

alter tablespace MGMT_ECM_DEPOT_TS read write;
```

3. Submit Enterprise Manager RDBMS jobs.

   Reset back job_queue_processes to original value and run:

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
submit_dbms_jobs.sql
```

4. Update OMS properties and start the OMS.

   Update emoms.properties to reflect the migrated repository. Update host name - *oracle.sysman.eml.mntr.emdRepServer* and port with the correct value and start the OMS.

5. Relocate the Management Services and Repository target.

   If the Management Services and Repository target needs to be migrated to the destination host, run *em_assoc. handle_relocated_target* to relocate the target or recreate the target on the target host.

6. Discover/relocate database and database listener targets.

   Discover the target database and listener in Enterprise Manager or relocate the targets from source agent to destination agent.

19.6.2.2 Data Pump

Oracle Data Pump technology enables high-speed, parallel movement of bulk data and metadata from one database to another. Data Pump uses APIs to load and unload

data instead of usual SQL commands. Data pump operations can be run via Enterprise Manager interface and is very useful for cross platform database migration.

The migration of the database using the Data Pump export and Data Pump import tools comprises these steps: export the data into a dump file on the source server with the *expdp* command; copy or move the dump file to the target server; and import the dump file into Oracle on the target server by using the *impdp* command; and run post import Enterprise Manager specific steps.

Tuning parameters that were used in original Export and Import, such as BUFFER and RECORDLENGTH, are neither required nor supported by Data Pump Export and Import

19.6.2.2.1 Prepare for Data Pump

Use the following steps to prepare for data pump:

1. Review the following prerequisite for using Data pump for Enterprise Manager repository:

   Impdp fails for Enterprise Manager repository because of data pump bug - Bug 4386766 - IMPDP WITH COMPRESSED INDEXES FAILS WITH ORA-14071 AND ORA-39083. This bug is fixed in 10.2. Backport is available for 10.1.0.4. This RDBMS patch has to be applied to use expdp/impdp for the Enterprise Manager repository migration or workaround is to use exp/imp for extract and import.

2. Create the data pump directory:

   ```
   create directory data_pump_dir as
   '/scratch/user/EM102/ttsdata';
   ```

3. Shutdown OMS instances and prepare for migration.

   Shutdown the OMS, set job *queue_processes* to 0 and run `@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_remove_dbms_jobs.sql`

   To improve throughput of a job, the PARALLEL parameter should be used to set a degree of parallelism that takes maximum advantage of the current conditions. In general, the degree of parallelism should be set to more than twice the number of CPUs on an instance.

   All data pump actions are performed by multiple jobs (server processes not DBMS_JOB jobs). These jobs are controlled by a master control process which uses Advanced Queuing. At runtime an advanced queue table, named after the job name, is created and used by the master control process. The table is dropped on completion of the data pump job. The job and the advanced queue can be named using the JOB_NAME parameter.

   DBMS_DATAPUMP APIs can also be used to do data pump export/import. Please refer to Data pump section in 10g administration manual for all the options.

19.6.2.2.2 Data Pump Export

Use these steps to run data pump export:

1. Run data pump export:

   ```
   expdp FULL=y DUMPFILE=data_pump_dir:dpfull1%U.dmp, data_pump_dir:dpfull2%U.dmp
   PARALLEL=4 LOGFILE=data_pump_dir:dpexpfull.log JOB_NAME=dpexpfull
   Verify the logs for any errors during export
   ```

Data pump direct path export sometimes fails for mgmt_metrics_raw and raises ORA 600. This is due to Bug 4221775 (4233303). This bug is fixed in release 10.2. The workaround: if using expdp data pump for mgmt_metrics_raw, run expdp with ACCESS_METHOD+EXTERNAL_TABLE parameter.

```
expdp directory=db_export dumpfile=exp_st2.dmp logfile=exp_
st2.log tables=sysman.mgmt_metrics_raw access_
method=external_table
```

19.6.2.2.3 Data Pump Import

Use these steps to run data pump import:

1. Run RepManager to drop target repository (if target database has Enterprise Manager repository installed):

   ```
   RepManager repository_host repository_port repository_SID
   -sys_password password_for_sys_account -action drop
   ```

2. Prepare the target database:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_tablespaces.sql
   ```

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_repos_user.sql
   ```

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   pre_import.sql
   ```

3. Run data pump import:

   ```
   Impdp FULL=y DUMPFILE=data_pump_dir:dpfull1%U.dmp, data_pump_
   dir:dpfull2%U.dmp PARALLEL=4 LOGFILE=data_pump_
   dir:dpimpfull.log JOB_NAME=dpimpfull
   ```

   Verify the logs for any issues with the import.

19.6.2.2.4 Post Import Enterprise Manager Steps

Use the following steps for post import Enterprise Manager steps:

1. Run post plugin steps to recompile any invalids, create public synonyms, create other users, enable VPD policy, repin packages:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_synonyms.sql
   ```

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   post_import.sql
   ```

   Check for invalid objects - compare source and destination schemas for any discrepancy in counts and invalids.

2. Submit Enterprise Manager dbms jobs.

   Reset back *job_queue_processes* to original value and run:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   submit_dbms_jobs.sql
   ```

3. Update OMS properties and startup the OMS.

   Update emoms.properties to reflect the migrated repository. Update host name - *oracle.sysman.eml.mntr.emdRepServer* and port with the correct value and start the OMS.

4. Relocate Management Services and Repository target.

   If Management Services and the repository target needs to be migrated to the destination host, run *em_assoc. handle_relocated_target* to relocate the target or recreate the target on the target host.

5. Discover/relocate Database and database Listener targets.

   Discover the target database and listener in Enterprise Manager or relocate the targets from source agent to destination agent.

19.6.2.3 Export/Import

If the source and destination database is non-10g, then export/import is the only option for cross platform database migration.

For performance improvement of export/import, set higher value for BUFFER and RECORDLENGTH. Do not export to NFS as it will slow down the process considerably. Direct path can be used to increase performance. Note - As Enterprise Manager uses VPD, conventional mode will only be used by Oracle on tables where policy is defined.

Also User running export should have EXEMPT ACCESS POLICY privilege to export all rows as that user is then exempt from VPD policy enforcement. SYS is always exempted from VPD or Oracle Label Security policy enforcement, regardless of the export mode, application, or utility that is used to extract data from the database.

19.6.2.3.1 Prepare for Export/Import

Use the following steps to prepare for Export/Import:

1. Mgmt_metrics_raw partitions check:

   ```
   select table_name,partitioning_type type,
   partition_count count, subpartitioning_type subtype from
   dba_part_tables where table_name = 'MGMT_METRICS_RAW'
   ```

   If MGMT_METRICS_RAW has more than 3276 partitions please see Bug 4376351 - This is fixed in release 10.2. The work around is to export mgmt_metrics_raw in conventional mode.

2. Shutdown OMS instances and prepare for migration

   Shutdown OMS, set job *queue_processes* to 0 and run `@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_remove_dbms_jobs.sql`

19.6.2.3.2 Export

Follow these steps for export:

1. Export data:

   ```
   exp full=y constraints=n indexes=n compress=y file=fullem102_
   1.dmp log=fullem102exp_1.log
   ```

2. Export without data and with constraints:

   ```
   exp full=y constraints=y indexes=y rows=n ignore=y
   file=fullem102_2.dmp log=fullem102exp_2.log
   ```

19.6.2.3.3 Import

Follow these steps to import:

1. Run RepManager to drop the target repository (if the target database has the Enterprise Manager repository installed):

   ```
   RepManager repository_host repository_port repository_SID
   -sys_password password_for_sys_account -action drop
   ```

2. Pre-create the tablespaces and the users in target database:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_tablespaces.sql
   ```

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_repos_user.sql
   ```

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   pre_import.sql
   ```

3. Import data:

   ```
   imp full=y constraints=n indexes=n file=fullem102_1.dmp
   log=fullem102imp_1.log
   ```

4. Import without data and with constraints:

   ```
   imp full=y constraints=y indexes=y rows=n ignore=y
   file=fullem102_2.dmp log=fullem102imp_2.log
   ```

19.6.2.3.4 Post-Import Enterprise Manager Steps

Follow these steps for post-import Enterprise Manager steps:

1. Run post plug-in steps to recompile any invalids, create public synonyms, create other users, enable VPD policy, repin packages:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   create_synonyms.sql
   ```

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   post_import.sql
   ```

   Check for invalid objects -- compare source and destination schemas for any discrepancy in counts and invalids.

2. Submit the Enterprise Manager dbms jobs.

   Reset back *job_queue_processes* to its original value and run:

   ```
   @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_
   submit_dbms_jobs.sql
   ```

3. Update the OMS properties and startup the OMS:

   Update emoms.properties to reflect the migrated repository. Update host name *oracle.sysman.eml.mntr.emdRepServer* and port with the correct value and start the OMS.

4. Relocate Management Services and the Repository target.

   If Management Services and the repository target need to be migrated to the destination host, run *em_assoc. handle_relocated_target* to relocate the target or recreate the target on the target host.

5. Discover/relocate Database and database Listener targets.

   Discover the target database and listener in Enterprise Manager or relocate the targets from the source agent to the destination agent.

19.6.3 Post Migration Verification

These verification steps should be carried out post migration to ensure that the migration was completely successful:

- Verify any discrepancy in objects by comparing source and target databases through Enterprise Manager.
- Verify the migrated database through Enterprise Manager to determine whether the database is running without any issues.
- Verify the repository operations, dbms jobs and whether any management system errors are reported.
- Verify that all Enterprise Manager functionalities are working correctly after the migration.
- Make sure Management Services and the Repository target is properly relocated by verifying it through Enterprise Manager.

Part VI

Configuring Enterprise Manager for High Availability

This section covers Enterprise Manager high availability best practices and strategies that allow you to safeguard your Oracle Enterprise Manager installation.

- Chapter 20, "High Availability Solutions"
- Chapter 21, "Setting Up High Availability"
- Chapter 22, "Backing Up Enterprise Manager"
- Chapter 23, "Enterprise Manager Outages"

20

High Availability Solutions

Highly Available systems are critical to the success of virtually every business today. It is equally important that the management infrastructure monitoring these mission-critical systems are highly available. The Enterprise Manager Cloud Control architecture is engineered to be scalable and available from the ground up. It is designed to ensure that you concentrate on managing the assets that support your business, while it takes care of meeting your business Service Level Agreements.

When you configure Cloud Control for high availability, your aim is to protect each component of the system, as well as the flow of management data in case of performance or availability problems, such as a failure of a host or a Management Service.

Maximum Availability Architecture (MAA) provides a highly available Enterprise Manager implementation by guarding against failure at each component of Enterprise Manager.

The impacts of failure of the different Enterprise Manager components are:

- Management Agent failure or failure in the communication between Management Agents and Management Service

  Results in targets no longer monitored by Enterprise Manager, though the Enterprise Manager console is still available and one can view historical data from the Management Repository.

- Management Service failure

  Results in the unavailability of Enterprise Manager console, as well as unavailability of almost all Enterprise Manager services.

- Management Repository failure

  Results in failure on the part of Enterprise Manager to save the uploaded data by the Management Agents as well as unavailability of almost all Enterprise Manager services.

Overall, failure in any component of Enterprise Manager can result in substantial service disruption. Therefore it is essential that each component be hardened using a highly available architecture.

20.1 Latest High Availability Information

Because of rapidly changing technology, and the fact that high availability implementations extend beyond the realm of Oracle Enterprise Manager, the following resources should be checked regularly for the latest information on third-party

integration with Oracle's high availability solutions (F5 or third-party cluster ware, for example).

- Oracle Maximum Availability Architecture Web site

  http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

- Support Note 330072.1: "How To Configure Grid Control Components for High Availability "

20.2 Defining High Availability

Oracle Enterprise Manager's flexible, distributed architecture permits a wide range of deployment configurations, allowing it to meet the monitoring and management needs of your business, as well as allowing for expansion as business needs dictate.

For this reason, high availability for Enterprise Manager cannot be narrowly defined as a singular implementation, but rather a range of protection levels based on your available resources, Oracle technology and best practices that safeguard the investment in your IT infrastructure. Depending on your Enterprise Manager deployment and business needs, you can implement the level of high availability necessary to sustain your business. High availably for Enterprise Manager can be categorized into four levels, each level building on the previous and increasing in implementation cost and complexity, but also incrementally increasing the level of availability.

20.2.1 Levels of High Availability

Each high availability solution level is driven by your business requirements and available IT resources. However, it is important to note that the levels represent a subset of possible deployments that are useful in presenting the various options available. Your IT organization will likely deploy its own configuration which need not exactly match one of the levels.

The following table summarizes the four example high availability levels for Oracle Enterprise Manager installations as well as general resource requirements.

Table 20–1 Enterprise Manager High Availability Levels

Level	Description	Minimum Number of Nodes	Recommended Number of Nodes	Load Balancer Requirements
Level 1	OMS and repository database. Each resides on their own host with no failover.	1	2	None
Level 2	OMS installed on shared storage with a VIP based failover. Database is using Local Data Guard.	2	4	None
Level 3	OMS in Active/Active configuration. The database is using RAC + Local Data Guard	3	5	Local Load Balancer

Table 20–1 (Cont.) Enterprise Manager High Availability Levels

Level	Description	Minimum Number of Nodes	Recommended Number of Nodes	Load Balancer Requirements
Level 4	OMS in Active/Active configuration on the primary site. The standby RAC database (DataGuard) is located at the disaster recovery site. Multiple standby OMS instances at remote site. Data Guard RAC database at the primary site **Note**: Level 4 is a MAA Best Practice, achieving highest availability in the most cost effective, simple architecture.	4	8	**Required**: Local Load Balancer for each site. **Optional**: Global Load Balancer

20.3 Comparing Availability Levels

The following tables compare the protection levels and recovery times for the various high-availability levels.

Table 20–2 High Availability Levels of Protection

Level	OMS Host Failure	OMS Storage Failure	Database Host Failure	Database Storage Failure	Site Failure/Disaster Recovery
Level 1	No	No	No	No	No
Level 2	Yes	No	Yes	Yes	No
Level 3	Yes	Yes	Yes	Yes	No
Level 4	Yes	Yes	Yes	Yes	Yes

Table 20–3 High Availability Level Recovery Times

Level	Node Failure	Local Storage Failure	Site Failure	Cost
Level 1	Hours-Days	Hours-Days	Hours-Days	$
Level 2	Minutes	Hours-Days	Hours-Days	$$
Level 3	No Outage	Minutes	Hours-Days	$$$
Level 4	No Outage	Minutes	Minutes	$$$$

One measure that is not represented in the tables is that of scalability. Levels three and four provide the ability to scale the Enterprise Manager installation as business needs grow. The repository, running as a RAC database, can easily be scaled upwards by adding new nodes to the RAC cluster and it is possible to scale the Management Service tier by simply adding more OMS servers.

If you need equalized performance in the event of failover to a standby deployment, whether that is a local standby database or a Level four standby site including a standby RAC database and standby OMS servers, it is essential to ensure that the deployments on both sites are symmetrically scaled. This is particularly true if you

want to run through planned failover routines where you actively run on the primary or secondary site for extended periods of time. For example, some finance institutions mandate this as part of operating procedures.

If you need survivability in the event of a primary site loss you need to go with a Level four architecture.

20.4 Implementing High Availability Levels

Once you have determined the high availability requirements for your enterprise, you are ready to begin implementing one of the high availability levels that is suitable for your environment. Use the following information roadmap to find implementation instructions for each level.

Level	Where to find information
Level 1	*Oracle Enterprise Manager Basic Installation Guide* and the *Oracle Enterprise Manager Advanced Installation and Configuration Guide*
Level 2	*Oracle Enterprise Manager Basic Installation Guide* and the *Oracle Enterprise Manager Advanced Installation and Configuration Guide* PLUS ■ Configuring the Cloud Control OMS in an Active/Passive Environment for High Availability Failover Using Virtual Host Names ■ Configuring a Standby Database for the Management Repository
Level 3	*Oracle Enterprise Manager Basic Installation Guide* and the *Oracle Enterprise Manager Advanced Installation and Configuration Guide* PLUS ■ Configuring the First Management Service for High Availability ■ Configuring the Software Library ■ Configuring a Load Balancer ■ Adding Additional Oracle Management Service (Oracle Enterprise Manager Basic Installation Guide, Chapter 7) ■ Configuring a Standby Database for the Management Repository
Level 4	*Oracle Enterprise Manager Basic Installation Guide* and the *Oracle Enterprise Manager Advanced Installation and Configuration Guide* PLUS ■ Configuring a Standby Database for the Management Repository ■ Configuring Standby Management Services on a Standby Site

21

Setting Up High Availability

This chapter discusses best practices for installation and configuration of each Cloud Control component and covers the following topics:

- Installation Best Practices for Enterprise Manager High Availability
- Configuring a Standby Database for the Management Repository
- Configuring Management Service to RAC Management Repository Communication
- Configuring the First Management Service for High Availability
- Configuring the Cloud Control OMS in an Active/Passive Environment for High Availability Failover Using Virtual Host Names
- Configuring the Software Library
- Configuring a Load Balancer
- Configuring Standby Management Services on a Standby Site

21.1 Installation Best Practices for Enterprise Manager High Availability

The following sections document best practices for installation and configuration of each Cloud Control component.

21.1.1 Configuring the Management Agent to Automatically Start on Boot and Restart on Failure

The Management Agent is started manually. It is important that the Management Agent be automatically started when the host is booted to insure monitoring of critical resources on the administered host. To that end, use any and all operating system mechanisms to automatically start the Management Agent. For example, on UNIX systems this is done by placing an entry in the UNIX `/etc/init.d` that calls the Management Agent on boot or by setting the Windows service to start automatically.

21.1.2 Configuring Restart for the Management Agent

Once the Management Agent is started, the watchdog process monitors the Management Agent and attempts to restart it in the event of a failure. The behavior of the watchdog is controlled by environment variables set before the Management Agent process starts. The environment variables that control this behavior follow. All testing discussed here was done with the default settings.

- EM_MAX_RETRIES – This is the maximum number of times the watchdog will attempt to restart the Management Agent within the EM_RETRY_WINDOW. The default is to attempt restart of the Management Agent three times.
- EM_RETRY_WINDOW - This is the time interval in seconds that is used together with the EM_MAX_RETRIES environmental variable to determine whether the Management Agent is to be restarted. The default is 600 seconds.

The watchdog will not restart the Management Agent if the watchdog detects that the Management Agent has required restart more than EM_MAX_RETRIES within the EM_RETRY_WINDOW time period.

21.1.3 Installing the Management Agent Software on Redundant Storage

The Management Agent persists its configuration, intermediate state and collected information using local files in the `agent_inst` directory.

In the event that these files are lost or corrupted before being uploaded to the Management Repository, a loss of monitoring data and any pending alerts not yet uploaded to the Management Repository occurs.

To protect from such losses, configure these sub-directories on striped redundant or mirrored storage. The Management Agent agent_inst directory is shown by entering the command 'emctl getemhome' on the command line, or from the Management Services and Repository > Agents tab in the Cloud Control console.

21.2 Installation Best Practices for Enterprise Manager Management Repository High Availability

Before installing Enterprise Manager, you should prepare the database, which will be used for setting up Management Repository. Install the database using Database Configuration Assistant (DBCA) to make sure that you inherit all Oracle install best practices.

For both high availability and scalability, you should configure the Management Repository in the latest certified database version, with the RAC option enabled. Check for the latest version of database certified for Enterprise Manager from the Certify tab on the My Oracle Support Web site.

- Choose Automatic Storage Management (ASM) as the underlying storage technology.

When the database installation is complete:

- Go to $ORACLE_HOME/rbdms/admin directory of the database home and execute the 'dbmspool.sql'

  This installs the DBMS_SHARED_POOL package, which will help in improving throughput of the Management Repository.

Management Repository Configuration to perform after installing the Management Service

There are some parameters that should be configured during the Management Repository database install (as previously mentioned) and some parameters that should be set after the Management Service has been installed.

Start by installing Management Agents on each Management Repository node. Once the Management Agents are installed and the Management Repository database is

discovered as a target, the Enterprise Manager console can be used to configure these best practices in the Management Repository.

- Enable ARCHIVELOG Mode
- Enable Block Checksums
- Configure the Size of Redo Log Files and Groups Appropriately
- Use a Flash Recovery Area
- Enable Flashback Database
- Use Fast-Start Fault Recovery to Control Instance Recovery Time
- Enable Database Block Checking
- Set DISK_ASYNCH_IO

The details of these settings are available in Oracle Database High Availability Best Practices.

Use the MAA Advisor for additional high availability recommendations that should be applied to the Management Repository. To access the MAA Advisor:

1. On the Database Target Home page, locate the High Availability section.
2. Click **Details** next to the Console item.
3. In the Availability Summary section of the High Availability Console page, click **Details** located next to the MAA Advisor item.

21.3 Configuring a Standby Database for the Management Repository

Prerequisites

The standby site must be similar to the primary site in terms of hardware and network resources to ensure there is no loss of performance when failover happens.

There must be sufficient network bandwidth between the primary and standby sites to handle peak redo data generation.

Configuration Steps

The following steps describe the procedure for setting up a standby Management Repository database.

1. Prepare Standby Management Repository hosts for Data Guard

   Install a Management Agent on each of the standby Management Repository hosts. Configure the Management Agents to upload by the SLB on the primary site. Install Grid infrastructure and RAC Database software on the standby Management Repository hosts. The version used must be the same as that on the primary site.

2. Prepare the Primary Management Repository database for Data Guard

   If the primary Management Repository database is not already configured, enable archive log mode, setup flash recovery area and enable flashback database on the primary Management Repository database.

3. Create the Physical Standby Database

   Use the Enterprise Manager console to set up a physical standby database in the standby environment. The Standby Management Repository database must be a

Physical Standby. Logical standby Management Repository databases are not supported.

The Enterprise Manager Console does not support creating a standby RAC database. If the standby database has to be RAC, configure the standby database using a single instance and then use the 'Convert to RAC' option from the Enterprise Manager Console to convert the single instance standby database to RAC. Note that the Convert to RAC option is available for Oracle Database releases 10.2.0.5, 11.1.0.7, and above. Oracle Database release 11.1.0.7 requires patch 8824966 for the Convert to RAC option to work.

During single instance standby creation, best practice is to create the database files on shared storage, ideally ASM, to facilitate conversion to RAC later.

4. Add Static Service to the Listener.

   To enable Data Guard to restart instances during the course of broker operations, a service with a specific name must be statically registered with the local listener of each instance. The value for the GLOBAL_DBNAME attribute must be set to a concatenation of <db_unique_name>_DGMGRL.<db_domain>. For example, in the LISTENER.ORA file:

   ```
   SID_LIST_LISTENER=(SID_LIST=(SID_DESC=(SID_NAME=sid_name)
        (GLOBAL_DBNAME=db_unique_name_DGMGRL.db_domain)
        (ORACLE_HOME=oracle_home)))
   ```

5. Enable Flashback Database on the Standby Database.

   To allow re-instate of an old primary database as a standby database after a failover, flashback database must be enabled. Hence do so for both the primary and the standby databases.

6. To allow Enterprise Manager to monitor a Physical Standby database (which is typically in a mounted state), specify sysdba monitoring privileges. This can be specified either during the Standby creation wizard itself or post creation by modifying the Monitoring Configuration for the standby database target.

7. Verify the Physical Standby

   Verify the Physical Standby database through the Enterprise Manager Console. Click the Log Switch button on the Data Guard page to switch log and verify that it is received and applied to the standby database.

21.4 Configuring Management Service to RAC Management Repository Communication

If the Management Repository is a Real Application Cluster (RAC) database, the Management Service processes need to be configured to communicate with each node of the RAC database in a redundant fashion.

Note that RAC nodes are referred to by their virtual IP (vip) names. The service_name parameter is used instead of the system identifier (SID) in connect_data mode and failover is turned on. Refer to the *Oracle Database Net Services Administrator's Guide* for details.

Configure the repository connect descriptor by running the emctl command from any Management Service:

```
emctl config oms -store_repos_details -repos_conndesc '"(DESCRIPTION=
(ADDRESS_LIST=(FAILOVER=ON)
(ADDRESS=(PROTOCOL=TCP)(HOST=node1-vip.example.com)(PORT=1521))
```

```
(ADDRESS=(PROTOCOL=TCP)(HOST=node2-vip.example.com)(PORT=1521)))
(CONNECT_DATA=(SERVICE_NAME=EMREP)))"' -repos_user sysman
```

After making the previous change, run the following command from any one OMS to make the same change to the monitoring configuration used for the Management Services and Repository target:

```
emctl config emrep -conn_desc <repository_connect descrptor as above>
```

21.5 Configuring the First Management Service for High Availability

The Management Service itself has a built-in restart mechanism based on the Oracle Weblogic Node Manager and the Oracle Process Management and Notification Service (OPMN). These services will attempt to restart a Management Service that is down. It is advisable to run OPMN and Node Manager as operating system services, so that they restart automatically if their host machine is restarted.

21.5.1 Management Service Install Location

If you are managing a large environment with multiple Management Services and Management Repository nodes, then consider installing the Management Services on hardware nodes that are different from Management Repository nodes (Figure 21–1). This allows you to scale out the Management Services in the future.

Figure 21–1 Management Service and Management Repository on Separate Hardware

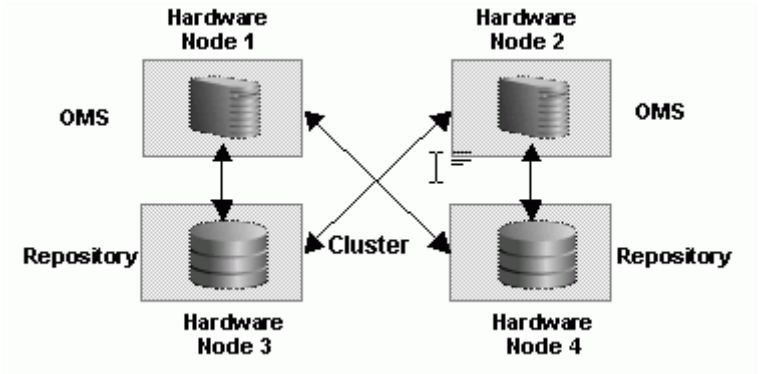

Also consider the network latency between the Oracle Management Service and the Management Repository while determining the Management Service install location. The distance between the Management Service and the Management Repository may be one of the factors that affect network latency and hence determine Enterprise Manager performance.

If the network latency between the Oracle Management Service and Management Repository tiers is high or the hardware available for running Enterprise Manager is limited, then the Management Service can be installed on the same hardware as the Management Repository (Figure 21–2). This allows for Enterprise Manager high availability, as well as keep the costs down.

Figure 21–2 Management Service and Management Repository on Same Hardware

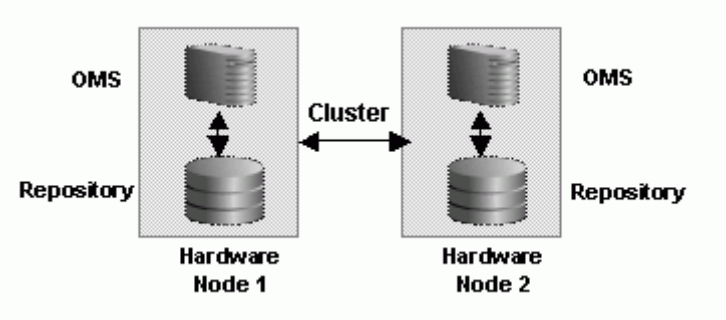

If you plan ahead, you can configure your Enterprise Manager deployment for high availability by choosing the correct options during the first Management Service install. You can also retrofit the MAA best practices into an existing Enterprise Manager deployment configured initially using the default install options.

21.6 Configuring the Cloud Control OMS in an Active/Passive Environment for High Availability Failover Using Virtual Host Names

This section provides a general reference for Cloud Control administrators who want to configure Enterprise Manager Cloud Control in Cold Failover Cluster (CFC) environments.

21.6.1 Overview and Requirements

The following conditions must be met for Cloud Control to fail over to a different host:

- The installation must be done using a Virtual Host Name and an associated unique IP address.
- Install on a shared disk/volume which holds the binaries and the gc_inst directory.
- The Inventory location must failover to the surviving node.
- The software owner and time zone parameters must be the same on all cluster member nodes that will host this Oracle Management Service (OMS).

21.6.2 Installation and Configuration

To override the physical host name of the cluster member with a virtual host name, software must be installed using the parameter ORACLE_HOSTNAME.

The software must be installed using the command line parameter -invPtrLoc to point to the shared inventory location file, which includes the path to the shared inventory location.

If you are using an NFS mounted volume for the installation, please ensure that you specify rsize and wsize in your mount command to prevent running into I/O issues.

For example:

```
oms.example.com:/u01/app/share1 /u01/app/share1 nfs
rw,bg,rsize=32768,wsize=32768,hard,nointr,tcp,noac,vers=3,timeo=600 0 0
```

Note: Any reference to shared failover volumes could also be true for non-shared failover volumes which can be mounted on active hosts after failover.

21.6.3 Setting Up the Virtual Host Name/Virtual IP Address

You can set up the virtual host name and virtual IP address by either allowing the clusterware to set it up, or manually setting it up yourself before installation and startup of Oracle services. The virtual host name must be static and resolvable consistently on the network. All nodes participating in the setup must resolve the virtual IP address to the same host name. Standard TCP tools such as nslookup and traceroute can be used to verify the host name. Validate using the following commands:

```
nslookup <virtual hostname>
```

This command returns the virtual IP address and full qualified host name.

```
nslookup <virtual IP>
```

This command returns the virtual IP address and fully qualified host name.

Be sure to try these commands on every node of the cluster and verify that the correct information is returned.

21.6.4 Setting Up Shared Storage

Storage can be managed by the clusterware that is in use or you can use any shared file system (FS) volume, such as NFS, as long as it is not an unsupported type, such as OCFS V1.

Note: Only OCFS V1 is not supported. **All other versions of OCFS are suppored.**

If the OHS directory is on a shared storage, the LockFile directive in the httpd.conf file should be modified to point to a local disk, otherwise there is a potential for locking issues.

21.6.5 Setting Up the Environment

Some operating system versions require specific operating system patches be applied prior to installing 12c. The user installing and using the 12c software must also have sufficient kernel resources available. Refer to the operating system's installation guide for more details. Before you launch the installer, certain environment variables need to be verified. Each of these variables must be identically set for the account installing the software on ALL machines participating in the cluster:

- **OS variable TZ**

  Time zone setting. You should unset this variable prior to installation.

- **PERL variables**

  Variables such as PERL5LIB should also be unset to avoid association to the incorrect set of PERL libraries

21.6.6 Synchronizing Operating System IDs

The user and group of the software owner should be defined identically on all nodes of the cluster. This can be verified using the 'id' command:

```
$ id -a

uid=550(oracle) gid=50(oinstall) groups=501(dba)
```

21.6.7 Setting Up Shared Inventory

Use the following steps to set up shared inventory:

1. Create your new ORACLE_HOME directory.
2. Create the Oracle Inventory directory under the new ORACLE_HOME:

   ```
   $ cd <shared oracle home>

   $ mkdir oraInventory
   ```

3. Create the oraInst.loc file. This file contains the Oracle Inventory directory path information needed by the Universal Installer.

   ```
   vi oraInst.loc
   ```

   Enter the path information to the Oracle Inventory directory and specify the group of the software owner as the oinstall user. For example:

   ```
   inventory_loc=/app/oracle/share1/oraInventory

