

Oracle® Enterprise Manager

Cloud Control Administrator's Guide

12c Release 3 (12.1.0.3)

E24473-24

July 2013

Copyright © 2013, 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	xxiii
Audience	xxiii
Documentation Accessibility	xxiii
Related Documents	xxiii
Conventions	xxiv

Part I Monitoring and Managing Targets

1 Overview of Enterprise Monitoring

1.1	Monitoring Overview	1-1
1.2	Comprehensive Out-of-Box Monitoring	1-1
1.3	Monitoring: Basics	1-2
1.3.1	Metric Thresholds: Determining When a Monitored Condition is an Issue	1-3
1.3.2	Metric Baselines: Determining Valid Metric Thresholds	1-3
1.3.3	Events: Defining What Conditions are of Interest	1-4
1.3.4	Corrective Actions: Resolving Issues Automatically	1-4
1.3.5	Metric Extensions: Customizing Monitoring	1-4
1.3.6	Blackouts	1-5
1.4	Monitoring: Advanced Setup	1-5
1.4.1	Monitoring Templates	1-6
1.4.2	Administration Groups and Template Collections	1-6
1.5	Notifications	1-7
1.5.1	Customizing Notifications	1-7
1.6	Managing Events, Incidents, and Problems	1-8
1.6.1	Incident Manager	1-9
1.6.2	Incident Rules and Rule Sets	1-10
1.6.3	Connectors	1-11
1.7	Accessing Monitoring Information	1-11

2 Discovering and Monitoring Targets

2.1	About Discovering, Promoting, and Monitoring Hosts and Targets	2-1
2.1.1	What is Discovery?	2-2
2.1.2	What is Promotion?	2-2
2.1.3	What is Monitoring?	2-3
2.2	Discovery and Monitoring in Enterprise Manager Lifecycle	2-3

2.3	Overview of the Discovery and Monitoring Process.....	2-3
2.4	Automatically Discovering and Promoting Targets.....	2-4
2.4.1	Prerequisites	2-5
2.4.2	Step 1: Discovering Unmanaged Hosts Using Network Scan	2-6
2.4.3	Step 2: Converting Unmanaged Hosts to Managed Hosts.....	2-8
2.4.4	Step 3: Discovering Targets on Managed Hosts.....	2-8
2.4.5	Step 4: Promoting Targets	2-9
2.5	Manually Adding Targets.....	2-11
2.5.1	Add Host Targets.....	2-11
2.5.2	Add Targets Using Guided Process.....	2-11
2.5.3	Add Targets Declaratively by Specifying Target Monitoring Properties.....	2-13
2.6	Discovering and Promoting Oracle Homes	2-14

3 Using Incident Management

3.1	Events, Incidents, and Problems.....	3-1
3.1.1	Event Management.....	3-1
3.1.2	Incident Management	3-1
3.1.3	Problem Management.....	3-2
3.1.4	Summing Up	3-2
3.1.5	Working with Events, Incidents, and Problems.....	3-3
3.2	Key Concepts: Events, Incidents, and Problems	3-4
3.2.1	Events	3-4
3.2.2	Incidents	3-7
3.2.2.1	Working with Incidents.....	3-7
3.2.2.2	Incident Composed of a Single Event.....	3-10
3.2.2.3	Incident Composed of Multiple Events	3-11
3.2.2.4	How are Incidents Created?.....	3-12
3.2.3	Problems.....	3-12
3.2.4	Rule Sets	3-13
3.2.4.1	Out-of-Box Rule Sets	3-14
3.2.4.2	Rule Set Types	3-15
3.2.4.3	Rules	3-16
3.2.5	Incident Manager.....	3-20
3.2.5.1	Views	3-22
3.3	Setting Up Your Incident Management Environment	3-22
3.3.1	Setting Up Your Monitoring Infrastructure.....	3-23
3.3.1.1	Rule Set Development.....	3-23
3.3.2	Setting Up Administrators and Privileges.....	3-27
3.3.3	Setting Up Rule Sets	3-30
3.3.3.1	Creating a Rule Set	3-30
3.3.3.2	Creating a Rule to Create an Incident.....	3-31
3.3.3.3	Creating a Rule to Manage Escalation of Incidents.....	3-32
3.3.3.4	Creating a Rule to Escalate a Problem.....	3-34
3.3.3.5	Subscribing to a Rule	3-34
3.3.3.6	Receiving E-mail for Private Rules	3-35
3.4	Working with Incidents	3-36
3.4.1	Finding What Needs to be Worked On.....	3-37

3.4.2	Searching for Incidents	3-39
3.4.3	Setting Up Custom Views	3-40
3.4.4	Responding and Working on a Simple Incident	3-41
3.4.5	Responding to and Managing Multiple Incidents, Events and Problems in Bulk..	3-42
3.4.6	Searching My Oracle Support Knowledge	3-44
3.4.7	Open Service Request (Problems-only)	3-44
3.4.8	Suppressing Incidents and Problems	3-45
3.4.9	Managing Workload Distribution of Incidents	3-45
3.4.10	Reviewing Events on a Periodic Basis	3-46
3.4.10.1	Creating an Incident Manually	3-46
3.5	Advanced Topics	3-47
3.5.1	Defining Custom Incident Statuses	3-47
3.5.2	Clearing Stateless Alerts for Metric Alert Event Types	3-47
3.5.3	User-reported Events	3-48
3.5.3.1	Format	3-49
3.5.3.2	Options	3-49
3.5.3.3	Examples	3-50
3.5.4	Additional Rule Applications	3-50
3.5.4.1	Setting Up a Rule to Send Different Notifications for Different Severity States of an Event 3-51	
3.5.4.2	Creating a Rule to Notify Different Administrators Based on the Event Type	3-52
3.5.4.3	Creating a Rule to Create a Ticket for Incidents	3-53
3.5.4.4	Creating a Rule to Send SNMP Traps to Third Party Systems	3-54
3.5.5	Event Prioritization	3-55
3.5.6	Root Cause Analysis (RCA) and Target Down Events	3-56
3.5.6.1	How RCA Works	3-56
3.5.6.2	Leveraging RCA Results in Incident Rule Sets	3-58
3.5.6.3	Leveraging RCA Results in Incident Manager	3-59
3.5.6.4	Leveraging RCA Results in the System Dashboard	3-61
3.5.6.5	Creating a Rule to Update Incident Priority for Non-symptom Events	3-62
3.5.6.6	Creating Incidents On Non-symptom Events	3-63
3.5.6.7	Introducing a Time Delay	3-64
3.6	Moving from Enterprise Manager 10/11g to 12c	3-65

4 Using Notifications

4.1	Setting Up Notifications	4-2
4.1.1	Setting Up a Mail Server for Notifications	4-2
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	4-6
4.1.2.3	Subscribe to Receive E-mail for Incident Rules	4-7
4.1.3	Setting Up E-mail for Other Administrators	4-9
4.1.4	E-mail Customization	4-10
4.1.4.1	E-mail Customization Reference	4-11
4.1.5	Setting Up Repeat Notifications	4-14
4.2	Extending Notification Beyond E-mail	4-15
4.3	Sending Notifications Using OS Commands and Scripts	4-16