   inst_group=oinstall
   ```

21.6.8 Installing the Software

Refer to the following steps when installing the software:

1. Create the shared disk location on both the nodes for the software binaries.
2. Point to the inventory location file oraInst.loc (under the ORACLE_BASE in this case), as well as specifying the host name of the virtual group. For example:

   ```
   $ runInstaller -invPtrloc /app/oracle/share1/oraInst.loc
   ORACLE_HOSTNAME=lxdb.example.com -debug
   ```

3. Install Oracle Management Services on cluster member Host1.
4. Continue the remainder of the installation normally.
5. Once completed, copy the files oraInst.loc and oratab to /etc on all cluster member hosts (Host2, Host3, ...)

Windows Specific Configuration Steps

On Windows environments, an additional step is required to copy over service and keys required by the Oracle software. Note that these steps are required if your clustering software does not provide a shared windows registry.

1. Using regedit on the first host, export each Oracle service from under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services.
2. Using regedit on the first host, export HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE.
3. Use regedit to import the files created in step 1 and 2 to the failover host.

21.6.9 Starting Up Services

Ensure that you start your services in the proper order. Use the order listed below:

1. Establish the IP address on the active node.
2. Start the TNS listener (if it is part of the same failover group).
3. Start the database (if it is part of the same failover group).
4. Start Cloud Control using `emctl start oms`
5. Test functionality.

In case of failover, refer to Chapter 23, "Enterprise Manager Outages".

21.7 Configuring the Software Library

The software library location must be accessed by all Management Services. The configuration of software library is not performed during installation and needs to be configured post-install using the Enterprise Manager Console:

1. On the Enterprise Manager home page, from the **Setup** menu, select **Provisioning and Patching**, and then select **Software Library**.
2. Click the **Provisioning** subtab.
3. On the Provisioning page, click the **Administration** subtab.
4. In the Software Library Configuration section, click **Add** to set the Software Library Directory Location to a shared storage that can be accessed by any host running the Management Service.

21.8 Configuring a Load Balancer

This section describes the guidelines for setting up a Server Load Balancer (SLB) to balance the agent and browser traffic to multiple Management Services.

Server Load Balancer Requirements

In order to configure your OMS's in an active/active configuration behind an SLB, your SLB must meet the following requirements:

- The SLB must provide support for multiple virtual server ports.

  Enterprise Manager typically requires that up to 4 ports are configured on the SLB (Secure Upload, Agent Registration, Secure Console, Unsecure Console)

- Support for persistence.

  HTTP and HTTPS traffic between the browser and the OMS requires persistence.

- Support for application monitoring.

  The SLB must be capable of monitoring the health of the OMSs and detecting failures, so that requests will not be routed to OMSs that are not available.

SLB configuration is a two-step process:

1. Configure the SLB
2. Make requisite changes on the Management Services.

21.8.1 SLB Setup

Use the following table as reference for setting up the SLB with Cloud Control Management Services.

Table 21–1 Management Service Ports

Cloud Control Service	TCP Port	Monitor Name	Persistence	Pool Name	Load Balancing	Virtual Server Name	Virtual Server Port
Secure Upload	4900	mon_gcsu4900	None	pool_ gcsu4900	Round Robin	vs_ gcsu4900	4900
Agent Registration	4889	mon_gcar4889	Active Cookie Insert	pool_ gcar4889	Round Robin	vs_ gcar4889	4889
Secure Console	7799	mon_gcsc7799	Source IP	pool_ gcsc7799	Round Robin	vs_ gcsc443	443
Unsecure Console (optional)	7788	mon_gcuc7788	Source IP	pool_ gcuc7788	Round Robin	vs_gcuc80	80

Use the administration tools that are packaged with your SLB. A sample configuration follows. This example assumes that you have two Management Services running on host A and host B using the default ports as listed in Table 33–1.

1. Create Pools

   A *pool* is a set of servers grouped together to receive traffic on a specific TCP port using a load balancing method. Each pool can have its own unique characteristic for a persistence definition and the load-balancing algorithm used.

Table 21–2 Pools

Pool Name	Usage	Members	Persistence	Load Balancing
pool_gcsu4900	Secure upload	HostA:4900 HostB:4900	None	Round Robin
pool_gcar4889	Agent registration	HostA:4889 HostB:4889	Active cookie insert; expiration 60 minutes	Round Robin
pool_gcsc7799	Secured console access	HostA:7799 HostB:7799	Source IP; expiration 60 minutes	Round Robin
pool_gcuc7788 (optional)	Unsecured console access	HostA:7788 HostB:7788	Source IP; expiration 60 minutes	Round Robin

2. Create Virtual Servers

   A *virtual server*, with its virtual IP Address and port number, is the client addressable hostname or IP address through which members of a load balancing pool are made available to a client. After a virtual server receives a request, it directs the request to a member of the pool based on a chosen load balancing method.

Table 21–3 Virtual Servers

Virtual Server Name	Usage	Virtual Server Port	Pool
vs_gcsu4900	Secure upload	4900	pool_gcsu4900
vs_gcar4889	Agent registration	4889	pool_gcar4889
vs_gcsc443	Secure console access	443	pool_gcsc7799
vs_gcuc80 (optional)	Unsecure console access	80	pool_gcuc7788

3. Create Monitors

Monitors are used to verify the operational state of pool members. Monitors verify connections and services on nodes that are members of load-balancing pools. A monitor is designed to check the status of a service on an ongoing basis, at a set interval. If the service being checked does not respond within a specified timeout period, the load balancer automatically takes it out of the pool and will choose the other members of the pool. When the node or service becomes available again, the monitor detects this and the member is automatically accessible to the pool and able to handle traffic.

Table 21–4 Monitors

Monitor Name	Configuration	Associate With
mon_gcsu4900	Type: https Interval: 60 Timeout: 181 Send String: GET /empbs/upload Receive String: Http Receiver Servlet active!	HostA:4900 HostB:4900
mon_gcar4889	Type: http Interval: 60 Timeout: 181 Send String: GET /empbs/genwallet Receive String: GenWallet Servlet activated	HostA:4889 HostB:4889
mon_gcsc7799	Type: https Interval: 5 Timeout: 16 Send String: GET /em/console/home HTTP/1.0\n Receive String: /em/login.jsp	HostA:7799 HostB:7788
mon_gcuc7788 (optional)	Type: https Interval: 5 Timeout: 16 Send String: GET /em/console/home HTTP/1.0\n Receive String: /em/login.jsp	HostA:7788 HostB:7788

Note: If you have SSO configured, use the following alternate definitions for the mon_gcsc7799 and mon_gcuc7788 monitors.

Table 21–5 Monitors for SSO Configuration

Monitor Name	Configuration	Associate With
mon_gcsc7799	Type: https Interval: 5 Timeout: 16 Send String: GET /empbs/genwallet Receive String: GenWallet Servlet activated	HostA:7799 HostB:7788
mon_gcuc7788 (optional)	Type: https Interval: 5 Timeout: 16 Send String: GET /empbs/genwallet Receive String: GenWallet Servlet activated	HostA:7788 HostB:7788

Note: F5 SLB monitors expect the "Receive String" within the first 5120 characters of a response. For SSO environments, the "Receive String" may be returned at some point beyond the 5120 limit. The monitor will not function in this situation.

21.8.2 Enterprise Manager Side Setup

Perform the following steps:

1. Resecure the Oracle Management Service

   By default, the service name on the Management Service-side certificate uses the name of the Management Service host. Management Agents do not accept this certificate when they communicate with the Oracle Management Service through a load balancer. You must run the following command to regenerate the certificate on each Management Service:

   ```
   emctl secure oms
     -sysman_pwd <sysman_pwd>
     -reg_pwd <agent_reg_password>
     -host slb.example.com
     -secure_port 4900
     -slb_port 4900
     -slb_console_port 443
     -console
     [-lock]  [-lock_console]
   ```

2. Resecure all Management Agents

   Management Agents that were installed prior to SLB setup, including the Management Agent that comes with the Management Service install, would be uploading directly to the Management Service. These Management Agents will not be able to upload after SLB is setup. Resecure these Management Agents to upload to the SLB by running the following command on each Management Agent:

   ```
   emctl secure agent –emdWalletSrcUrl https://slb.example.com:<upload port>/em
   ```

21.9 Configuring Standby Management Services on a Standby Site

Consider the following before installing the standby Management Services.

Oracle recommends that this activity be done during a lean period or during a planned maintenance window. When new Oracle Management Service instances are installed on the standby site, they are initially configured to connect to the Management Repository database on the primary site. Some workload will be taken up by the new Management Service. This could result in temporary loss in performance if the standby site Management Services are located far away from the primary site Management Repository database. However there would be no data loss and the performance would recover once the standby Management Services are shutdown post configuration.

Prerequisites

- The primary site must be configured as per Cloud Control MAA guidelines described in previous sections. This includes Management Services fronted by an SLB and all Management Agents configured to upload to Management Services by the SLB.
- The standby site must be similar to the primary site in terms of hardware and network resources to ensure there is no loss of performance when failover happens.
- Configure storage used by the software library to be replicated at the primary and standby site. In the event of a site outage, the contents of this storage must be made available on the standby site using hardware vendor disk level replication technologies.
- The shared storage used for the software library must be made available on the standby site using the same paths as the primary site.
- For complete redundancy in a disaster recovery environment, a second load balancer must be installed at the standby site. The secondary SLB must be configured in the same fashion as the primary. Some SLB vendors (such as F5 Networks) offer additional services that can be used to pass control of the Virtual IP presented by the SLB on the primary site to the SLB on the standby site in the event of a site outage. This can be used to facilitate automatic switching of Management Agent traffic from the primary site to the standby site.

21.9.1 Installing the First Standby Management Service

Install the first standby Management Service using the following steps:

1. Copy the emkey to the Management Repository by running the following command on the first Management Service on the primary site:

   ```
   emctl config emkey -copy_to_repos
   ```

2. Export the configuration from the first Management Service on the primary site using:

   ```
   emctl exportconfig oms -dir <location for the export file>
   ```

   After the configuration is exported, do not make any configuration changes to the primary site still the standby management service is configured.

3. Install a Management Agent on the standby host if one does not already exist.

4. Perform a software-only install of the Enterprise Manager software using a modified version of the "Add Management Service" Deployment Procedure.

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Procedure Library**.
2. Select **Add Management Service procedure** and click **Create Like**.
3. Go to the **Procedure Steps** tab and select and disable the steps - “Configure Management Service”, “Targets Discovery” and “Post Configuration Tasks”.
4. Save the modified deployment procedure and use it to install the Enterprise Manager software on the standby OMS host.
5. After the Deployment Procedure completes, delete the file emInstanceMapping.properties from <OMS Oracle Home>/sysman/config on the standby OMS host.

5. Configure the Management Service by running omsca in standby mode. Choose a different domain name for the standby. For example, if the primary WebLogic domain is GCDomain, choose GCDomainStby.

   ```
   omsca standby –EM_DOMAIN_NAME GCDomainStby –NM_USER
   nodemanager -AS_USERNAME weblogic –nostart
   ```

   When prompted for the Administration Server host and EM Instance host, enter the standby OMS hostname (or accept the default).

   When prompted for the passwords, provide the same passwords as the primary site.

   When prompted for Management Repository details, provide the Primary database details.

6. Configure the required plugins by running the following command:

   ```
   pluginca -action deploy -isFirstOMS true -plugins
   <plugin-list> -oracleHome <oms oracle home> -middlewareHome
   <wls middleware home>
   ```

   where plugin-list is the list of plugins returned by the SQL query

   ```
   SELECT epv.plugin_id, epv.version FROM em_plugin_version epv,
   em_current_deployed_plugin ecp WHERE epv.plugin_type NOT IN (
   'BUILT_IN_TARGET_TYPE' , 'INSTALL_HOME') AND ecp.dest_
   type='2' AND epv.plugin_version_id = ecp.plugin_version_id;
   ```

   and is a comma separate list in the following format:
   `<plugin-id>=<plugin-version>,<plugin-id>=<plugin-version>,…`

   Example:

   ```
   "oracle.sysman.empa=12.1.0.1.0,oracle.sysman.mos=12.1.0.1.0,o
   racle.sysman.emas=12.1.0.1.0,oracle.sysman.emfa=12.1.0.1.0,or
   acle.sysman.db=12.1.0.1.0,oracle.sysman.emct=12.1.0.1.0,oracl
   e.sysman.vt=12.1.0.1.0,oracle.sysman.ssa=12.1.0.1.0"
   ```

7. Copy over the configuration exported from the Primary Management Service in step 2 above to the standby Management Service host. Import the exported configuration on the standby Management Service using:

   ```
   emctl importconfig oms -file <full path of the export file>
   ```

   Note this command emits a warning about a failed export and prompts for confirmation to proceed. The warning can be ignored by entering "y" to proceed.

   Note this command will start the Management Service.

8. Stop the Management Service but leave the Administration Server running using:

   ```
   emctl stop oms
   ```

9. Add the standby Weblogic Domain and associated targets:

   The standby Weblogic Domain and associated targets can be added using the Guided Discovery process.

   From the **Setup** menu, select **Add Target** , then select **Add Targets Manually**. Select *Oracle Fusion Middleware* from the Target Types menu. Use the secure port (typically 7101) and, under **Advanced**, set the JMX Protocol to **t3s**.

   Note that the WebLogic targets, except the Administration Server, will be shown as down as the standby OMS is down at this stage.

10. If you have Single Sign On configured on the primary site, follow the same steps to configure SSO on the standby OMS.

11. If you have Real User Experience Insight, AD4J Manager or BI Publisher configured on the primary site, follow the same steps to configure them on the standby OMS.

21.9.2 Installing Additional Standby Management Services

It is recommended that your standby site be similar in configuration as your primary site. This means configuring multiple OMS on your standby site, similar to your primary site. Install additional standby Management Services as per the procedure listed below under "Additional Standby Management Services".

If, however, you choose to start with a single OMS on the standby site initially, you may skip this step and continue with the next section "Validating your installation and Complete the Setup". If you decide to add an additional standby OMS later after having run the "Validating your installation and Complete the Setup" steps, the steps listed under "Additional Standby Management Services" can be followed after executing the following additional step:

1. Start the standby Administration Server by running the following command on the first standby Management Service:

   ```
   emctl start oms –admin_only
   ```

2. Export the configuration from the first Management Service on the primary site using:

   ```
   emctl exportconfig oms -dir <location for the export file>
   ```

   After the configuration is exported, do not make any configuration changes to the primary site still the standby management service is configured.

3. Install a Management Agent on the standby host.

4. Perform a software-only install of the Enterprise Manager software using a modified version of "Add Management Service" Deployment Procedure.

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

   Select **Add Management Service** procedure and then click **Create Like**.

   Go to the Procedure Steps tab and select and disable the steps - "Configure Management Service", "Targets Discovery" and "Post Configuration Tasks".

   Save the modified deployment procedure and use it to install the Enterprise Manager software on the standby OMS host.

After the Deployment Procedure completes, delete the file emInstanceMapping.properties from <OMS Oracle Home>/sysman/config on the standby OMS host.

5. Configure the Management Service by running omsca.

   ```
   omsca add –nostart
   ```

   When prompted for Management Repository details, provide the Primary database details.

   When prompted for Administration Server details, provide the standby administration server details.

6. Configure the required plugins by running the following command:

   ```
   pluginca -action deploy –isFirstOMS false -plugins
   <plugin-list> -oracleHome <oms oracle home> -middlewareHome
   <wls middleware home>
   ```

   where plugin-list is the list of plugins returned by the SQL query above and is a comma separate list in the following format:
   `<plugin-id>=<plugin-version>,<plugin-id>=<plugin-version>,…`

   Example

   ```
   "oracle.sysman.empa=12.1.0.1.0,oracle.sysman.mos=12.1.0.1.0,o
   racle.sysman.emas=12.1.0.1.0,oracle.sysman.emfa=12.1.0.1.0,or
   acle.sysman.db=12.1.0.1.0,oracle.sysman.emct=12.1.0.1.0,oracl
   e.sysman.vt=12.1.0.1.0,oracle.sysman.ssa=12.1.0.1.0"
   ```

7. Copy over the configuration exported from the Primary Management Service in step 1 above to the standby Management Service host. Import the exported configuration on the standby Management Service using:

   ```
   emctl importconfig oms -file <full path of the export file>
   ```

   Note this command emits a warning about a failed export and prompts for confirmation to proceed. The warning can be ignored by entering "y" to proceed.

   Note this command will start the Management Service.

8. Stop the Management Service using:

   ```
   emctl stop oms
   ```

9. Refresh the standby domain target from the console. This will present a guided workflow to discover and add the new managed server and associated targets.

10. If you have Single Sign On configured on the primary site, follow the same steps to configure SSO on the standby OMS.

11. If you have Real User Experience Insight, AD4J Manager or BI Publisher configured on the primary site, follow the same steps to configure them on the standby OMS.

21.9.3 Validating Your Installation and Complete the Setup

Update the standby SLB configuration by adding the standby Management Service(s) to the different pools on the SLB. Setup monitors for the new Management Service.

21.9.3.1 Keeping the Standby Site in Sync

After the initial setup of the standby site, the standby site has to be kept in sync with the changes done on primary site. Transactions on the Primary Management

Repository get propagated to the Standby Management Repository automatically through Data Guard but the OMS side changes have to be redone manually on the standby site. The following sections describe this procedure for typical activities.

Applying patches

When patches are applied on the primary site Management Services, they have to be applied on the standby site Management Services too. Note that patches typically update the Oracle Homes (via the opatch apply command) and optionally might require scripts to be run against the Management Repository. On the standby site, it is sufficient to update the Oracle Homes (via the opatch apply command) and skip the running of scripts on the Management Repository because database changes are automatically propagated to the standby site using Data Guard.

Managing Plugins

When new Plugins are deployed on the Primary Site or existing Plugins upgraded or un-deployed on the Primary Site, the following procedures needs to be run on the standby site too to keep the Standby Management Services in sync. Note if the Standby Management Services are not kept in sync, they would fail to start when a switchover or failover operation is attempted.

The procedure below assume that the standby site was setup as per the documented process and the standby management services are currently down and point to the standby repository. The plugin(s) deployment on the Primary site has been completed successfully.

Deploying a New Plugin or Upgrading a Plugin on Standby Site

1. Extract the Plugin archives from the Primary site

   Go to the Self Update Home, click on Plugins, select the required plugin and select export from the Action table menu. Note the EM CLI command from the popup that gets displayed.

   ```
   emcli export_update -id=<update id> -deep -host=<standby OMS
   host> -dir=<directory to export archives> <host credential
   options>
   ```

   Note that an additional option "-deep" is required. This command would create 4 files on the destination directory specified. The filename <version>_OMS_ <platform>_<revision>.zip is the one to be used in next step.

2. Start the Standby Administration Server, if it is down.

   ```
   emctl start oms -admin_only
   ```

3. Install the OMS archive to First Standby OMS Oracle Home

   ```
   pluginia –archives <path to plugin archive>
   ```

4. Configure the Plug-in on First Standby OMS Oracle Home

   ```
   pluginca -action deploy -isFirstOMS true -plugins
   <plugin-list> -oracleHome <oms oracle home> -middlewareHome
   <wls middleware home>
   ```

   where <plugin-list> is the plugin name in the format <plugin-id>=<plugin-version>

5. Repeat steps 3 and 4 for each Standby additional OMS

   ```
   pluginia –archives <path to plugin archive>
   ```

```
pluginca -action deploy -isFirstOMS false -plugins
<plugin-list> -oracleHome <oms oracle home> -middlewareHome
<wls middleware home>
```

This completes the plugin deployment on Standby site.

22

Backing Up Enterprise Manager

As the monitoring and management framework for your ecosystem, an important part of your high availability strategy is to ensure Enterprise Manager is regularly backed up so that it can be restored in the event of failure.

This chapter covers the following topics:

- Backing Up Your Deployment
- Management Repository Backup
- Oracle Management Service Backup
- Management Agent Backup

22.1 Backing Up Your Deployment

Although Enterprise Manager functions as a single entity, technically, it is built on a distributed, multi-tier software architecture composed of the following software components:

- Oracle Management Services (OMS)
- Management Agent
- Management Repository

Each component, being uniquely different in composition and function, requires different approaches to backup and recovery. For this reason, the backup strategies are discussed on a per-tier basis in this chapter. For an overview of Enterprise Manager architecture, refer to the Oracle® Enterprise Manager Cloud Control Basic Installation Guide.

22.2 Management Repository Backup

The Management Repository is the storage location where all the information collected by the Management Agent gets stored. It consists of objects such as database jobs, packages, procedures, views, and tablespaces. Because it is configured in an Oracle Database, the backup and recovery strategies for the Management Repository are essentially the same as those for the Oracle Database. Backup procedures for the database are well established standards and can be implemented using the RMAN backup utility, which can be accessed via the Cloud Control console.

Management Respository Backup

Oracle recommends using High Availability Best Practices for protecting the Management Repository database against unplanned outages. As such, use the following standard database backup strategies.

- Database should be in *archivelog* mode. Not running the repository database in *archivelog* mode leaves the database vulnerable to being in an unrecoverable condition after a media failure.
- Perform regular hot backups with RMAN using the *Recommended Backup Strategy* option via the Cloud Control console. Other utilities such as DataGuard and RAC can also be used as part of a comprehensive strategy to prevent data loss.

Adhering to these strategies will create a full backup and then create incremental backups on each subsequent run. The incremental changes will then be rolled up into the baseline, creating a new full backup baseline.

Using the *Recommended Backup Strategy* also takes advantage of the capabilities of Enterprise Manager to execute the backups: Jobs will be automatically scheduled through the Job sub-system of Enterprise Manager. The history of the backups will then be available for review and the status of the backup will be displayed on the repository database target home page. This backup job along with archiving and flashback technologies will provide a restore point in the event of the loss of any part of the repository. This type of backup, along with archive and online logs, allows the repository to be recovered to the last completed transaction.

You can view when the last repository backup occurred on the Management Services and Repository Overview page under the Repository details section.

A thorough summary of how to configure backups using Enterprise Manager is available in the *Oracle Database 2 Day DBA* guide. For additional information on Database high availability best practices, review the *Oracle Database High Availability Best Practices* documentation.

22.3 Oracle Management Service Backup

The Oracle Management Service (OMS) orchestrates with Management Agents to discover targets, monitor and manage them, and store the collected information in a repository for future reference and analysis. The OMS also renders the Web interface for the Enterprise Manager console.

Backing Up the OMS

The OMS is generally stateless. Some configuration data is stored on the OMS file system.

A snapshot of OMS configuration can be taken using the `emctl exportconfig oms` command.

```
$ <OMS_HOME>/bin/emctl exportconfig oms [-sysman_pwd <sysman password>]
[-dir <backup dir>] Specify directory to store backup file
[-keep_host] Specify this parameter if the OMS was installed using a virtual
hostname (using
ORACLE_HOSTNAME=<virtual_hostname>)
```

Running *exportconfig* captures a snapshot of the OMS at a given point in time, thus allowing you to back up the most recent OMS configuration on a regular basis. *exportconfig* should always be run on the OMS running the WebLogic Admin Server. If required, the most recent snapshot can then be restored on a fresh OMS installation on the same or different host.

Backup strategies for the OMS components are as follows:

- **Software Homes**

  Composed of Fusion Middleware Home, the OMS Oracle Home and the WebTier (OHS) Oracle Home and multiple Management Plug-in Oracle Homes.

  Software Homes change when patches or patchsets are applied or updates are applied through the new Self Update feature. For this reason, filesystem-level backups should be taken after each patch/patchset application or application of updates through Self Update. You should back up the Oracle inventory files along with the Software Homes and save the output of opatch lsinventory –detail to make it easy to determine which patches are applied to the backed up Oracle Homes.

  Note: If you do not have filesystem-level backups, you can also reinstall the software homes using the "Installing Software Only" install method.

  Important: The location of the OMS Oracle Home must be the same for all OMS instances in your Cloud Control deployment.

- **Instance Home**

  The gc_inst directory, composed of WebLogic Server, OMS and web tier configuration files.

  The Instance Home can be backed up using the `emctl exportconfig oms` command.

- **Administration Server**

  The Administration Server operates as the central control entity for the configuration of the entire OMS instance domain. The Administration Server is an integral part of the first OMS installed in your Cloud Control deployment and shares the Software Homes and Instance Home.

  The Administration Server is backed up at the same time as the Instance Home, the `emctl exportconfig oms` command (only run on the first OMS with the Administration Server).

22.4 Management Agent Backup

The Management Agent is an integral software component that is deployed on each monitored host. It is responsible for monitoring all the targets running on those hosts, communicating that information to the middle-tier OMS and managing and maintaining the hosts and its targets.

Backing Up Management Agents

There are no special considerations for backing up Management Agents. As a best practice, reference Management Agent installs should be maintained for different platforms and kept up-to-date in terms of customizations in the emd.properties file and patches applied. Use Deployment options from the Cloud Control console to install and maintain reference Agent installs.

If a Management Agent is lost, it should be reinstalled by cloning from a reference install.

23

Enterprise Manager Outages

Outages can be planned as might be the case when performing upgrades or periodic maintenance, or unplanned as can happen in the event of hardware/software failure, or perhaps some environmental catastrophe. Regardless of the type of outage, you want to ensure that your IT infrastructure can be restored an running as soon as possible.

This chapter covers the following:

Recovery of Failed Enterprise Manager Components

- Repository Recovery
- Recovering the OMS
- Recovering Management Agents
- Recovering from a Simultaneous OMS-Management Repository Failure

Switching or Failing over to Standby Enterprise Manager Configurations

- Switchover
- Failover

23.1 Recovery of Failed Enterprise Manager Components

Recovering Enterprise Manager means restoring any of the three fundamental components of the Enterprise Manager architecture.

- Management Repository
- Management Service
- Management Agent

23.1.1 Repository Recovery

Recovery of the Repository database must be performed using RMAN since Cloud Control will not be available when the repository database is down. There are two recovery cases to consider:

- **Full Recovery**: No special consideration is required for Enterprise Manager.
- **Point-in-Time/Incomplete Recovery**: Recovered repository may be out of sync with Agents because of lost transactions. In this situation, some metrics may show up incorrectly in the Cloud Control console unless the repository is synchronized with the latest state available on the Agents.

A repository resync feature allows you to automate the process of synchronizing the Enterprise Manager repository with the latest state available on the Management Agents.

To resynchronize the repository with the Management Agents, you use Enterprise Manager command-line utility (emctl) `resync repos` command:

```
emctl resync repos -full -name "<descriptive name for the operation>"
```

You must run this command from the OMS Oracle Home AFTER restoring the Management Repository, but BEFORE starting the OMS. After submitting the command, start up all OMS instances and monitor the progress of repository resychronization from the Enterprise Manager console's Repository Resynchronization page, as shown in the following figure.

Figure 23–1 Repository Synchronization Page

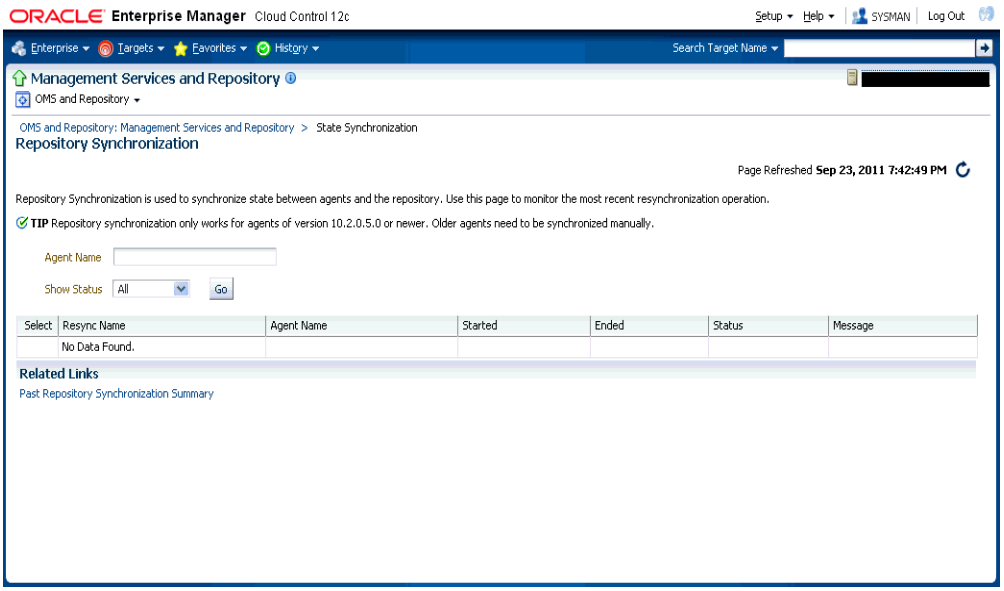

Management Repository recovery is complete when the resynchronization jobs complete on all Management Agents.

Oracle strongly recommends that the Management Repository database be run in *archivelog* mode so that in case of failure, the database can be recovered to the latest transaction. If the database cannot be recovered to the last transaction, *Repository Synchronization* can be used to restore monitoring capabilities for targets that existed when the last backup was taken. Actions taken after the backup will not be recovered automatically. Some examples of actions that will not be recovered automatically by *Repository Synchronization* are:

- Incident Rules
- Preferred Credentials
- Groups, Services, Systems
- Jobs/Deployment Procedures
- Custom Reports
- New Agents

23.1.2 Recovery Scenarios

A prerequisite for repository (or any database) recovery is to have a valid, consistent backup of the repository. Using Enterprise Manager to automate the backup process ensures regular, up-to-date backups are always available if repository recovery is ever required. Recovery Manager (RMAN) is a utility that backs up, restores, and recovers Oracle Databases. The RMAN recovery job syntax should be saved to a safe location. This allows you to perform a complete recovery of the Enterprise Manager repository database. In its simplest form, the syntax appears as follows:

```
run {
restore database;
recover database;
}
```

Actual syntax will vary in length and complexity depending on your environment. For more information on extracting syntax from an RMAN backup and recovery job, or using RMAN in general, see the *Oracle Database Backup and Recovery Advanced User's Guide*.

The following scenarios illustrate various repository recovery situations along with the recovery steps.

23.1.2.1 Full Recovery on the Same Host

Repository database is running in *archivelog* mode. Recent backup, archive log files and redo logs are available. The repository database disk crashes. All datafiles and control files are lost.

Resolution:

1. Stop all OMS instances using `emctl stop oms`.
2. Recover the database using RMAN
3. Bring the site up using the command `emctl start oms` on all OMS instances.
4. Verify that the site is fully operational.

23.1.2.2 Incomplete Recovery on the Same Host

Repository database is running in *noarchivelog* mode. Full offline backup is available. The repository database disk crashes. All datafiles and control files are lost.

Resolution:

1. Stop the OMS instances using `emctl stop oms`.
2. Recover the database using RMAN.
3. Initiate Repository Resync using `emctl resync repos -full -name` "<resync name>" from one of the OMS Oracle Home.
4. Start the OMS(instances using `emctl start oms`.
5. Log into Cloud Control. Navigate to the Management Services and Repository Overview page. Click **Repository Synchronization** under **Related Links**. Monitor the status of resync jobs. Resubmit failed jobs, if any, after fixing the error.
6. Verify that the site is fully operational.

23.1.2.3 Full Recovery on a Different Host

The Management Repository database is running on host "A" in *archivelog* mode. Recent backup, archive log files and redo logs are available. The repository database crashes. All datafiles and control files are lost.

Resolution:

1. Stop the OMS instances using the command `emctl stop oms`.
2. Recover the database using RMAN on a different host (host "B").
3. Correct the connect descriptor for the repository by running the following command on each OMS.

   ```
   $emctl config oms –store_repos_details -repos_conndesc <connect descriptor>
   -repos_user sysman
   ```

4. Start the OMS instances using the command `emctl start oms`.
5. Relocate the Management Repository database target to the Agent running on host "B" by running the following command from the OMS:

   ```
   $emctl config repos -host <hostB> -oh <OH of repository on hostB>  -conn_desc
   "<TNS connect descriptor>"
   ```

   Note: This command can only be used to relocate the repository database under the following conditions:

   - An Agent is already running on this machine.
   - No database on host "B" has been discovered.

6. Change the monitoring configuration for the OMS and Repository target: by running the following command from the OMS:

   ```
   $emctl config emrep -conn_desc "<TNS connect descriptor>"
   ```

7. Verify that the site is fully operational.

23.1.2.4 Incomplete Recovery on a Different Host

The Management Repository database is running on host "A" in *noarchivelog* mode. Full offline backup is available. Host "A" is lost due to hardware failure. All datafiles and control files are lost.

Resolution:

1. Stop the OMS instances using `emctl stop oms`.
2. Recover the database using RMAN on a different host (host "B").
3. Correct the connect descriptor for the repository in credential store.

   ```
   $emctl config oms –store_repos_details -repos_conndesc <connect descriptor>
   -repos_user sysman
   ```

4. Initiate Repository Resync:

   ```
   $emctl resync repos -full -name "<resync name>"
   ```

   from one of the OMS Oracle Homes.

5. Start the OMS using the command `emctl start oms`.

6. Run the command to relocate the repository database target to the Management Agent running on host "B":

   ```
   $emctl config repos -agent <agent on host B> -host <hostB>
   -oh <OH of repository on hostB> -conn_desc "<TNS connect
   descriptor>"
   ```

7. Run the command to change monitoring configuration for the OMS and Repository target:

   ```
   emctl config emrep -conn_desc "<TNS connect descriptor>"
   ```

8. Log in to Cloud Control. Navigate to **Management Services and Repository Overview** page.

9. Choose on **Repository Synchronization** under **Related Links**. Monitor the status of resync jobs. Resubmit failed jobs, if any, after fixing the error mentioned.

10. Verify that the site is fully operational.

23.1.3 Recovering the OMS

If an Oracle Management Service instance is lost, recovering it essentially consists of two steps: Recovering the Software Homes, then configuring the Instance Home.

When restoring on the same host, the software homes can be restored from filesystem backup. In case a backup does not exist, or if installing to a different host, the Software Homes can be reconstructed using the “Install Software Only" option from the Cloud Control software distribution. Care should be taken to select and install all Management Plug-ins that existed in your environment prior to crash.

The following SQL command can be run against the repository database as the “sysman” user to obtain the list of plug-ins already deployed:

```
SELECT epv.display_name, epv.plugin_id, epv.version FROM em_plugin_version epv,
em_current_deployed_plugin ecp WHERE epv.plugin_type NOT IN ( 'BUILT_IN_TARGET_
TYPE' , 'INSTALL_HOME') AND ecp.dest_type='2' AND epv.plugin_version_id =
ecp.plugin_version_id;
```

Note that some plug-ins might have not shipped with Cloud Control and might not be present in the install media. Such plug-ins should be downloaded from OTN and their location passed to the Oracle Installer. Choose all plug-ins returned by the SQL query above in the plug-ins page of the Installer. Recovery will fail if all the required plug-ins are not selected.

After running the installer in software only mode, all patches that were installed prior to the crash must be re-applied. Assuming the Management Repository is intact, the post scripts that run SQLs against the repository can be skipped as the repository already has those patches applied.

As stated earlier, the location of the OMS Oracle Home is fixed and cannot be changed. Hence, ensure that the OMS Oracle Home is restored in the same location that was used previously.

Once the Software Homes are recovered, the instance home can be reconstructed using the omsca command in recovery mode:

```
omsca recover –as –ms -nostart –backup_file <exportconfig file>
```

Use the export file generated by the `emctl exportconfig` command shown in the previous section.

23.1.4 OMS Recovery Scenarios

The following scenarios illustrate various OMS recovery situations along with the recovery steps.

Important: A prerequisite for OMS recovery is to have recent, valid OMS configuration backups available. Oracle recommends that you back up the OMS using the `emctl exportconfig oms` command whenever an OMS configuration change is made. This command must be run on the primary OMS running the WebLogic AdminServer.

Alternatively, you can run this command on a regular basis using the Enterprise Manager Job system.

Each of the following scenarios cover the recovery of the Software homes using either a filesystem backup (when available and only when recovering to the same host) or using the Software only option from the installer. In either case, the best practice is to recover the instance home (gc_inst) using the omsca recover command, rather than from a filesystem backup. This guarantees that the instance home is valid and up to date.

23.1.4.1 Single OMS, No Server Load Balancer (SLB), OMS Restored on the same Host

Site hosts a single OMS. No SLB is present. The OMS configuration was backed up using the `emctl exportconfig oms` command on the primary OMS running the AdminServer. The OMS Oracle Home is lost.

Resolution:

1. Perform cleanup on failed OMS host.

   Make sure there are no processes still running from the Middleware home using a command similar to the following:

   ```
   ps -ef | grep -i -P "(Middleware|gc_inst)" | grep -v grep | awk '{print $2}' |
   xargs kill -9
   ```

   Note: Change *Middleware*|*gc_inst* to strings that match your own middleware and instance homes.

   If recovering the software homes using the software only install method, first de-install the existing Oracle Homes using the Cloud Control software distribution installer. This is required even if the software homes are no longer available as it is necessary to remove any record of the lost Oracle Homes from the Oracle inventory.

   If they exist, remove the 'Middleware' and 'gc_inst' directories.

2. Ensure that software library locations are still accessible.

3. Restore the Software Homes.

   If restoring from a filesystem backup, delete the file *OMS_HOME*/sysman/config/emInstanceMapping.properties and any gc_inst directory that may have been restored, if they exist.

Alternatively, if a backup does not exist, use the software only install method to reconstruct the software homes:

1. Select the 'Install Software Only' option from the 'Install Types' step page within the Cloud Control software installer.
2. Ensure all previously deployed plug-ins are selected on the 'Select Plug-ins' step page.

   It is possible to determine which plug-ins were deployed previously by running the following SQL against the repository database:

   ```
   SELECT epv.display_name, epv.plugin_id, epv.version FROM em_plugin_version
   epv, em_current_deployed_plugin ecp WHERE epv.plugin_type NOT IN ( 'BUILT_
   IN_TARGET_TYPE' , 'INSTALL_HOME') AND ecp.dest_type='2' AND epv.plugin_
   version_id = ecp.plugin_version_id;
   ```

   Note: At the end of the Software only installation, do NOT run *ConfigureGC.pl* when told to do so by the installer. This step should only be performed as part of a fresh install, not as part of a recovery operation.

3. Apply any patches that were previously applied to the OMS software homes.

4. Run omsca in recovery mode specifying the export file taken earlier to configure the OMS:

   ```
   <OMS_HOME>/bin/omsca recover –as –ms –nostart –backup_file <exportconfig file>
   ```

   Note: The -backup_file to be passed must be the latest file generated from emctl exportconfig oms command.

5. Start the OMS.

   ```
   OMS_HOME/bin/emctl start oms
   ```

6. Recover the Agent (if necessary).

   If the Management Agent Software Home was recovered along with the OMS Software Homes (as is likely in a single OMS install recovery where the Management Agent and agent_inst directories are commonly under the Middleware home), the Management Agent instance directory should be recreated to ensure consistency between the Management Agent and OMS.

   1. Remove the agent_inst directory if it was restored from backup
   2. Use agentDeploy.sh to configure the agent:

      ```
      <AGENT_HOME>/core/12.1.0.0.0/sysman/install/agentDeploy.sh AGENT_BASE_
      DIR=<AGENT_BASE_DIR> AGENT_INSTANCE_HOME=<AGENT_INSTANCE_HOME> ORACLE_
      HOSTNAME=<AGENT_HOSTNAME> AGENT_PORT=<AGENT_PORT> -configOnly OMS_HOST=<oms
      host> EM_UPLOAD_PORT=<OMS_UPLOAD_PORT> AGENT_REGISTRATION_PASSWORD=<REG_
      PASSWORD>
      ```

   3. The OMS automatically blocks the Management Agent. Resync the Management Agent from the Management Agent homepage.

If the Management Agent software home was not recovered along with the OMS but the Agent still needs to be recovered, follow the instructions in section *Agent Reinstall Using the Same Port*.

Note: This is only likely to be needed in the case where a filesystem recovery has been performed that did not include a backup of the Agent software homes. If the OMS software homes were recovered using the Software only install method, this step will not be required because a Software only install installs an Agent software home under the Middleware home.

---

7. Verify that the site is fully operational.

23.1.4.2 Single OMS, No SLB, OMS Restored on a Different Host

Site hosts a single OMS. The OMS is running on host "A." No SLB is present. The OMS configuration was backed up using the `emctl exportconfig oms` command. Host "A" is lost.

Resolution:

1. Ensure that software library locations are accessible from “Host B”.
2. Restore the software homes on “Host B”.

   Oracle does not support restoring OMS Oracle Homes from filesystem backup across different hosts. Use the software-only install method to reconstruct the software homes:

   1. Select the ‘Install Software Only’ option from the ‘Install Types’ step page within the Cloud Control software installer.
   2. Ensure all previously deployed plug-ins are selected on the ‘Select Plug-ins’ step page.

      It is possible to determine which plug-ins were deployed previously by running the following SQL against the repository database:

      ```
      SELECT epv.display_name, epv.plugin_id, epv.version FROM em_plugin_version
      epv, em_current_deployed_plugin ecp WHERE epv.plugin_type NOT IN ( 'BUILT_
      IN_TARGET_TYPE' , 'INSTALL_HOME') AND ecp.dest_type='2' AND epv.plugin_
      version_id = ecp.plugin_version_id;
      ```

      Note: At the end of the Software only installation, do NOT run *ConfigureGC.pl* when told to do so by the installer. This step should only be performed as part of a fresh install, not as part of a recovery operation.

      ---

   3. Apply any patches that were previously applied to the OMS software homes.
3. Run omsca in recovery mode specifying the export file taken earlier to configure the OMS:

   ```
   <OMS_HOME>/bin/omsca recover -as -ms -nostart -backup_file <exportconfig file>
   ```

   Note: The -backup_file to be passed must be the latest file generated from emctl exportconfig oms command.

   ---

4. Start the OMS.

   ```
   <OMS_HOME>/bin/emctl start oms
   ```

   An agent is installed as part of the Software only install and needs to be configured using the agentDeploy.sh command:

5. Configure the Agent.

   ```
   <AGENT_HOME>/core/12.1.0.0.0/sysman/install/agentDeploy.sh AGENT_BASE_
   DIR=<AGENT_BASE_DIR> AGENT_INSTANCE_HOME=<AGENT_INSTANCE_HOME> ORACLE_
   HOSTNAME=<AGENT_HOSTNAME> AGENT_PORT=<AGENT_PORT> -configOnly OMS_HOST=<oms
   host> EM_UPLOAD_PORT=<OMS_UPLOAD_PORT> AGENT_REGISTRATION_PASSWORD=<REG_
   PASSWORD>
   ```

   The OMS automatically blocks the Management Agent. Resync the Management Agent from the Management Agent homepage

6. Relocate the oracle_emrep target to the Management Agent of the new OMS host using the following commands:

   ```
   <OMS_HOME>/bin/emcli login –username=sysman
   <OMS_HOME>/bin/emcli sync
   <OMS_HOME>/bin/emctl config emrep -agent <agent on host "B", e.g
   myNewOMSHost.example.com:3872>
   ```

7. In the Cloud Control console, locate the 'WebLogic Domain' target for the Cloud Control Domain. Go to 'Monitoring Credentials' and update the adminserver host to host B. Then do a Refresh Weblogic Domain to reconfigure the domain with new hosts.

8. Locate duplicate targets from the Management Services and Repository Overview page of the Enterprise Manager console. Click the Duplicate Targets link to access the Duplicate Targets page. To resolve duplicate target errors, the duplicate target must be renamed on the conflicting Agent. Relocate duplicate targets from Agent "A" to Agent "B".

9. Change the OMS to which all Management Agents point and then resecure all Agents.

   Because the new machine is using a different hostname from the one originally hosting the OMS, all Agents in your monitored environment must be told where to find the new OMS. On each Management Agent, run the following command:

   ```
   <AGENT_INST_DIR>/bin/emctl secure agent -emdWalletSrcUrl "http://hostB:<http_
   port>/em"
   ```

10. Assuming the original OMS host is no longer in use, remove the Host target (including all remaining monitored targets) from Cloud Control by selecting the host on the Targets > Hosts page and clicking ‘Remove’. You will be presented with an error that informs you to remove all monitored targets first. Remove those targets then repeat the step to remove the Host target successfully.

11. Verify that the site is fully operational.

23.1.4.3 Single OMS, No SLB, OMS Restored on a Different Host using the Original Hostname

Site hosts a single OMS. The OMS is running on host "A." No SLB is present. The OMS configuration was backed up using the `emctl exportconfig oms` command. Host "A" is lost. Recovery is to be performed on “Host B” but retaining the use of “Hostname A”.

Resolution:

1. Ensure that loader receive directory and software library locations are accessible from Host "B".

   Oracle does not support restoring OMS Oracle Homes from filesystem backup across different hosts. Use the software-only install method to reconstruct the software homes:

   1. Select the 'Install Software Only' option from the 'Install Types' step page within the Cloud Control software installer.
   2. Ensure all previously deployed plug-ins are selected on the 'Select Plug-ins' step page.

      It is possible to determine which plug-ins were deployed previously by running the following SQL against the Management Repository database:

      ```
      SELECT epv.display_name, epv.plugin_id, epv.version FROM em_plugin_version
      epv, em_current_deployed_plugin ecp WHERE epv.plugin_type NOT IN ( 'BUILT_
      IN_TARGET_TYPE' , 'INSTALL_HOME') AND ecp.dest_type='2' AND epv.plugin_
      version_id = ecp.plugin_version_id;
      ```

      Note: At the end of the Software only installation, do NOT run *ConfigureGC.pl* when told to do so by the installer. This step should only be performed as part of a fresh install, not as part of a recovery operation.

   3. Apply any patches that were previously applied to the OMS software homes.

2. Modify the network configuration such that "Host B" also responds to hostname of "Host A". Specific instructions on how to configure this are beyond the scope of this document. However, some general configuration suggestions are:

   Modify your DNS server such that both "Hostname B" and "Hostname A" network addresses resolve to the physical IP of "Host B".

   Multi-home "Host B". Configure an additional IP on "Host B" for the IP address that "Hostname A" resolves to. For example, on "Host B" run the following commands:

   ```
   ifconfig eth0:1 <IP assigned to "Hostname A"> netmask <netmask>
   /sbin/arping -q -U -c 3 -I eth0 <IP of HostA>
   ```

3. Run omsca in recovery mode specifying the export file taken earlier to configure the OMS:

   ```
   <OMS_HOME>/bin/omsca recover -as -ms -nostart -backup_file <exportconfig file>
   -AS_HOST <hostA> -EM_INSTANCE_HOST <hostA>
   ```

   Note: The -backup_file to be passed must be the latest file generated from emctl exportconfig oms command.

4. Start the OMS

   ```
   <OMS_HOME>/bin/emctl start oms
   ```

5. Configure the agent.

An agent is installed as part of the Software only install and needs to be configured using the agentDeploy.sh command:

```
<AGENT_HOME>/core/12.1.0.0.0/sysman/install/agentDeploy.sh AGENT_BASE_
DIR=<AGENT_BASE_DIR> AGENT_INSTANCE_HOME=<AGENT_INSTANCE_HOME> ORACLE_
HOSTNAME=<AGENT_HOSTNAME> AGENT_PORT=<AGENT_PORT> -configOnly OMS_HOST=<oms
host> EM_UPLOAD_PORT=<OMS_UPLOAD_PORT> AGENT_REGISTRATION_PASSWORD=<REG_
PASSWORD>
```

6. The OMS automatically blocks the Management Agent. Resync the Management Agent from the Management Agent homepage.

   Run the command to relocate Management Services and Management Repository target to Management Agent "B":

   ```
   emctl config emrep -agent <agent on host B>
   ```

7. In the Cloud Control console, locate the 'WebLogic Domain' target for the Cloud Control Domain. Go to 'Monitoring Credentials' and update the adminserver host to host B. Then do a Refresh Weblogic Domain to reconfigure the domain with new hosts.
8. Locate duplicate targets from the Management Services and Repository Overview page of the Enterprise Manager console. Click the Duplicate Targets link to access the Duplicate Targets page. To resolve duplicate target errors, the duplicate target must be renamed on the conflicting Management Agent. Relocate duplicate targets from Management Agent "A" to Management Agent "B".
9. Verify that the site is fully operational.

23.1.4.4 Multiple OMS, Server Load Balancer, Primary OMS Recovered on the Same Host

Site hosts multiple OMS instances. All OMS instances are fronted by a Server Load Balancer. OMS configuration backed up using the `emctl exportconfig oms` command on the primary OMS running the WebLogic AdminServer. The primary OMS is lost.

Resolution:

1. Perform a cleanup on the failed OMS host.

   Make sure there are no processes still running from the Middleware home using a command similar to the following:

   ```
   ps -ef | grep -i -P "(Middleware|gc_inst)" | grep -v grep | awk '{print $2}' |
   xargs kill -9
   ```

   Note: Change *Middleware|gc_inst* to strings that match your own middleware and instance homes.

   If recovering the software homes using the software only install method, first de-install the existing Oracle Homes using the Cloud Control software distribution installer. This is required even if the software homes are no longer available as it is necessary to remove any record of the lost Oracle Homes from the Oracle inventory.

   If they exist, remove the 'Middleware' and 'gc_inst' directories.

2. Ensure that software library locations are still accessible.

3. Restore the software homes.

   If restoring from a filesystem backup, delete the file <OMS_HOME>/sysman/config/emInstanceMapping.properties and any gc_inst directory that may have been restored, if they exist.

   Alternatively, if a backup does not exist, use the software only install method to reconstruct the software homes:

   1. Select the 'Install Software Only' option from the 'Install Types' step page within the Cloud Control software installer.
   2. Ensure all previously deployed plug-ins are selected on the 'Select Plug-ins' step page.

      It is possible to determine which plugins were deployed previously by running the following SQL against the Management Repository database:

      ```
      SELECT epv.display_name, epv.plugin_id, epv.version FROM em_plugin_version
      epv, em_current_deployed_plugin ecp WHERE epv.plugin_type NOT IN ( 'BUILT_
      IN_TARGET_TYPE' , 'INSTALL_HOME') AND ecp.dest_type='2' AND epv.plugin_
      version_id = ecp.plugin_version_id;
      ```

      Note: At the end of the Software only installation, do NOT run *ConfigureGC.pl* when told to do so by the installer. This step should only be performed as part of a fresh install, not as part of a recovery operation.

   3. Apply any patches that were previously applied to the OMS software homes.

4. Run omsca in recovery mode specifying the export file taken earlier to configure the OMS:

   ```
   <OMS_HOME>/bin/omsca recover –as –ms –nostart –backup_file <exportconfig file>
   ```

   Note: The -backup_file to be passed must be the latest file generated from emctl exportconfig oms command.

5. Start the OMS.

   ```
   <OMS_HOME>/bin/emctl start oms
   ```

6. Recover the Management Agent.

   If the Management Agent software home was recovered along with the OMS software homes (as is likely in a Primary OMS install recovery where the agent and agent_inst directories are commonly under the Middleware home), the Management Agent instance directory should be recreated to ensure consistency between the Management Agent and OMS.

   1. Remove the agent_inst directory if it was restored from backup.
   2. Use agentDeploy.sh to configure the Management Agent:

      ```
      <AGENT_HOME>/core/12.1.0.0.0/sysman/install/agentDeploy.sh AGENT_BASE_
      DIR=<AGENT_BASE_DIR> AGENT_INSTANCE_HOME=<AGENT_INSTANCE_HOME> ORACLE_
      HOSTNAME=<AGENT_HOSTNAME> AGENT_PORT=<AGENT_PORT> -configOnly OMS_HOST=<oms
      host> EM_UPLOAD_PORT=<OMS_UPLOAD_PORT> AGENT_REGISTRATION_PASSWORD=<REG_
      PASSWORD>
      ```

3. The OMS automatically blocks the Management Agent. Resync the Management Agent from the Management Agent homepage.

If the Management Agent software home was not recovered along with the OMS but the Management Agent still needs to be recovered, follow the instructions in section *Agent Reinstall Using the Same Port.*

Note: This is only likely to be needed in the case where a filesystem recovery has been performed that did not include a backup of the Management Agent software homes. If the OMS software homes were recovered using the Software only install method, this step will not be required because a Software only install installs an Management Agent software home under the Middleware home.

---

7. Re-enroll the additional OMS, if any, with the recovered Administration Server by running emctl enroll oms on each additional OMS.
8. Verify that the site is fully operational.

23.1.4.5 Multiple OMS, Server Load Balancer configured, Primary OMS Recovered on a Different Host

Site hosts multiple OMS instances. OMS instances fronted by a Server Load Balancer. OMS Configuration backed up using emctl exportconfig oms command. Primary OMS on host "A" is lost and needs to be recovered on Host "B".

1. If necessary, perform cleanup on failed OMS host.

   Make sure there are no processes still running from the Middleware home using a command similar to the following:

   ```
   ps -ef | grep -i -P "(Middleware|gc_inst)" | grep -v grep | awk '{print $2}' |
   xargs kill -9
   ```

2. Ensure that software library locations are accessible from “Host B”.
3. Restore the software homes on “Host B”.

   Oracle does not support restoring OMS Oracle Homes from filesystem backup across different hosts. Use the software-only install method to reconstruct the software homes:

   1. Select the ‘Install Software Only’ option from the ‘Install Types’ step page within the Cloud Control software installer.
   2. Ensure all previously deployed plug-ins are selected on the ‘Select Plug-ins’ step page.

      It is possible to determine which plugins were deployed previously by running the following SQL against the Management Repository database:

      ```
      SELECT epv.display_name, epv.plugin_id, epv.version FROM em_plugin_version
      epv, em_current_deployed_plugin ecp WHERE epv.plugin_type NOT IN ( 'BUILT_
      IN_TARGET_TYPE' , 'INSTALL_HOME') AND ecp.dest_type='2' AND epv.plugin_
      version_id = ecp.plugin_version_id;
      ```

> **Note:** At the end of the Software only installation, do NOT run *ConfigureGC.pl* when told to do so by the installer. This step should only be performed as part of a fresh install, not as part of a recovery operation.

3. Apply any patches that were previously applied to the OMS software homes.

4. Run omsca in recovery mode specifying the export file taken earlier to configure the OMS:

```
<OMS_HOME>/bin/omsca recover -as -ms -nostart -backup_file <exportconfig file>
```

> **Note:** The -backup_file to be passed must be the latest file generated from emctl exportconfig oms command.

5. Start the OMS.

```
<OMS_HOME>/bin/emctl start oms
```

6. Configure the Management Agent.

   An agent is installed as part of the Software only install and needs to be configured using the agentDeploy.sh command:

```
<AGENT_HOME>/core/12.1.0.0.0/sysman/install/agentDeploy.sh AGENT_BASE_
DIR=<AGENT_BASE_DIR> AGENT_INSTANCE_HOME=<AGENT_INSTANCE_HOME> ORACLE_
HOSTNAME=<AGENT_HOSTNAME> AGENT_PORT=<AGENT_PORT> -configOnly OMS_HOST=<oms
host> EM_UPLOAD_PORT=<OMS_UPLOAD_PORT> AGENT_REGISTRATION_PASSWORD=<REG_
PASSWORD>
```

   The OMS automatically blocks the Management Agent. Resync the Management Agent from the Management Agent homepage

7. Add the new OMS to the SLB virtual server pools and remove the old OMS.

8. Relocate the oracle_emrep target to the Management Agent of the new OMS host using the following commands:

```
<OMS_HOME>/bin/emcli sync
<OMS_HOME>/bin/emctl config emrep -agent <agent on host "B", e.g
myNewOMSHost.example.com:3872>
```

9. In the Cloud Control console, locate the 'WebLogic Domain' target for the Cloud Control Domain. Go to 'Monitoring Credentials' and update the adminserver host to host B. Then do a Refresh Weblogic Domain to reconfigure the domain with new hosts.

10. Locate duplicate targets from the Management Services and Repository Overview page of the Enterprise Manager console. Click the Duplicate Targets link to access the Duplicate Targets page. To resolve duplicate target errors, the duplicate target must be renamed on the conflicting Management Agent. Relocate duplicate targets from Management Agent "A" to Management Agent "B".

11. Assuming the original OMS host is no longer in use, remove the Host target (including all remaining monitored targets) from Cloud Control by selecting the host on the Targets > Hosts page and clicking ‘Remove’. You will be presented with an error that informs you to remove all monitored targets first. Remove those targets then repeat the step to remove the Host target successfully.

12. Verify that the site is fully operational.

23.1.4.6 Multiple OMS, SLB configured, additional OMS recovered on same or different host

Multiple OMS site where the OMS instances are fronted by an SLB. OMS configuration backed up using the `emctl exportconfig oms` command on the first OMS. Additional OMS is lost and needs to be recovered on the same or a different host.

1. If recovering to the same host, ensure cleanup of the failed OMS has been performed:

   Make sure there are no processes still running from the Middleware home using a command similar to the following:

   ```
   ps -ef | grep -i -P "(Middleware|gc_inst)" | grep -v grep | awk '{print $2}' |
   xargs kill -9
   ```

   First de-install the existing Oracle Homes using the Cloud Control software distribution installer. This is required even if the software homes are no longer available as it is necessary to remove any record of the lost Oracle Homes from the Oracle inventory.

   If they exist, remove the 'Middleware' and 'gc_inst' directories.

2. Ensure that shared software library locations are accessible.
3. Install an Management Agent on the required host (same or different as the case may be).
4. Use the Additional OMS deployment procedure to configure a new additional OMS.
5. Verify that the site is fully operational.

23.1.5 Recovering Management Agents

If an Management Agent is lost, it should be reinstalled by cloning from a reference install. Cloning from a reference install is often the fastest way to recover an Management Agent install as it is not necessary to track and reapply customizations and patches. Care should be taken to reinstall the Management Agent using the same port. Using the Enterprise Manager's Management Agent Resynchronization feature, a reinstalled Management Agent can be reconfigured using target information present in the Management Repository. When the Management Agent is reinstalled using the same port, the OMS detects that it has been re-installed and blocks it temporarily to prevent the auto-discovered targets in the re-installed Management Agent from overwriting previous customizations.

> **Blocked Management Agents:** A This is a condition in which the OMS rejects all heartbeat or upload requests from the blocked Management Agent. Hence, a blocked Agent will not be able to upload any alerts or metric data to the OMS. However, blocked Management Agents continue to collect monitoring data.

The Management Agent can be resynchronized and unblocked from the Management Agent homepage by clicking on the **Resynchronize Agent** button. Resynchronization pushes all targets from the Management Repository to the Management Agent and then unblocks the Agent.

23.1.6 Management Agent Recovery Scenarios

The following scenarios illustrate various Management Agent recovery situations along with the recovery steps. The Management Agent resynchronization feature requires that a reinstalled Management Agent use the same port and location as the previous Management Agent that crashed.

23.1.6.1 Management Agent Reinstall Using the Same Port

A Management Agent is monitoring multiple targets. The Agent installation is lost.

1. De-install the Agent Oracle Home using the Oracle Universal Installer.

   Note: This step is necessary in order to clean up the inventory.

2. Install a new Management Agent or use the Management Agent clone option to reinstall the Management Agent though Enterprise Manager. Specify the same port that was used by the crashed Agent. The location of the install must be same as the previous install.

   The OMS detects that the Management Agent has been re-installed and blocks the Management Agent.

3. Initiate Management Agent Resynchronization from the Management Agent homepage.

   All targets in the Management Repository are pushed to the new Management Agent. The Agent is instructed to clear backlogged files and then do a clearstate. The Agent is then unblocked.

4. Reconfigure User-defined Metrics if the location of User-defined Metric scripts have changed.

5. Verify that the Management Agent is operational and all target configurations have been restored using the following emctl commands:

   ```
   emctl status agent
   emctl upload agent
   ```

   There should be no errors and no XML files in the backlog.

23.1.6.2 Management Agent Restore from Filesystem Backup

A single Management Agent is monitoring multiple targets. File system backup for the Agent Oracle Home exists. The Agent install is lost.

1. Restore the Management Agent from the filesystem backup then start the Management Agent.

   The OMS detects that the Management Agent has been restored from backup and blocks the Management Agent.

2. Initiate Management Agent Resynchronization from the Management Agent homepage.

   All targets in the Management Repository are pushed to the new Management Agent. The Agent is instructed to clear backlogged files and performs a clearstate. The Management Agent is unblocked.

3. Verify that the Management Agent is functional and all target configurations have been restored using the following emctl commands:

```
emctl status agent
emctl upload agent
```

There should be no errors and no XML files in the backlog.

23.2 Recovering from a Simultaneous OMS-Management Repository Failure

When both OMS and Management Repository fail simultaneously, the recovery situation becomes more complex depending upon factors such as whether the OMS and Management Repository recovery has to be performed on the same or different host, or whether there are multiple OMS instances fronted by an SLB. In general, the order of recovery for this type of compound failure should be Management Repository first, followed by OMS instances following the steps outlined in the appropriate recovery scenarios discussed earlier. The following scenarios illustrate two OMS-Management Repository failures and the requisite recovery steps.

23.2.1 Collapsed Configuration: Incomplete Management Repository Recovery, Primary OMS on the Same Host

Management Repository and the primary OMS are installed on same host (host "A"). The Management Repository database is running in noarchivelog mode. Full cold backup is available. A recent OMS backup file exists (emctl exportconfig oms). The Management Repository, OMS and the Management Agent crash.

1. Follow the Management Repository recovery procedure shown in Incomplete Recovery on the Same Host with the following exception:

   Since the OMS OracleHome is not available and Management Repository resynchronization has to be initiated before starting an OMS against the restored Management Repository, submit "resync" via the following PL/SQL block. Log into the Management Repository as SYSMAN using SQLplus and run:

   ```
   begin emd_maintenance.full_repository_resync('<resync name>'); end;
   ```

2. Follow the OMS recovery procedure shown in Section 23.1.4.1, "Single OMS, No Server Load Balancer (SLB), OMS Restored on the same Host."
3. Verify that the site is fully operational.

23.2.2 Distributed Configuration: Incomplete Management Repository Recovery, Primary OMS and additional OMS on Different Hosts, SLB Configured

The Management Repository, primary OMS, and additional OMS all reside on the different hosts. The Management Repository database was running in noarchivelog mode. OMS backup file from a recent backup exists (emctl exportconfig oms). Full cold backup of the database exists. All three hosts are lost.

1. Follow the Management Repository recovery procedure shown in Section 23.1.2.2, "Incomplete Recovery on the Same Host." with the following exception:

   Since OMS Oracle Home is not yet available and Management Repository resync has to be initiated before starting an OMS against the restored Management Repository, submit resync via the following PL/SQL block. Log into the Management Repository as SYSMAN using SQLplus and run the following:

   ```
   begin emd_maintenance.full_repository_resync('resync name'); end;
   ```

2. Follow the OMS recovery procedure shown in Section 23.1.4.5, "Multiple OMS, Server Load Balancer configured, Primary OMS Recovered on a Different Host" with the following exception:

   Override the Management Repository connect description present in the backup file by passing the additional omsca parameter:

   ```
   -REPOS_CONN_STR <restored repos descriptor>
   ```

   This needs to be added along with other parameters listed in Section 23.1.4.5, "Multiple OMS, Server Load Balancer configured, Primary OMS Recovered on a Different Host."

3. Follow the OMS recovery procedure shown in Section 23.1.4.6, "Multiple OMS, SLB configured, additional OMS recovered on same or different host."

4. Verify that the site is fully operational.

23.3 Switching Over or Failing Over to Standby Enterprise Manager Configurations

23.3.1 Switchover

Switchover is a planned activity where operations are transferred from the Primary site to a Standby site. This is usually done for testing and validation of Disaster Recovery (DR) scenarios and for planned maintenance activities on the primary infrastructure.

Switchover to the Passive OMS in a Level 2 Active/Passive Configuration

Switchover follows the same steps as Failover, See Section 'Failover to the Passive OMS in a Level 2 Active/Passive Configuration'

Switchover to the Standby Site in a Level 4 MAA Configuration

This section describes the steps to switchover to the standby site. The same procedure is applied to switchover in either direction.

Enterprise Manager Console cannot be used to perform switchover of the Management Repository database. Use the Data Guard Broker command line tool DGMGRL instead.

1. Prepare the Standby Database

   Verify that recovery is up-to-date. Using the Enterprise Manager Console, you can view the value of the ApplyLag column for the standby database in the Standby Databases section of the Data Guard Overview Page.

2. Shut down the Primary Enterprise Manager Application Tier.

   Shutdown all the Management Service instances in the primary site by running the following command on each Management Service:

   ```
   emctl stop oms -all
   ```

3. Verify Software Library Availability

   Ensure all files from the primary site are available on the standby site.

4. Switch over to the Standby Database

Use DGMGRL to perform a switchover to the standby database. The command can be run on the primary site or the standby site. The switchover command verifies the states of the primary database and the standby database, affects switchover of roles, restarts the old primary database, and sets it up as the new standby database.

```
SWITCHOVER TO <standby database name>;
```

Verify the post switchover states. To monitor a standby database completely, the user monitoring the database must have SYSDBA privileges. This privilege is required because the standby database is in a mounted-only state. A best practice is to ensure that the users monitoring the primary and standby databases have SYSDBA privileges for both databases.

```
SHOW CONFIGURATION;
SHOW DATABASE <primary database name>;
SHOW DATABASE <standby database name>;
```

5. Start the Admin Server if it is not already running.

   ```
   emctl start oms -admin_only
   ```

6. Make the standby Management Services point to the Standby Database which is now the new Primary by running the following on each standby Management Service.

   ```
   emctl config oms -store_repos_details -repos_conndesc
   <connect descriptor of new primary database> -repos_user
   sysman
   ```

7. Startup the Enterprise Manager Application Tier

   Startup all the Management Services on the standby site:

   ```
   emctl start oms
   ```

8. Relocate Management Services and Management Repository target

   The Management Services and Management Repository target is monitored by a Management Agent on one of the Management Services on the primary site. To ensure that the target is monitored after switchover/failover, relocate the target to a Management Agent on the standby site by running the following command on one of the Management Service standby sites.

   ```
   emctl config emrep -agent <agent name> -conn_desc
   ```

9. Switchover to Standby SLB.

   Make appropriate network changes to failover your primary SLB to standby SLB that is, all requests should now be served by the standby SLB without requiring any changes on the clients (browser and Management Agents).

10. Establish the old primary Management Services as the new standby Management Services to complete the switchover process.

    Start the Administration Server on old primary site

    ```
    emctl start oms -admin_only
    ```

    Point the old primary Management Services to the new Primary Repository database by running the following command on each Management Service on the old primary site.