4.3.1	Script Examples.....	4-18
4.3.2	Migrating pre-12c OS Command Scripts	4-21
4.3.2.1	Migrating Metric Alert Event Types.....	4-21
4.3.2.2	Migrating Target Availability Event Types	4-22
4.3.2.3	Migrating Job Status Change Event Types	4-22
4.3.2.4	Migrating Corrective Action-Related OS Scripts	4-23
4.3.2.5	Notification Type Mapping.....	4-23
4.4	Sending Notifications Using PL/SQL Procedures.....	4-24
4.4.1	Defining a PL/SQL-based Notification Method	4-24
4.4.2	Migrating Pre-12c PL/SQL Advanced Notification Methods	4-32
4.4.2.1	Mapping for MGMT_NOTIFY_SEVERITY	4-32
4.4.2.2	Mapping for MGMT_NOTIFY_JOB.....	4-36
4.4.2.3	Mapping for MGMT_NOTIFY_CORRECTIVE_ACTION	4-37
4.5	Sending SNMP Traps to Third Party Systems.....	4-38
4.5.1	SNMP Traps: Moving from Previous Enterprise Manager Releases to 12c	4-42
4.6	Management Information Base (MIB).....	4-42
4.6.1	About MIBs.....	4-42
4.6.2	MIB Definition.....	4-43
4.6.3	Reading the MIB Variable Descriptions	4-43
4.6.3.1	Variable Name	4-43
4.7	Passing Corrective Action Status Change Information.....	4-44
4.7.1	Passing Corrective Action Execution Status to an OS Command or Script.....	4-44
4.7.2	Passing Corrective Action Execution Status to a PLSQL Procedure.....	4-45
4.8	Passing Job Execution Status Information.....	4-46
4.8.1	Passing Job Execution Status to a PL/SQL Procedure.....	4-46
4.8.2	Passing Job Execution Status to an OS Command or Script.....	4-48
4.9	Passing User-Defined Target Properties to Notification Methods	4-49
4.10	Notification Reference.....	4-50
4.10.1	EMOMS Properties.....	4-50
4.10.2	Passing Event, Incident, Problem Information to an OS Command or Script.....	4-52
4.10.2.1	Environment Variables Common to Event, Incident and Problem	4-52
4.10.2.2	Event Notification-Specific Environment Variables.....	4-53
4.10.2.3	Environment Variables Specific to Event Types	4-55
4.10.2.4	Environment Variables Specific to Incident Notifications	4-58
4.10.2.5	Environment Variables Specific to Problem Notifications	4-59
4.10.2.6	Environment Variables Common to Incident and Problem Notifications.....	4-60
4.10.3	Passing Information to a PL/SQL Procedure.....	4-61
4.10.3.1	Notification Payload Elements Specific to Event Types	4-69
4.10.4	Troubleshooting Notifications	4-72
4.10.4.1	General Setup	4-72
4.10.4.2	Notification System Errors	4-73
4.10.4.3	Notification System Trace Messages	4-73
4.10.4.4	E-mail Errors	4-75
4.10.4.5	OS Command Errors.....	4-75
4.10.4.6	SNMP Trap Errors.....	4-76
4.10.4.7	PL/SQL Errors	4-76

5 Managing Groups

5.1	Introduction to Groups	5-1
5.1.1	Overview of Groups	5-2
5.1.2	Overview of Privilege Propagating Groups	5-2
5.1.3	Overview of Dynamic Groups	5-3
5.1.4	Overview of Administration Groups	5-3
5.2	Managing Groups	5-4
5.2.1	Creating and Editing Groups	5-4
5.2.2	Creating Dynamic Groups	5-5
5.2.3	Adding Members to Privilege Propagating Groups	5-7
5.2.4	Converting Conventional Groups to Privilege Propagating Groups	5-7
5.2.5	Viewing and Managing Groups	5-8
5.2.6	Overview of Group Charts	5-9
5.2.7	Overview of Group Members	5-10
5.2.8	Viewing Group Status History	5-10
5.2.9	About the System Dashboard	5-10
5.3	Using Out-of-Box Reports	5-11

6 Using 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	6-9
6.3.1	Creating an Administration Group	6-10
6.3.1.1	Accessing the Administration Group Home Page	6-10
6.3.1.2	Defining a Hierarchy	6-11
6.3.1.3	Defining Template Collections	6-15
6.3.1.4	Associating Template Collections with Administration Groups	6-18
6.4	Changing the Administration Group Hierarchy	6-27
6.4.1	Adding a New Hierarchy Level	6-28
6.4.2	Removing a Hierarchy Level	6-28
6.4.3	Merging Administration Groups	6-29
6.4.4	Removing Administration Groups	6-32

7 Using Monitoring Templates

7.1	About Monitoring Templates	7-1
7.2	Definition of a Monitoring Template	7-2
7.3	Default Templates (Auto Apply Templates)	7-2
7.4	Viewing a List of Monitoring Templates	7-2
7.5	Creating a Monitoring Template	7-3
7.6	Editing a Monitoring Template	7-4
7.7	Applying Monitoring Templates to Targets	7-4
7.7.1	Applying a Monitoring Template	7-5
7.7.2	Monitoring Template Application Options	7-5
7.7.2.1	Apply Options	7-5
7.7.2.2	Metrics with Key Value Settings	7-6

7.8	Comparing Monitoring Templates with Targets	7-7
7.8.1	When is a metric between a template and a target considered "different"?	7-8
7.9	Comparing Metric Settings Using Information Publisher	7-8
7.10	Exporting and Importing Monitoring Templates	7-9
7.11	Changing the Monitoring Template Apply History Retention Period	7-10