```
emctl config oms -store_repos_details -repos_conndesc
<connect descriptor of new primary database> -repos_user
sysman
```

This completes the switchover operation. Access and test the application to ensure that the site is fully operational and functionally equivalent to the primary site. Repeat the same procedure to switchover in the other direction.

23.3.2 Failover

Failover to the Passive OMS in a Level 2 Active/Passive Configuration

1. Establish the IP address on failover host.
2. If the Database and Listener are also part of the same failover group:
    1. Start the TNS listener using the command `lsnrctl start`.
    2. Start the database using the command `dbstart`.
3. Start Cloud Control using the command `emctl start oms`.
4. Test the functionality.

Failover of the OMS in a Level 3 Active/Active Configuration

OMS failover is handled transparently in a Level 3, Active/Active OMS, configuration. The SLB monitors determine that the failed OMS is no longer available and route traffic to available OMS servers.

Manual Failover to the Standby Site in a Level 4 MAA Configuration

This section describes the steps to failover to a standby database, recover the Enterprise Manager application state by resynchronization the Management Repository database with all Management Agents, and enabling the original primary database as a standby using flashback database.

The word *manual* is used here to contrast this type of failover with a fast-start failover described later in Section 23.3.3, "Automatic Failover to the Standby Site in a Level 4 MAA Configuration".

1. Verify Software Library Availability

    Ensure all files from the primary site are available on the standby site.

2. Failover to Standby Database.

    Shutdown the database on the primary site. Use DGMGRL to connect to the standby database and execute the FAILOVER command:

    ```
    FAILOVER TO <standby database name>;
    ```

    Verify the post failover states:

    ```
    SHOW CONFIGURATION;
    SHOW DATABASE <primary database name>;
    SHOW DATABASE <standby database name>;
    ```

    Note that after the failover completes, the original primary database cannot be used as a standby database of the new primary database unless it is re-enabled.

3. Start the Admin Server if it is not already running.

```
emctl start oms -admin_only
```

4. Make the standby Management Services point to the Standby Database which is now the new Primary by running the following on each standby Management Service.

   ```
   emctl config oms -store_repos_details -repos_conndesc <connect descriptor of
   new primary database> -repos_user sysman
   ```

5. Resync the New Primary Database with Management Agents.

   Skip this step if you are running in Data Guard Maximum Protection or Maximum Availability level as there is no data loss on failover. However, if there is data loss, synchronize the new primary database with all Management Agents.

   On any one Management Service on the standby site, run the following command:

   ```
   emctl resync repos -full -name "<name for recovery action>"
   ```

   This command submits a resync job that would be executed on each Management Agent when the Management Services on the standby site are brought up.

   Repository resynchronization is a resource intensive operation. A well tuned Management Repository will help significantly to complete the operation as quickly as possible. Specifically if you are not routinely coalescing the IOTs/indexes associated with Advanced Queueing tables as described in My Oracle Support note 271855.1, running the procedure before resync will significantly help the resync operation to complete faster.

6. Start up the Enterprise Manager Application Tier

   Start up all the Management Services on the standby site by running the following command on each Management Service.

   ```
   emctl start oms
   ```

7. Relocate Management Services and Management Repository target.

   The Management Services and Management Repository target is monitored by a Management Agent on one of the Management Services on the primary site. To ensure that target is monitored after switchover/failover, relocate the target to a Management Agent on the standby site by running the following command on one of the standby site Management Service.

   ```
   emctl config emrep -agent <agent name> -conn_desc
   ```

8. Switchover to the Standby SLB.

   Make appropriate network changes to failover your primary SLB to the standby SLB, that is, all requests should now be served by the standby SLB without requiring any changes on the clients (browser and Management Agents).

9. Establish Original Primary Database as Standby Database Using Flashback

   Once access to the failed site is restored and if you had flashback database enabled, you can reinstate the original primary database as a physical standby of the new primary database.

   1. Shut down all the Management Services in the original primary site.

      ```
      emctl stop oms -all
      ```

   2. Restart the original primary database in mount state:

      ```
      shutdown immediate;
      ```

```
startup mount;
```

3. Reinstate the Original Primary Database

   Use DGMGRL to connect to the old primary database and execute the REINSTATE command

   ```
   REINSTATE DATABASE <old primary database name>;
   ```

4. The newly reinstated standby database will begin serving as standby database to the new primary database.

5. Verify the post reinstate states.

   ```
   SHOW CONFIGURATION;

   SHOW DATABASE <primary database name>;

   SHOW DATABASE <standby database name>;
   ```

10. Establish Original Primary Management Service as the standby Management Service.

    Start the Administration Server on old primary site

    ```
    emctl start oms -admin_only
    ```

    Point the old primary Management Service to the new Primary Repository database by running the following command on each Management Service on the old primary site.

    ```
    emctl config oms -store_repos_details -repos_conndesc <connect descriptor of
    new primary database> -repos_user sysman
    ```

11. Monitor and complete Repository Resynchronization

    Navigate to the Management Services and Repository Overview page of Cloud Control Console. Under Related Links, click Repository Synchronization. This page shows the progress of the resynchronization operation on a per Management Agent basis. Monitor the progress.

    Operations that fail should be resubmitted manually from this page after fixing the error mentioned. Typically, communication related errors are caused by Management Agents being down and can be fixed by resubmitting the operation from this page after restarting the Management Agent.

    For Management Agents that cannot be started due to some reason, for example, old decommissioned Management Agents, the operation should be stopped manually from this page. Resynchronization is deemed complete when all the jobs have a completed or stopped status.

This completes the failover operation. Access and test the application to ensure that the site is fully operational and functionally equivalent to the primary site.

Perform a switchover procedure if the site operations have to be moved back to the original primary site.

23.3.3 Automatic Failover to the Standby Site in a Level 4 MAA Configuration

This section details the steps to achieve complete automation of failure detection and failover procedure by utilizing Fast-Start Failover and Observer process. At a high level the process works as follows:

- Fast-Start Failover (FSFO) determines that a failover is necessary and initiates a failover to the standby database automatically

- When the database failover has completed the DB_ROLE_CHANGE database event is fired
- The event causes a trigger to be fired which calls a script that configures and starts Enterprise Manager Application Tier

Perform the following steps:

1. Develop Enterprise Manager Application Tier Configuration and Startup Script

   Develop a script that will automate the Enterprise Manager Application configuration and startup process. See the sample shipped with Cloud Control in the OH/sysman/ha directory. A sample script for the standby site is included here and should be customized as needed. Make sure ssh equivalence is setup so that remote shell scripts can be executed without password prompts. Place the script in a location accessible from the standby database host. Place a similar script on the primary site.

```
#!/bin/sh
# Script: /scratch/EMSBY_start.sh
# Primary Site Hosts
# Repos: earth, OMS: jupiter1, jupiter2
# Standby Site Hosts
# Repos: mars, # OMS: saturn1, saturn2
LOGFILE="/net/mars/em/failover/em_failover.log"
OMS_ORACLE_HOME="/scratch/OracleHomes/em/oms11"
CENTRAL_AGENT="saturn1.example.com:3872"

#log message
echo "################################" >> $LOGFILE
date >> $LOGFILE
echo $OMS_ORACLE_HOME >> $LOGFILE
id >>  $LOGFILE 2>&1

#switch all OMS to point to new primary and startup all OMS
ssh orausr@saturn1 "$OMS_ORACLE_HOME/bin/emctl oms -store_repos_details -repos_
conndesc <connect descriptor of new primary database> -repos_user sysman
-repos_pwd <password>" >> $LOGFILE 2>&1
ssh orausr@saturn1 "$OMS_ORACLE_HOME/bin/emctl start oms" >>  $LOGFILE 2>&1

#Repeat the above two lines for each OMS in a multiple OMS setup. E.g.
ssh orausr@saturn2 "$OMS_ORACLE_HOME/bin/emctl oms -store_repos_details -repos_
conndesc <connect descriptor of new primary database> -repos_user sysman
-repos_pwd <password>" >> $LOGFILE 2>&1
ssh orausr@saturn2 "$OMS_ORACLE_HOME/bin/emctl start oms" >>  $LOGFILE 2>&1

#relocate Management Services and Repository target
#to be done only once in a multiple OMS setup
#allow time for OMS to be fully initialized
ssh orausr@saturn1 "$OMS_ORACLE_HOME/bin/emctl config emrep -agent $CENTRAL_
AGENT -conn_desc -sysman_pwd <password>" >> $LOGFILE 2>&1

#always return 0 so that dbms scheduler job completes successfully
exit 0
```

2. Automate Execution of Script by Trigger

   Create a database event "DB_ROLE_CHANGE" trigger, which fires after the database role changes from standby to primary. See the sample shipped with Cloud Control in OH/sysman/ha directory.

```
--
```

```
--
-- Sample database role change trigger
--
--
CREATE OR REPLACE TRIGGER FAILOVER_EM
AFTER DB_ROLE_CHANGE ON DATABASE
DECLARE
    v_db_unique_name varchar2(30);
    v_db_role varchar2(30);
BEGIN
    select upper(VALUE) into v_db_unique_name
    from v$parameter where NAME='db_unique_name';
    select database_role into v_db_role
    from v$database;

    if v_db_role = 'PRIMARY' then

      -- Submit job to Resync agents with repository
      -- Needed if running in maximum performance mode
      -- and there are chances of data-loss on failover
      -- Uncomment block below if required
      -- begin
      --  SYSMAN.setemusercontext('SYSMAN', SYSMAN.MGMT_USER.OP_SET_
IDENTIFIER);
      --  SYSMAN.emd_maintenance.full_repository_resync('AUTO-FAILOVER to '||v_
db_unique_name||' - '||systimestamp, true);
      --  SYSMAN.setemusercontext('SYSMAN', SYSMAN.MGMT_USER.OP_CLEAR_
IDENTIFIER);
      -- end;