8 Using Metric Extensions

8.1	What are Metric Extensions?	8-1
8.2	Metric Extension Lifecycle.....	8-3
8.3	Working with Metric Extensions	8-6
8.3.1	Administrator Privilege Requirements	8-6
8.3.2	Granting Create Metric Extension Privilege	8-7
8.3.3	Managing Administrator Privileges	8-7
8.3.4	Managing Administrator Access to Metric Extensions.....	8-8
8.3.4.1	Granting Full/Edit Privileges on a Metric Extension	8-8
8.3.4.2	Revoking Access Privileges on a Metric Extension	8-8
8.3.4.3	Transferring Metric Extension Ownership	8-9
8.3.5	Creating a New Metric Extension	8-9
8.3.6	Creating a New Metric Extension (Create Like)	8-14
8.3.7	Editing a Metric Extension	8-14
8.3.8	Creating the Next Version of an Existing Metric Extension.....	8-15
8.3.9	Importing a Metric Extension	8-15
8.3.10	Exporting a Metric Extension.....	8-16
8.3.11	Deleting a Metric Extension	8-16
8.3.12	Deploying Metric Extensions to a Group of Targets	8-16
8.3.13	Updating Older Versions of Metric Extensions Already deployed to a Group of Targets	8-17
8.4	Adapters	8-18
8.4.1	OS Command Adapter - Single Column.....	8-18
8.4.2	OS Command Adapter- Multiple Values.....	8-21
8.4.3	OS Command Adapter - Multiple Columns.....	8-22
8.4.4	SQL Adapter.....	8-23
8.4.5	SNMP (Simple Network Management Protocol) Adapter	8-24
8.4.6	JMX Adapter.....	8-24
8.5	Converting User-defined Metrics to Metric Extensions	8-25
8.5.1	Overview	8-26
8.5.2	Commands.....	8-26
8.6	Metric Extension Command Line Verbs.....	8-30

9 Utilizing the Job System and Corrective Actions

9.1	Job System Purpose and Overview	9-1
9.1.1	What Are Job Executions and Job Runs?.....	9-2
9.1.2	Operations on Job Executions and Job Runs	9-2
9.2	Preliminary Considerations.....	9-3
9.2.1	Creating Scripts	9-3
9.2.2	Sharing Job Responsibilities	9-4
9.2.3	Submitting Jobs for Groups.....	9-4

9.3	Creating Jobs.....	9-4
9.3.1	Selecting a Job Type.....	9-4
9.3.2	Creating an OS Command Job.....	9-5
9.3.2.1	Specifying a Single Operation.....	9-10
9.3.2.2	Specifying a Script	9-11
9.3.2.3	Access Level Rules.....	9-12
9.3.3	Creating a SQL Script Job	9-12
9.3.3.1	Specifying Targets	9-12
9.3.3.2	Specifying Options for the Parameters Page	9-13
9.3.3.3	Specifying Host and Database Credentials.....	9-13
9.3.3.4	Returning Error Codes from SQL Script Jobs.....	9-14
9.3.4	Creating a Multi-task Job.....	9-15
9.3.4.1	Job Capabilities	9-15
9.3.4.2	Specifying Targets for a Multi-task Job	9-15
9.3.4.3	Adding Tasks to the Job.....	9-15
9.4	Analyzing Job Activity	9-16
9.5	Generating Job Event Criteria	9-16
9.5.1	Enabling Events For Job Status, Status Severity, and Targetless Jobs	9-17
9.5.2	Adding Targets To Generate Events For Job Status	9-18
9.6	Creating Event Rules For Job Status Change.....	9-18
9.6.1	Creating Job Status Change Event Rules For Jobs	9-18
9.6.2	Creating Job Status Change Event Rules For Targets	9-21
9.7	Creating Corrective Actions	9-24
9.7.1	Providing Credentials	9-25
9.7.2	Creating Corrective Actions for Metrics.....	9-25
9.7.3	Creating a Library Corrective Action	9-26
9.7.4	Specifying Access to Corrective Actions	9-27
9.7.4.1	Defining or Modifying Access	9-27
9.7.4.2	Access Level Rules.....	9-27
9.7.5	Setting Up Notifications for Corrective Actions	9-28
9.7.6	Providing Agent-side Response Actions.....	9-29
9.7.6.1	Specifying Commands and Scripts	9-29
9.7.6.2	Using Target Properties in Commands.....	9-29
9.7.6.3	Using Advanced Capabilities	9-30

10 Configuring and Using Services

10.1	Introduction to Services	10-1
10.1.1	Defining Services in Enterprise Manager.....	10-2
10.1.2	Modeling Services.....	10-2
10.1.2.1	Availability of a Service.....	10-2
10.1.2.2	Performance and Usage.....	10-3
10.1.2.3	Service Level Agreements (SLA).....	10-3
10.2	Creating and Configuring a Service.....	10-4
10.2.1	Creating a Generic Service	10-4
10.2.2	Configuring a Service.....	10-6
10.2.2.1	Root Cause Analysis Configuration	10-6
10.2.2.2	Service Tests and Beacons	10-8

10.2.2.3	Availability Definition	10-10
10.2.2.4	Performance Metrics	10-11
10.2.2.5	Usage Metrics	10-12
10.2.3	Configuring an Aggregate Service	10-13
10.3	Creating a Test Based Service.....	10-14
10.3.1	Creating an ATS Service Test Using OATS Load Script	10-15
10.3.1.1	Creating a Self Contained Zip File	10-16
10.3.1.2	Creating an ATS Service Test.....	10-17
10.3.1.3	Databanking and Parameterization	10-18
10.3.1.4	Parameterizing URLs	10-19
10.3.1.5	Success or Failure Validation.....	10-19
10.3.1.6	Using Beacon Override	10-20
10.3.1.7	Updating the Databank File.....	10-20
10.3.1.8	Using SLM Header for RUEI Integration.....	10-20
10.3.2	Using the Transaction Recorder	10-20
10.4	Creating a System Based Service	10-21
10.5	Setting Up and Using Service Level Agreements	10-22
10.5.1	Actionable Item Rules for SLAs.....	10-23
10.5.2	Creating a Service Level Objective	10-24
10.5.3	Lifecycle of an SLA	10-25
10.5.4	Viewing the Status of SLAs for a Service	10-26
10.5.5	Defining Custom SLA Business Calendars.....	10-27
10.6	Using the Services Dashboard	10-27
10.6.1	Viewing the All Dashboards Page	10-28
10.6.2	Viewing the Dashboard Details Page	10-28
10.6.3	Customizing and Personalizing the Dashboard	10-29
10.6.4	Viewing the Dashboard Service Details Page.....	10-29
10.7	Monitoring Settings	10-30
10.8	Setting Up Monitoring Templates	10-31
10.8.1	Configuring Service Tests and Beacons for the Monitoring Template	10-32
10.9	Configuring Service Levels.....	10-32
10.9.1	Defining Service Level Rules	10-33
10.9.2	Viewing Service Level Details	10-34
10.10	Configuring a Service Using the Command Line Interface.....	10-34
10.11	Troubleshooting Service Tests	10-38
10.11.1	Verifying and Troubleshooting Forms Transactions.....	10-38
10.11.1.1	Troubleshooting Forms Transaction Playback.....	10-38
10.11.1.2	Troubleshooting Forms Transaction Recording	10-40
10.11.2	Verifying and Troubleshooting Web Transactions.....	10-40

11 Configuring a Software Library

11.1	Overview of Software Library	11-1
11.2	Users, Roles, and Privileges.....	11-3
11.3	What's New in Software Library	11-5
11.4	Performing Software Library Tasks Using EM CLI Verbs or in Graphical Mode	11-5
11.5	Software Library Storage	11-6
11.5.1	Upload File Locations	11-8

11.5.2	Referenced File Location.....	11-9
11.6	Prerequisites for Configuring Software Library.....	11-10
11.7	Configuring Software Library Storage Location	11-10
11.7.1	Configuring an OMS Shared Filesystem Location.....	11-11
11.7.2	Configuring an OMS Agent Filesystem Location	11-11
11.7.3	Configuring a Referenced File Location	11-13
11.8	Configuring Software Library on a Multi-OMS System	11-14
11.9	Using Software Library Entities.....	11-15
11.10	Tasks Performed Using the Software Library Home Page	11-16
11.10.1	Organizing Entities.....	11-16
11.10.2	Creating Entities.....	11-17
11.10.2.1	Creating Generic Components	11-17
11.10.2.2	Creating Directives.....	11-19
11.10.3	Customizing Entities	11-21
11.10.4	Managing Entities.....	11-22
11.10.4.1	Accessing Software Library Home Page.....	11-22
11.10.4.2	Accessing Software Library Administration Page.....	11-22
11.10.4.3	Granting or Revoking Privileges	11-22
11.10.4.4	Moving Entities.....	11-23
11.10.4.5	Changing Entity Maturity	11-23
11.10.4.6	Adding Notes to Entities	11-24
11.10.4.7	Adding Attachments to Entities.....	11-24
11.10.4.8	Viewing, Editing, and Deleting Entities	11-24
11.10.4.9	Purging Deleted Entities.....	11-25
11.10.4.10	Searching Entities	11-26
11.10.4.11	Exporting Entities.....	11-27
11.10.4.12	Importing Entities.....	11-28
11.10.4.13	Staging Files Associated With an Entity	11-28
11.11	Maintaining Software Library	11-29
11.11.1	Periodic Maintenance Tasks.....	11-30
11.11.2	Re-Importing Oracle Owned Entity Files.....	11-30
11.11.3	Removing (and Migrating) Software Library Storage Location	11-30
11.11.4	Scheduling Purge Job	11-31
11.11.5	Backing Up Software Library.....	11-32