      -- Start the EM mid-tier
      dbms_scheduler.create_job(
          job_name=>'START_EM',
          job_type=>'executable',
          job_action=> '<location>' || v_db_unique_name|| '_start_oms.sh',
          enabled=>TRUE
      );
    end if;
EXCEPTION
WHEN OTHERS
THEN
    SYSMAN.mgmt_log.log_error('LOGGING', SYSMAN.MGMT_GLOBAL.UNEXPECTED_ERR,
SYSMAN.MGMT_GLOBAL.UNEXPECTED_ERR_M || 'EM_FAILOVER: ' ||SQLERRM);
END;
/
```

> **Note:** Based on your deployment, you might require additional steps to synchronize and automate the failover of SLB and shared storage used for software library. These steps are vendor specific and beyond the scope of this document. One possibility is to invoke these steps from the Enterprise Manager Application Tier startup and configuration script.

3. Configure Fast-Start Failover and Observer.

   Use the Fast-Start Failover configuration wizard in Enterprise Manager Console to enable FSFO and configure the Observer.

This completes the setup of automatic failover..

Part VII

Engineered Systems Management

This section contains the following chapters:

- Discovering and Managing Exadata Targets and Systems
- Using Oracle Exalogic Elastic Cloud

24

Discovering and Managing Exadata Targets and Systems

Oracle Database Machine is an integrated data warehousing solution that eases data warehousing by integrating the whole hardware and software stack into one solution. DB Machine management simplifies monitoring and managing the DB Machine by integrating all hardware and software components into one entity. You do not need to monitor each target individually but instead you can view the whole DB Machine as a single target. You can view all critical issues in the system, monitor performance, and drill down to individual targets from the DB Machine target homepage.

Enterprise Manager automatically or manually discovers the components of the DB Machine and adds these components as managed targets. You can modify the member targets of the DB Machine by adding and removing targets.

The DB Machine homepage includes the following components:

- The hardware schematic allows you to view hardware components of the DB Machine (for example, compute nodes, Exadata cells, and Infiniband switches), monitor critical hardware metrics and view aggregated alerts and faults from all components.
- Alerts sourced from hardware components.
- Easily accessible links and flows to other key feature such as viewing the topology or modifying the schematic

You can also monitor all components of the DB Machine. DB Machine monitors all subcomponent targets, whether hardware or software. This includes the database, ASM, CRS, hosts, Exadata and the Infiniband network. The two targets available to facilitate DB Machine monitoring are:

- Exadata Storage Server
- Infiniband Network Fabric

Note: The following are the minimum requirements for a DB Machine to be monitored and managed by Enterprise Manager Release 12c:

- Minimum required Exadata Software version 11.2.2.3 and above
- Supported hardware version
- Pre-requirement to verify the existence of the Schematic file and the minimum required version. If that does not exist, you must download and run the latest version of Configurator to generate the schematic file.
- Required Schematic file generated with the DB Machine Configurator of version 120 and above on the compute node where the Enterprise Manager agent in installed.

For more information about discovering and managing Exadata targets, see the following topics:

- Automatically Discovering an Oracle Database Machine
- Viewing the Topology of an Existing DB Machine Target
- Drilling Down to Individual Targets
- Viewing Critical Hardware Information for the DB Machine
- Viewing DB Machine Alerts
- Adding Exadata Components Manually

24.1 Automatically Discovering an Oracle Database Machine

To automatically discover a DB Machine target, follow these steps:

1. From the **Setup** menu, select **Add Target**, then select **Add Targets Manually**.

   Enterprise Manager displays the Add Targets Manually page.

2. Choose **Add Non-Host Targets Using Guided Process (Also Adds Related Targets)**.

3. From the Target Types drop-down, choose **Oracle Exadata Database Machine**, then click **Add Using Guided Discovery**.

   Enterprise Manager displays the Oracle Database Machine Discovery page. From here you can add the hardware components such as Exadata Storage Servers and Infiniband Switches in the Oracle Database Machine as managed targets. You can choose to discover a new DB Machine and its hardware components as targets or instead discover newly added hardware components in an existing DB Machine as targets.

4. Choose **Discover a new Database Machine and its hardware components as targets**, then click **Discover Targets**.

   Enterprise Manger displays the Database Machine Discovery wizard that steps you through the process of discovering the Database Machine. The first page of the wizard, the Discovery Inputs page, appears.

5. In the Discovery Agent section of the Discovery Inputs page, choose or enter the **Agent URL** that exists on one of the compute nodes to perform the discovery and

then enter the **Database Oracle Home** of the database installation (release 11.2 or later) on the agent host. In the Infiniband Discovery section, specify the host name of one of the Infiniband switches in the Database Machine and specify the password of the on the Infiniband switch. Click **Next**.

6. On the Components page, the new components discovered in the Oracle Database Machine display.

   In the Compute Node section, select the hosts that are compute nodes in the Oracle Database Machine. The hosts must be added as managed targets before the Oracle Database Machine target promotion can proceed.

   In the Infiniband Switch section, select the Infiniband Switches that are part of the Oracle Database Machine. These also will be added as managed targets.

   In the Ethernet Switch section, select the Ethernet switches that are part of the Oracle Database Machine. The Ethernet switches will be added as managed targets.

   In the Compute Node ILOM section, select the Integrated Lights Out Managers (ILOM) of the compute nodes that are part of this Oracle Database Machine. Theses Integrated Lights Out Managers will be added as managed targets.

   In the KVM section, select the KVM switches that are part of the Oracle Database Machine. The KVM switches will be added as managed targets.

   Click **Next** to continue.

   Enterprise Manager displays the Monitoring Agents page.

7. On the Monitoring Agents page, choose whether you want the monitoring agents to be selected automatically or manually. If you choose **Manually select the agents**, you must add the **Monitoring Agent** and optionally a **Backup Monitoring Agent** for each of the Exadata Cells and Infiniband Switches.

   Click **Next** to continue.

   Enterprise Manager displays the Agent Credential page.

8. On the Agent Credential page, specify whether the Management Agent host users and passwords are the same for all agents. The agent users and passwords are required to set up password-less SSH between the agents and the cells monitored by the agents.

   If the users and passwords are the same, choose **Same for all agents** and enter the user and password combination. If they are not the same for each agent, choose **Different for all agents** and enter each combination for each agent.

   Click **Next**.

9. On the Component Credential page, enter the credentials for the components. In all cases you can choose to enter the same user/password combinations for all components or you can enter the credentials separately for each occurrence.

   Click **Next** to continue.

   Enterprise Manager displays the Component Properties page.

10. On the Component Properties page, specify the target properties of the different components.

    For each component, enter the required information.

    Click **Next** to continue.

    Enterprise Manager displays the Review page.

11. Use the Review page to view the selections you have made in the previous pages of the wizard. When you have verified your selections, click **Submit**.

    Enterprise Manager displays the Target Promotion Summary page that displays the targets that are now managed targets.

24.2 Viewing the Topology of an Existing DB Machine Target

You can view the DB Machine topology of the DB Machine target.

DB Machine management simplifies monitoring and managing the DB Machine by integrating all hardware and software components into one entity. You do not need to monitor each target individually but instead you can view the whole DB Machine as a single target. You can view all critical issues in the system, monitor performance, and drill down to individual targets from the DB Machine target homepage.

Use the Topology page of DB Machine to view the topology of the system by Cluster or by Database. Clusters are a complete software system starting with a RAC database, the underlying ASM, and CRS. Clusters define one logical entity that is interconnected. The DB Machine could include several clusters, one cluster, or could just be a number of individual databases. While cabinets define the hardware topology of the DB Machine, clusters define the logical or system topology of the DB Machine.

You can view the Topology by Cluster or Database. Click an element in the Topology and view alert data associated with the element.

You can monitor all components of the DB Machine. DB Machine monitors all subcomponent targets, whether hardware or software. This includes the database, ASM, CRS, hosts, Exadata and the Infiniband network.

To view the topology of an existing DB Machine target, follow these steps:

1. From the **Targets** menu, select **Exadata**.

   Enterprise Manager displays the Oracle Exadata Database Machines page showing all the available DB Machine targets. From this page you can add hardware components (such as Oracle Exadata Storage Servers, Infiniband switches, Ethernet Switches, KVM switches, PDU, and compute node ILOM) in the Oracle Database Machine as managed targets.

2. From the Oracle Exadata Database Machines page, select the Oracle Database Machine target whose topology you want to view.

3. From the Oracle Database Machine Home page, click **Target**, then select **Members Topology** from the drop-down menu.

   Cloud Control displays the Configuration Topology page.

24.3 Drilling Down to Individual Targets

You can drill down immediately to a subcomponent target of the DB Machine (such as RAC, a database instance, or an Exadata cell).

To drill down to individual targets, follow these tasks:

1. From the **Targets** menu, select **Exadata**.

   Enterprise Manager displays the Oracle Exadata Database Machines page showing all the available DB Machine targets. From this page you can add hardware components (such as Oracle Exadata Storage Servers, Infiniband switches, Ethernet Switches, KVM switches, PDU, and compute node ILOM) in the Oracle Database Machine as managed targets.

2. From the Oracle Exadata Database Machines page, select the Oracle Database Machine target whose components you want to view.

   Enterprise Manager displays the Oracle Database Machine Home page showing an Overview, Schematic, and Incident section for the selected DB Machine.

3. From the Oracle Database Machine Home page, use the left navigation panel to expand the list of available targets that comprise the Database Machine.

4. Click the target to which you want to drill down.

24.4 Viewing Critical Hardware Information for the DB Machine

You can view critical metrics for all the hardware subcomponents of the DB Machine such as DB hosts, Exadata cells, Infiniband switches and so on. These metrics vary for different component targets. For example, database server nodes and Exadata servers include the CPU, I/O, and Storage metrics.

To view critical hardware-centric information for the entire DB machine, follow these steps:

To drill down to individual targets, follow these tasks:

1. From the **Targets** menu, select **Exadata**.

   Enterprise Manager displays the Oracle Exadata Database Machines page showing all the available DB Machine targets. From this page you can add hardware components (such as Oracle Exadata Storage Servers, Infiniband switches, Ethernet Switches, KVM switches, PDU, and compute node ILOM) in the Oracle Database Machine as managed targets.

2. From the Oracle Exadata Database Machines page, select the Oracle Database Machine target whose hardware information you want to view.

3. From the Oracle Database Machine Home page, view the hardware schematic of the Database Machine.

24.5 Viewing DB Machine Alerts

You can view alerts on the DB Machine and drill down to details about each alert. These alerts may be performance/configuration metrics or hardware faults.

To view DB Machine alerts, follow these steps:

1. From the **Targets** menu, select **Exadata**.

   Enterprise Manager displays the Oracle Exadata Database Machines page showing all the available DB Machine targets. From this page you can add hardware components (such as Oracle Exadata Storage Servers, Infiniband switches, Ethernet Switches, KVM switches, PDU, and compute node ILOM) in the Oracle Database Machine as managed targets.

2. From the Oracle Exadata Database Machines page, select the Oracle Database Machine target whose machine configuration information you want to view.

   Enterprise Manager displays the Oracle Database Machine home page on which you can see all alerts associated with the current DB Machine.

24.6 Adding Exadata Components Manually

You can add Exadata components manually using the following steps:

1. From the **Setup** menu, select **Add Target**, then select **Add Targets Manually**.

   Enterprise Manager displays the Add Targets Manually page where you can choose the type of target you want to add.

2. From the Add Targets Manually section, choose **Add Non-Host Targets by Specifying Target Monitoring Properties**.

3. In the **Target Type** combo box, choose the appropriate target type, for example, KVM for kvm, PDU for pdu, Cisco switch for Cisco, and Oracle ILOM Server for ilom plug-in).

4. Choose the monitoring agent using the search option.

5. Click **Add Manually** and provide the required properties

24.7 About Oracle Exadata Storage Server

An Exadata cell is a network-accessible storage array with Exadata software installed on it. Use the Exadata Home page to manage and monitor the HP Oracle Exadata Storage Server (also known as Exadata cell) by managing the Exadata cells as Enterprise Manager Cloud Control targets. You can discover and consolidate management, monitoring and administration of a single or a group of Oracle Exadata Storage Servers in a datacenter using Enterprise Manager.

Exadata cells can be discovered automatically or manually. Once discovered, you can add the Exadata cells as Enterprise Manager targets. The individual Exadata cell is monitored and managed as an Enterprise Manager target and provides the exception, configuration and performance information.

Grouping of Exadata cells is used for easy management and monitoring of the set of Exadata cells. You can group Exadata cells both manually and automatically. The grouping function provides an aggregation of exceptions, configuration and performance information of the group of cells.

You can view performance analysis by linking Exadata performance both at a cell level and group level to ASM and database performance. You can drill down to Exadata configuration and performance issues from both the database and ASM targets.

Storage Grid (for example, multiple database/ASM instances sharing the same Exadata cell) is supported to the same extent as dedicated storage.

For more information about managing Exadata servers, see the following topics:

- Using Exadata As a Cloud Control Target
- Performing Administration Tasks on Exadata Cells
- Performing Administration Tasks on Infiniband Networks
- Launching the IORM Performance Page
- Viewing an Exadata Cell Configuration
- Managing a Single I/O Resource Management Allocation
- Accessing Oracle Support Workbench for Exadata Cell
- Changing the IORM Mode and Updating the IORM Objective

24.7.1 Using Exadata As a Cloud Control Target

Use Oracle Exadata to manage and monitor the HP Oracle Exadata Storage Server (also known as Exadata cell) by managing the Exadata cells as Enterprise Manager

Cloud Control targets. You can discover and consolidate management, monitoring and administration of a single or a group of Oracle Exadata Storage Servers in a datacenter using Enterprise Manager.

Exadata cells can be discovered automatically or manually. Once discovered, you can add the Exadata cells as Enterprise Manager targets.

The individual Exadata cell is monitored and managed as an Enterprise Manager target and provides the exception, configuration and performance information.

Grouping of Exadata cells is used for easy management and monitoring of the set of Exadata cells. You can group Exadata cells both manually and automatically. The grouping function provides an aggregation of exceptions, configuration and performance information of the group of cells.

You can view performance analysis by linking Exadata performance both at a cell level and group level to ASM and database performance. You can drill down to Exadata configuration and performance issues from both the database and ASM targets.

24.7.2 Performing Administration Tasks on Exadata Cells

To perform an administration operation on an Exadata cell, such as executing a cell command, follow these steps:

1. Navigate to the Exadata Cell home page by choosing the Exadata target for which you want to perform an administrative task from the All Targets page.

   Enterprise Manager displays the Exadata Cell Home page for the target you selected.

2. Click **Target**, then select **Administration**.

   From this menu you can choose either **Execute Cell Command, Support Workbench,** or **Manage I/O Resources**.

3. Click **Execute Cell Command**.

   The Command page of the Exadata Cell Administration wizard appears. Enter a CELLCLI command as the administrative command to be executed on the cell. You must read the Command Instructions before you enter a command. Only a single CELLCLI command is allowed to execute. You must enter the command without the 'cellcli -e' prefix, which is automatically appended when you submit the command. Finally, you cannot use the following characters: ; / ' < > / |.

4. Click **Next** to continue.

   Enterprise Manager displays the Admin Credentials page. Select or enter the Administration credentials to execute the command. The credentials you enter are used when submitting the operation. You can choose between Preferred Credentials, Named Credentials, and New Credentials. You can also click **More Details** to view information about Credential Type, Last modified, Credential Name, Credential Owner, Last Modified Date, Last Modified By, and Preferred Credentials Set At.

5. Click **Next**.

   Enterprise Manager displays the Schedule page. Use the Schedule page to schedule the administration task. Enter the Job Name and the Job Description, then provide the job information in the Schedule the Administration Job section. You can choose to begin the job immediately or enter the time you want the job to begin.

6. Click **Next** to continue.

The Summary page displays. Use the Summary page to ensure you have entered the correct values and then submit the command. The Summary page lists the Job Name, Description, Command to Execute, when the job is Scheduled, and the Selected Cell.

7. Click **Submit Command** to submit the job.

   The Job Status page displays. Use the Job Status page to link to the Job Detail page of the administration task.

24.7.3 Performing Administration Tasks on Infiniband Networks

To perform an administration operation on an Infiniband Network, follow these steps:

1. Navigate to the DB Machine home page of the Infiniband Network by choosing the DB Machine for which you want to perform an administrative task from the All Targets page.

   Enterprise Manager displays the DB Machine Home page for the target you selected.

2. Select the IB Network for which you want to perform an administrative task.

3. From the Target menu item, choose **Administration**.

   The Target & Command page of the Infiniband Network Administration wizard appears.

4. Choose the **Target Type** and then select the target on which you want to perform the administrative task from the Target drop-down list. Enter the administrative command you want to execute. The available operations from which you can select are dependent on the target type and target you selected. Once you choose the operation, you may need to select a value that will appear after choosing the operation.

5. Click **Next** to continue.

   Enterprise Manager displays the Credentials & Schedule page. Select or enter the credentials to execute the command. The credentials you enter are used when submitting the operation. You can choose between Preferred Credentials, Named Credentials, and New Credentials. Schedule the administration task. Provide the job information in the **Administration Job Schedule** section. You can choose to begin the job immediately or enter the time you want the job to begin

6. Click **Next** to continue.

   The Review page appears. Use the Review page to ensure you have entered the correct values and then submit the command. The Review page lists the Job Name, Description, Command to Execute, when the job is Scheduled, the Target Type, and the Selected Target.

7. Click **Submit Command** to submit the job.

   When you click Submit Command, a popup is shown if the job is successful. You can go to the Job Detail Page or back to the page from where this wizard was launched.

24.7.4 Launching the IORM Performance Page

To launch the IORM Performance page, follow these steps:

1. Navigate to the Exadata Cell home page by choosing the Exadata target for which you want to view the IORM Performance page from the All Targets page.

Enterprise Manager displays the Exadata Cell Home page for the target you selected.

2. From the Administration menu item, choose **Administration**. and then **Perform IORM Management**.

   The IORM Performance page appears where you can view the status of the current IORM configuration.

24.7.5 Viewing an Exadata Cell Configuration

You can view the configuration of an Oracle Exadata target by following the steps below:

1. Navigate to the Exadata Cell home page by choosing the Exadata target for which you want to view the IORM Performance page from the All Targets page.

   Enterprise Manager displays the Exadata Cell Home page for the target you selected.

2. From the Target menu, choose **Configuration** and then **Topology**.

   Enterprise Manager displays the Configuration Topology page for the selected Exadata Cell. The topology page provides a visual layout of a the target's relationships with other targets. From this page you can:

   - Do a target search filtered by target status/events/target type)
   - Select from a set of relationships to represent in the graph
   - Select annotations to display in the graph, such as alerts and link labels
   - Select from a set of options: view navigator, expand or collapse all, toggle graph layout, reload topology
   - Print
   - Zoom via the slide control
   - Pan via the navigator control
   - Toggle the presentation from graph to table

   When you hover over a node or group member, a popup displays detailed information about the entity. A link can appear in the popup to more detailed information such as customer documentation.

24.7.6 Managing a Single I/O Resource Management Allocation

You can manage the I/O resource allocation for each cell in your Exadata environment. The Manage I/O Resource page displays the current mode (active or inactive) and allows you to change the mode by using the Update IORM Mode function. You cans also change the IORM Objective using the Update IORM Objective function. The page also shows statistics and metrics for the current Exadata cell.

To manage the I/O Resource Management for a single cell, follow these steps:

1. From the **Target** menu of the Exadata cell for which you want to manage I/O resources, choose **Administration**.

2. From the **Administration** menu, select **Manage I/O Resource**.

   The I/O Resource Management page displays where you can change the IORM Mode between Active and Inactive or you can change the IORM Objective.

3. Use the IORM Objective chart to display the historical value of objectives.
4. Use the IORM Wait chart to display the IORM Wait times for all databases or a selected database. You can select the database (or all databases) to display in the IORM Wait chart by choosing the database from the **Select Database** drop-down list.

For more information about changing the IORM mode or updating the IORM objective, see Changing the IORM Mode and Updating the IORM Objective.

24.7.7 Accessing Oracle Support Workbench for Exadata Cell

You can access the Oracle Support Workbench for the current Exadata cell to access diagnostic data for problems and incidents related to the cell.

To access the Support Workbench for a single Exadata cell, follow these steps:

1. From the **Target** menu of the Exadata cell for which you want to access the Oracle Support Workbench, choose **Administration**.
2. From the **Administration** menu, select **Support Workbench**.

   The Login to Support Workbench page displays where you can enter the credentials required to access the Workbench.
3. Choose the type of cell administration credential you want to use from the Credential section.

   You can choose **Preferred Credential**, Named **Credential**, or **New Credentials**.
4. Depending on what you chose in Step 3, enter the Credential Details (**Username** and **Password**) in the appropriate fields. You can turn on the More Details Option to see more details about the credentials.
5. Click **Continue** to display the Oracle Support Workbench.

24.7.8 Changing the IORM Mode and Updating the IORM Objective

Follow these steps to change the IORM Mode or update the IORM Objective.

Changing the IORM Mode

To change the IORM mode from the I/O Resource Management page, follow these steps:

1. From the **Target** menu of the Exadata cell for which you want to manage I/O resources, choose **Administration**.
2. From the **Administration** menu, select**Manage I/O Resource**.

   The I/O Resource Management page displays where you can change the IORM Mode between Active and Inactive. The current IORM Mode is displayed on the page.
3. In the IORM Mode and Objective section, click the mode (Active or Inactive) to which you want to change and then click the **Change IORM Mode** button.

Updating the IORM Objective

To update the IORM Objective from the I/O Resource Management page, follow these steps:

1. From the **Target** menu of the Exadata cell for which you want to manage I/O resources, select **Administration**.

2. From the **Administration** menu, choose **Manage I/O Resource**.

   The I/O Resource Management page displays where you can update the IORM Objective. The current IORM Objective is displayed on the page.

3. From the Change IORM Objective box within the IORM Mode and Objective section, use the Select drop-down list to update the IORM Objective and then click **Update IORM Objective**.

   You can choose one of the IORM Objectives from the list: Low Latency, Balanced, High Throughput, or Auto.

25

Using Oracle Exalogic Elastic Cloud

Use Oracle Exalogic Elastic Cloud to discover and monitor instances of Exalogic Elastic Cloud (Middleware Machine). You can then display status information, including alerts and key performance metrics, of the following target types in an Exalogic Elastic Cloud:

- Application Deployments
- WebLogic Domains
- IB Switch - An IB Switch system has IB Switches as members. This is a one to many association. Each IB switch can have many IB Switch Ports as target components.
- Coherence Clusters

Exalogic Elastic Cloud (Middleware Machine) is modeled as a System target rather than a group target.

For more information about using Oracle Exalogic Elastic Cloud, see the following sections:

- Using the Exalogic Elastic Cloud Discovery Wizard
- Displaying and Using the Exalogic Elastic Cloud Home Page and Dashboard
- Viewing Application Deployments in Exalogic Elastic Cloud Targets
- Viewing WebLogic Domains in Exalogic Elastic Cloud Targets
- Viewing Coherence Clusters in Exalogic Elastic Cloud Targets
- Viewing Hosts in Exalogic Elastic Cloud Targets

25.1 Using the Exalogic Elastic Cloud Discovery Wizard

You can use the Exalogic Elastic Cloud Discovery wizard to discover and monitor an exalogic target in Enterprise Manager.

One important prerequisite for discovery is that all the components running on the Exalogic Elastic Cloud must already exist in Enterprise Manager as targets. Exalogic Elastic Cloud discovery does not add or discover any of these targets. It simply identifies the targets present in the Exalogic Elastic Cloud and maps them to Enterprise Manager targets and then adds these Enterprise Manager targets as Exalogic Elastic Cloud system members.

To use the Exalogic Elastic Cloud Discover wizard, follow these steps:

1. Navigate to the Systems page.

2. From the Add drop-down menu, select **Exalogic Elastic Cloud,** then click **Go**.

   Enterprise Manager displays the Discover Exalogic Elastic Cloud page which allows you to enter the parameters and values required to discover an Oracle Exalogic target.

3. Specify a unique name for the Oracle Exalogic target you want to monitor in the Name field.

4. Select a Management Agent on one of the hosts in the Exalogic System to perform the discovery.

   If you choose a Management Agent that is not part of the Exalogic System, an error message appears stating that no Exalogic Property File can be found and indicating that you must choose a Management Agent which is on a Exalogic System Host.

5. Click **Next**. Enterprise Manager displays the Discover Oracle Exalogic Targets: Discovered Targets page. The page displays the Hosts that are discovered as part of the system.

6. Click **Finish** to complete discovery. You can also choose Back to return to the Discover Exalogic Elastic Cloud page or Cancel to terminate the discovery process.

   Enterprise Manager displays a confirmation that the Exalogic Elastic Cloud instance has been added and begins to monitor the Exalogic Elastic Cloud target. The new target is displayed on the Systems page.

25.2 Displaying and Using the Exalogic Elastic Cloud Home Page and Dashboard

Use the Exalogic Elastic Cloud Dashboard to display the status information including alerts and key performance metrics of the following targets in the Exalogic Elastic Cloud:

- Application Deployments
- WebLogic Domains
- Coherence Clusters
- Hosts

You can also use the Home page to access the Hardware tab where you can view information about the hardware and infrastructure of the Exalogic Elastic Cloud.

To display and use the Exalogic Elastic Cloud Dashboard, follow these steps:

1. Navigate to the Systems page. Filter the entries in the Search field by choosing **Exalogic Elastic Cloud** from the drop-down list, then click **Go**.

   In the Search Results table, choose the Exalogic Elastic Cloud you want to view.

   Enterprise Manager displays the Exalogic Elastic Cloud Home page where you can monitor the status of the Exalogic target and its components on the Exalogic Dashboard.

2. You can view detailed information about each component by choosing the component name from the **Exalogic Elastic Cloud** drop-down menu.

   Enterprise Manager displays the component page you selected. For example, select the WebLogic Domains Summary page to see the charts showing the status of WebLogic Servers, Request Processing Time metric information, CPU Usage,

Requests per minute, and Heap Usage data. For more information about each of those component pages, see the topics in the Related Topics section below.

3. You can return to this page at any time by choosing by choosing Home from the Exalogic Elastic Cloud drop-down menu.
4. You can display General Information about the Exalogic target by choosing General Information from the Exalogic Elastic Cloud drop-down menu.

25.3 Viewing Application Deployments in Exalogic Elastic Cloud Targets

Use the Application Deployments page in the Exalogic Elastic Cloud target area to view details about the applications hosted on the hosts running on the Exalogic Elastic Cloud target.

To display and use the Application Deployments page, follow these steps:

1. Navigate to the Systems page. Filter the entries in the Search field by choosing **Exalogic Elastic Cloud** from the drop-down list, then clicking **Go**.

   In the Search Results table, choose the Exalogic Elastic Cloud you want to view.

   Enterprise Manager displays the Exalogic Elastic Cloud Home page where you can monitor the status of the Exalogic target and its components.
2. Choose **Application Deployments** from the **Exalogic Elastic Cloud** drop-down menu.

   Enterprise Manager displays the Application Deployments page.
3. You can choose to show All Domains or filter by specific domains by choosing the domain from the Show menu.
4. You can drill down to specific applications, targets, domains, or dependencies by clicking its related value in each row.
5. You can filter the list of applications by choosing a value from the Status drop-down. You can select from Up, Down, Unknown, Blackout, and All.
6. You can change the column appearance of the table by clicking **View** and choosing which Columns to display, expanding or collapsing rows, or scrolling to the first or last row. You can also reorder columns.

25.4 Viewing WebLogic Domains in Exalogic Elastic Cloud Targets

Use the WebLogic Domains page in the Exalogic Elastic Cloud target area to view details about the domains hosted on the virtual machines running on Exalogic Elastic Cloud target.

To display and use the WebLogic Domain page, follow these steps:

1. Navigate to the Systems page. Filter the entries in the Search field by choosing **Exalogic Elastic Cloud** from the drop-down list, then click **Go**. In the Search Results table, choose the Exalogic Elastic Cloud you want to view.

   Enterprise Manager displays the Exalogic Elastic Cloud Home page where you can monitor the status of the Exalogic target and its components.
2. Choose **Weblogic Domain** from the Exalogic Elastic Cloud drop-down menu. You can choose to view either a Summary of the WebLogic Domains or specific information about Members.

   Enterprise Manager displays the related WebLogic Domain page.

3. On the Summary page you can view a chart that shows the status of the WebLogic Domains and displays the percentage of domains that are up and down. You can also view server information that shows the Server Status and alert and policy violation information for each. You can monitor charts that display metric information such as Request Processing Time and CPU Usage and you can drill down through these charts for more detailed information. Change the chart view to a table view by clicking **Table View** or **Chart View** beneath each table or chart. The Server table displays information about servers and domains showing host information and related metrics.
4. On the Members page, you can view the Status information along with alerts and policy violation and metric data for each WebLogic Server or Domain. Use the Performance Summary section to view metrics for each, such as Host and Cluster information and metrics such as Heap Usage and Request Processing Time.

25.5 Viewing Coherence Clusters in Exalogic Elastic Cloud Targets

Use the Coherence Clusters page in the Exalogic Elastic Cloud target area to view details about the Coherence targets hosted on the virtual machines running on the Exalogic Elastic Cloud target.

To display and use the Coherence Clusters page, follow these steps:

1. Navigate to the Systems page. Filter the entries in the Search field by choosing **Exalogic Elastic Cloud** from the drop-down list, then click **Go**.

   In the Search Results table, choose the Exalogic Elastic Cloud you want to view.

   Enterprise Manager displays the Exalogic Elastic Cloud Home page where you can monitor the status of the Exalogic target and its components.
2. Choose **Coherence Clusters** from the **Exalogic Elastic Cloud** drop-down menu.

   Enterprise Manager displays the Coherence Clusters page.
3. You can view a chart that shows the status of the Coherence Clusters and displays the percentage of clusters that are up and down.
4. You can drill down to specific values for each cluster such as Alerts and Policy Violations along with Node information.
5. You can filter the list of clusters by choosing a value from the Status drop-down. You can select from Up, Down, Unknown, Blackout, and All.
6. You can change the column appearance of the table by clicking **View** and choosing which Columns to display. You can also reorder columns.
7. The Coherence Clusters page displays two charts showing the Top Nodes With Lowest Available Memory and Caches With Lowest Hit To Get Ratio. You can drill down to specific node information by clicking on the Node name below the Top Nodes With Lowest Available Memory chart.
8. The Nodes table displays information about each Node, including Host and several metric values such as Memory Available, Gets, and Puts.
9. The Applications table displays information about applications such as Local Attribute Cache, Clustered Session Cache, and other metrics. You can drill down to specific information about each application by clicking on the Application name.

25.6 Viewing Hosts in Exalogic Elastic Cloud Targets

Use the Hosts page in the Exalogic Elastic Cloud target area to view details about the host targets hosted on the virtual machines running on the Exalogic Elastic Cloud target.

To display and use the Hosts page, follow these steps:

1. Navigate to the All Targets page. Filter the entries in the Search field by choosing **Exalogic Elastic Cloud** from the drop-down list, then click **Go**.

   In the Search Results table, choose the Exalogic Elastic Cloud you want to view.

   Enterprise Manager displays the Exalogic Elastic Cloud Home page where you can monitor the status of the Exalogic target and its components.

2. Choose **Hosts** from the **Exalogic Elastic Cloud** drop-down menu.

   Enterprise Manager displays the Hosts page.

3. You can view a chart that shows the status of the hosts and displays the percentage of hosts that are up and down.

4. You can view information about the Middleware Targets that lists the Type, Status, CPU Utilization percentage, Memory Utilization percentage, and Incident statistics along with Configuration Changes.

5. You can view charts showing the CPU Utilization percentage based on time and similarly, Memory Utilization based on time.

Index

A

accessing Software Library Administration page, 9-4
accessing Software Library console, 9-1
Add HTTP Location, 9-10
Add NFS Location, 9-10
Add OMS Agent file system, 9-9
Add OMS Agent file system location, 9-10
Add OMS Shared file system, 9-8
ADDM, 2-1
  purpose of, 2-1
administration group
  creating, 6-8
  definition, 6-1
  hierarchy, 6-10
  home page, 6-20
Agent patching directory structure, 16-3
Agent Registration Password, 10-31
  changing, 10-39
AGENT_HOME/network/admin, 10-44
Agents, updating, 15-5
aggregration and purging policies
  See data retention policies
alerts
  server-generated, 2-1
analyzing
  job activity, 8-16
Application Performance Management, 10-75
Application Server Control
  starting and stopping on Windows systems, 17-6
archive logging
  for Management Repository database, 19-1
authentication for Enterprise Manager, 10-2
automated patching
  Offline mode, 16-4
  Online mode, 16-4
automated patching advantages, 16-5
automated patching vs manual patching, 16-4
automatic target discovery
  configuring, 1-1
  promoting targets, 1-5
Availability History Report, picture of, 11-2

B

Beacons
  introduction, 2-2
  monitoring Web Applications over HTTPS, 10-75
benefits of Information Publisher, 11-1
blackouts
  controlling with emctl, 17-11
  examples, 17-13
Bundle Patch Updates, 16-1

C

Certificate dialog box
  Internet Explorer, 10-71, 10-76
changing Management Agent time zone, 17-14
Checkpoint Firewall, Oracle ecosystem and, 2-2
Cloud Control
  starting, 17-8
  starting all components of, 17-8
Cloud Control Mobile
  acknowledge issue, 12-8
  change views, 12-7
  force quit, 12-9
  iDevice, 12-1
  Incident Manager, 12-4
  incidents and problems, 12-6
  iTunes App Store, 12-1
  logging in, 12-2
  manage issue workflow, 12-8
  manage settings, 12-3
  My Oracle Support, 12-6
  requirements, 12-1
  setup, 12-1
  SR number, 12-6
  touch-and-hold, 12-8
Compare Monitoring Template feature, 2-3
configuring
  monitoring templates, 2-3
configuring Software Librar
  installation procedure
    OMS Agent storage, 9-9
    OMS shared file system, 9-8
    referenced storage location, 9-9
configuring Software Library, 9-1
  administrators privileges, 9-3
  installation procedure, 9-8
  maintenance procedure, 9-22
    deleting Software Library storage

location., 9-23
periodic maintenance tasks, 9-23
re-importing Oracle owned entity files, 9-23
overview, 9-1
prerequisites, 9-7
roles and Software Library privileges, 9-4
storage, 9-4
user roles and privileges, 9-3
connect descriptor
using to identify the Management Repository database, 19-9, 19-11
controlling Oracle Management Service on UNIX, 17-5
creating
administration groups, 6-8
custom reports, 11-2
report definitions, 11-2
Critical URL Monitoring, as substitute for Management Agent, 2-2
custom reports, 11-2
customizing Cloud Control pages, 13-1

D

dashboard
groups, 5-4
data retention policies
for Application Performance Management data, 19-3
for other Management data, 19-3
modifying default, 19-3
of the Management Repository, 19-2
when targets are deleted, 19-4
DBSNMP database user, 17-10
setting the password for, 17-10
deleting targets
data retention policies when, 19-4
Diagnostic Patches, 16-1
discovering targets, 1-1
disk mirroring and stripping
Management Repository guideline, 19-1
disk space management
controlling the size and number of log and trace files, 18-12
controlling the size of log and trace files, 18-13
downloading logs, 18-5
dropping the Management Repository, 19-9

E

E-mail Customization, 4-12
e-mail notifications, upper limits, 4-5
EMCLI, setting up, 15-3
emctl
controlling blackouts, 17-11
listing targets on a managed host, 17-11
secure agent utility, 10-35
secure agent utility, sample output, 10-36
secure oms utility, 10-30
secure oms utility,sample output, 10-31
security commands, 10-30
starting, stopping, and checking the Management Service, 17-5
emctl commands, 17-16
OMS, 17-16
secure setpwd, 10-40
secure unlock, 10-38
start oms, sample output, 10-57
emctl config agent listtargets, 17-11
emctl istop, 17-3
emctl patching tool, 16-2
emctl reload, 17-10
emctl start agent, 17-2
emctl start blackout, 17-13
emctl start oms, 17-5
emctl status agent, 17-2
emctl status blackout, 17-13
emctl status oms, 17-6
emctl stop agent, 17-2
emctl stop blackout, 17-13
emctl stop oms, 17-6
emctl upload, 17-10
emctl.log, 18-6, 18-12
emctl.log file, 17-22
emoms_pbs.log, 18-11
emoms_pbs.trc, 18-12
emoms.log, 18-11, 18-12
emomslogging.properties
MaxBackupIndex, 18-13
MaxFileSize, 18-13
emoms.properties
oracle.net.crypto_checksum_client, 10-43
oracle.net.crypto_checksum_types_client, 10-43
oracle.net.encryption_client, 10-42
oracle.net.encryption_types_client, 10-43
oracle.sysman.emRep.dbConn.enableEncryption, 10-42
emoms.trc, 18-11
Enterprise Manager
monitoring templates, 2-3
Enterprise Manager Framework Security
about, 10-30
enabling for Management Repository, 10-41
enabling for multiple Management Services, 10-37
restricting HTTP access, 10-37
types of secure connections, 10-30
Enterprise User Security Based Authentication, 10-2
Extended Network, as substitute for Management Agent, 2-2

F

For, 19-3
force quit
Cloud Control Mobile, 12-9

G

gcagent_errors.log, 18-6

gcagent_mdu.log, 18-6
gcagent.log, 18-6
generating HTML reports, 11-2
Grid Control
  stopping, 17-9
  stopping all components of, 17-9
Group Hierarchy, 6-3
Group Members page, picture of, 5-4
groups
  central monitoring location, 5-2
  dashboard, 5-4
  description and purpose, 5-1
  management features, 5-2
  member targets, 5-4
  redundancy, 5-6

H

home page
  setting, 13-3
HTTP access, restricting, 10-37
HTTPS, 10-30

I

IBM WebSphere
  Oracle ecosystem and, 2-2
iDevice
  supported by Cloud Control Mobile, 12-1
Incident Manager
  Cloud Control Mobile, 12-4
incidents and problems
  Cloud Control Mobile, 12-6
Information Publisher
  Create Like function, 11-2
  generating HTML reports, 11-2
  overview of, 11-1
  predefined reports, 5-5
  report
    definitions, 11-2
    elements, 11-3
  reporting framework, 11-1
  sharing reports, 11-4
  viewing reports, 11-4
Interim Patches, 16-1
Internet Explorer
  Certificate dialog box, 10-71, 10-76
  security alert dialog box, 10-71
istop
  emctl command, 17-3
iTunes App Store
  Cloud Control Mobile, 12-1

J

javax.net.ssl.SSLException
  SSL handshake failed, 10-75
Job Activity page, 8-1
jobs
  analyzing job activity, 8-16
  definition of, 8-1
  Job Activity page, 8-1
  job executions, 8-2
  job runs, 8-2
  multitask, 8-15
  notification rules for e-mail, 8-9
  operations on runs and executions, 8-2
  privileges for sharing job responsibilities, 8-4
  purpose of, 8-1

L

load balancer switches
  BIG-IP, Oracle ecosystem and, 2-2
local store, 15-3
log files
  controlling the size and number of, 18-12
  locating and configuring, 18-1
  locating Management Agent, 18-7
  locating Management Service, 18-12
  Management Agent, 18-6
  Oracle Management Service, 18-11
  searching, 18-3
log4j.appender.emlogAppender.
    MaxBackupIndex, 18-13
log4j.appender.emlogAppender. MaxFileSize, 18-13
log4j.appender.emtrcAppender.
    MaxBackupIndex, 18-13
log4j.appender.emtrcAppender. MaxFileSize, 18-13
logging in to Cloud Control Mobile, 12-2
LVM (Logical Volume Manager), 19-1

M

Management Agent, 18-6
  additional Management Agent commands, 17-10
  checking the status on UNIX, 17-2
  checking the status on Windows, 17-3
  Critical URL Monitoring as substitute, 2-2
  Extended Network as substitute, 2-2
  purpose of, 2-2
  starting and stopping on UNIX, 17-1
  starting and stopping on Windows, 17-3
Management Agent logs
  setting log levels, 18-7, 18-8, 18-9
  setting trace levels, 18-11
Management Agent time zone, 17-14
Management Information Base (MIB), 4-55
  definition, 4-55
  MIB variable descriptions, 4-55
Management Repository
  introduction of, 2-2
  See Oracle Management Repository
Management Repository Deployment
    Guideline, 19-1
Management Service
  starting and stopping on Windows systems, 17-6
managing
  groups, 5-2
managing Cloud Control Mobile sites, 12-3
managing logs, 18-1

MaxBackupIndex
property in emomslogging.properties, 18-13
MaxFileSize
property in emomslogging.properties, 18-13
metrics
threshold values, 2-3
MGMT_ADMIN.DISABLE_METRIC_DELETION, 19-4
MGMT_ADMIN.ENABLE_METRIC_DELETION, 19-4
MGMT_METRICS_1DAY table, 19-3
MGMT_METRICS_1HOUR table, 19-3
MGMT_METRICS_RAW table, 19-3
MIB
<italic>See Management Information Base (MIB)
modes of patching, 16-4
monitoring
alerts as they occur, 5-4
basics of, 2-1
templates
function of, 2-3
monitoring credentials
defined, 17-10
setting, 17-10
monitoring template, comparing, 2-3
multitask jobs, 8-15

N

NetApp Filers
Oracle ecosystem and, 2-2
network/admin, 10-41, 10-43, 10-44
notification methods
based on a PL/SQL Procedure, 4-33
based on an SNMP trap, 4-48
based on operating system commands, 4-17
definition, 4-16
notification rules
custom, 4-10
definition, 4-9
out-of-box, 4-10
out-of-the-box notification rules, 4-6
subscribing to, 4-9
notification schedules, 4-6
notification system
e-mail errors, 4-59
errors, 4-57
trace messages, 4-57
notifications
defining multiple mail servers, 4-3
for jobs, 8-9
long e-mail notifications, 4-6
mail server settings, 4-2
management information base (MIB), 4-55
notification schedules, 4-6
sample Operating System command script, 4-27
setting up, 4-1
short email notifications, 4-6
using custom notfication methods, 4-17

O

Offline mode, 16-4
OMS core patches, 16-2
OMS patch directory structure, 16-2
OMS plugin patches, 16-2
Online mode, 16-4
OPatch, 16-2
Operating System command
sample notification method for, 4-18
sample script, 4-27
Operating System scripts, 4-16
while creating notification methods, 4-17
ORA-12645 (security)
Parameter does not exist, 10-41
Oracle
ecosystem, 2-2
Oracle Access Manager, 10-4
Oracle Access Manager (OAM) SSO, 10-2
Oracle Advanced Security, 10-30, 10-41
enabling for Management Repository, 10-43
enabling for the Management Agent, 10-44
Oracle Enterprise Manager
log files, 18-1
Oracle Enterprise Manager Cloud Control, 17-8
Oracle HTTP Server logs, 18-14
Oracle Management Agent
about log and trace files, 18-6
enabling security for, 10-35, 10-44
location of log and trace files, 18-7
log and trace files, 18-6
starting and stopping, 17-1
Oracle Management Repository
data retention policies, 19-2
dropping, 19-9
enabling Oracle Advanced Security, 10-43
enabling security for, 10-41
identifying with a connect descriptor, 19-9, 19-11
recreating, 19-9, 19-10
reloading data, 17-10
starting the Management Repository database, 17-8
troubleshooting, 19-12
uploading data, 17-10
Oracle Management Service, 17-5
about the log and trace files, 18-11
enabling security for, 10-30
enabling security for multiple Management Services, 10-37
location the log and trace files, 18-12
log and trace files, 18-11
modifying monitoring credentials, 17-10
starting, stopping, and checking, 17-5
Oracle Management Service logs, 18-11, 18-12
Oracle Management Service trace files, 18-13
Oracle Process Management and Notification (OPMN)
using to start and stop the Management Service, 17-6
Oracle WebLogic
Oracle ecosystem and, 2-2

Oracle WebLogic Server logs, 18-14
ORACLE_HOME/network/admin, 10-41, 10-43
oracle.net.crypto_checksum_client
property in emoms.properties, 10-43
oracle.net.crypto_checksum_types_client
property in emoms.properties, 10-43
oracle.net.encryption_client
property in emoms.properties, 10-42
oracle.net.encryption_types_client
property in emoms.properties, 10-43
oracle.sysman.emRep.dbConn.enableEncryption
entry in emoms.properties, 10-42
OS scripts
See Operating System scripts
out-of-box reports, 11-2

P

Patch Recommendations, 16-5
Patch Search region, 16-6, 16-7
Patch Set Updates, 16-1
Patches and Updates, 16-5
Agent patching
Add All To Plan, 16-9
Add Patch to Plan, 16-8
Create Plan, 16-8
Patch Plans, 16-11
Deployment Options, 16-11
Patches page, 16-11
Plan Information, 16-11
Review and Deploy page, 16-12
Validation page, 16-11
View Plan, 16-9
warnings
Null Platform, 16-10
Target is Down, 16-10
Patch Recommendation region, 16-5
Patch Search region, 16-6
Advanced Search, 16-7
Basic Search, 16-6
Patches and Updates vs My Oracle Support, 16-4
patching Enterprise Manager
applying OMS and Repository patches, 16-2
Management Agents
automated patching, 16-4
accessing Patches and Updates, 16-5
applying Agent patches, 16-7
deinstalling Agent patches, 16-13
searching Patches, 16-6
validating applied agent patches, 16-12
verifying the applied agent patches, 16-12
viewing Patch recommendations, 16-5
manual patching, 16-13
overview, 16-3
OMS patches, 16-2
overview, 16-1
patch types, 16-1
Repository patches, 16-2
patching Enterprise Manager core components, 16-1
patching Management Agents, 16-1
patching OMS, 16-1
patching Repository, 16-1
patching tool
emctl, 16-2
OPatch utility, 16-2
Patches and Updates, 16-2
personalize, 13-1
personlizing Cloud Control pages, 13-1
PL/SQL procedures, 4-16
while creating notfication methods, 4-17
while creating notification methods, 4-33
plug-in un-deploy, 15-13
plug-ins
deploying, 15-7, 15-8
external, 15-8
overview, 15-7
removing, 15-13
updating, 15-7, 15-11
privileges
for sharing job responsibilities, 8-4
ProcessManager
service used to control the Management Service on Windows systems, 17-6
promoting discovered targets, 1-5
Public Key Infrastructure (PKI), 10-30, 10-75
purging policies
See data retention policies

R

RAID-capable disk
Management Repository guideline, 19-1
redundancy groups, 5-6
reevaluating metric collections, 17-14
Repeat Notifications, 4-3
Repeat Notifications for Rules, 4-4
RepManager script, 19-9, 19-10
reports
creating custom reports, 11-2
custom, 11-2
definitions, Information Publisher, 11-2
e-mailing, 11-4
generating HTML report, 11-2
Information Publisher, 11-1
out-of-box, Information Publisher, 11-2
predefined, 5-5
predefined report definitions, 11-2
report elements, 11-3
scheduling, 11-3
sharing, 11-4
storing and purging, 11-4

viewing, 11-4
Repository-Based Authentication, 10-2
requirements
Cloud Control Mobile, 12-1
root password
See also SYSMAN
when enabling security for the Management Service, 10-31

S

scheduling
reports, 11-3
reports, flexibility, 11-3
searching logs, 18-3
security
about Enterprise Manager security, 10-1
alert dialog box
Internet Explorer, 10-71
certificate alerts
responding to, 10-71
See also Enterprise Manager Framework Security
security features
See Enterprise Manager Framework Security
Security Patch Updates, 16-1
Self Update feature
setting up, 15-1
using, 15-1
Server Connection Hung
error while creating the repository, 19-12
Server Load Balancer, 10-40
server-generated alerts, 2-1
Services control panel
using to start the Management Service, 17-6
setting
metric threshold values, 2-3
setting your home page, 13-3
sharing reports, 11-4
SNMP traps, 4-16, 4-17, 4-48
sample, 4-48
Software Library, 15-3
designers, 9-4
Operators, 9-4
Super Administrators, 9-4
users, 9-4
Software Library Administration, 9-4, 9-5
referenced file locations, 9-7
Agent storage, 9-7
http storage, 9-7
NFS storage, 9-7
upload file locations, 9-5
OMS Agent file system, 9-6
OMS shared file system, 9-6
Software Library console, 9-2
Software Library referenced locations, 9-4
Software Library storage, 9-1
Software library upload locations, 9-4
SQLNET.CRYPTO_SEED
entry in sqlnet.ora, 10-43, 10-44
SQLNET.ENCRYPTION_SERVER
entry in sqlnet.ora, 10-43
sqlnet.ora, 10-41
SQLNET.CRYPTO_SEED, 10-43, 10-44
SQLNET.ENCRYPTION_SERVER, 10-43
SSO-Based Authentication, 10-2
status codes, corrective actions, 4-53
support of third-party components, 2-1
SYSMAN
checking for existence of, 19-12
entering SYSMAN password when enabling security, 10-31
system errors, notification, 4-57

T

target
definition of, 2-2
target monitoring credentials
defined, 17-10
setting, 17-10
targets
adding manually, 1-5
discovering, 1-1
listing targets on a managed host, 17-11
Template Collections, 6-8
with administration groups, 6-16
third-party
components, support of, 2-1
thresholds
definition of, 2-3
time zone, 17-14
touch-and-hold
Cloud Control Mobile, 12-8
trace files
controlling the contents of Management Service, 18-13
controlling the size and number of, 18-12
locating Management Agent, 18-7
locating Management Service, 18-12
Management Agent, 18-6
Oracle Management Service, 18-11
troubleshooting
general techniques while creating the Management Repository, 19-12
Management Agent startup errors, 17-4
Management Service, 17-6
Management Service startup errors, 17-6
notifications, 4-57
while creating the Management Repository, 19-11
troubleshooting Management Agent, 17-4

U

un-deploying a plug-in, 15-13
updates
applying in offline mode, 15-5
applying in online mode, 15-4
updating Cloud Control, 15-1

V

viewing
reports, 11-4
viewing logs, 18-3
VPN
Cloud Control Mobile requirement, 12-1

W

Web Applications
monitoring over HTTPS, 10-75