12 Introducing Enterprise Manager Support for SNMP

12.1	Benefits of SNMP Support.....	12-1
12.2	About the SNMP Management Station	12-2
12.3	How Enterprise Manager Supports SNMP	12-2
12.4	Sending SNMP Trap Notifications	12-4
12.4.1	About the Management Information Base (MIB).....	12-5
12.5	Monitoring External Devices Using SNMP	12-5
12.5.1	About SNMP Receivelets.....	12-5
12.5.2	About SNMP Fetchlets.....	12-6
12.6	About Metric Extensions.....	12-6

Part II Security

13 Configuring Security

Part III Generating Reports

14 Using Information Publisher

14.1	About Information Publisher	14-1
14.2	Out-of-Box Report Definitions	14-2
14.3	Custom Reports.....	14-2
14.3.1	Creating Custom Reports	14-2
14.3.2	Report Parameters	14-3
14.3.3	Report Elements	14-3
14.4	Scheduling Reports.....	14-3
14.4.1	Flexible Schedules.....	14-3
14.4.2	Storing and Purging Report Copies	14-4
14.4.3	E-mailing Reports	14-4
14.5	Sharing Reports	14-4

Part IV Accessing Enterprise Manager via Mobile Devices

15 Remote Access To Enterprise Manager

15.1	Reviewing System Requirements	15-1
15.2	Performing Initial Setup.....	15-1
15.3	Connecting the First Time.....	15-2
15.4	Encountering the Login Screen	15-2
15.5	Managing Settings	15-3
15.6	Using Cloud Control Mobile in Incident Manager	15-4
15.7	Working in Cloud Control Mobile	15-6
15.7.1	Viewing Incidents and Problems	15-6
15.7.2	Changing Views.....	15-8
15.7.3	Performing Actions	15-8
15.8	Learning Tips and Tricks	15-8
15.9	Connecting to Enterprise Manager Desktop Version.....	15-9

Part V Engineered Systems Management

16 Discovering and Managing Exadata Targets and Systems

16.1	Automatically Discovering an Oracle Database Machine	16-2
16.2	Viewing the Topology of an Existing DB Machine Target	16-4
16.3	Drilling Down to Individual Targets	16-4
16.4	Viewing Critical Hardware Information for the DB Machine.....	16-5
16.5	Viewing DB Machine Alerts.....	16-5
16.6	Adding Exadata Components Manually.....	16-5
16.7	About Oracle Exadata Storage Server.....	16-6

16.7.1	Using Exadata As a Cloud Control Target.....	16-6
16.7.2	Performing Administration Tasks on Exadata Cells	16-7
16.7.3	Performing Administration Tasks on Infiniband Networks	16-8
16.7.4	Launching the IORM Performance Page.....	16-8
16.7.5	Viewing an Exadata Cell Configuration.....	16-9
16.7.6	Manage IO Resource	16-9
16.7.6.1	Modify the Disk I/O Objective.....	16-10
16.7.6.2	The Inter-Database Plan	16-11
16.7.7	Accessing Oracle Support Workbench for Exadata Cell.....	16-12
16.7.8	Changing the IORM Mode and Updating the IORM Objective	16-12

Part VI Administering Cloud Control

17 Personalizing Cloud Control

17.1	Personalizing a Cloud Control Page	17-1
17.2	Customizing a Region	17-2
17.3	Setting Your Homepage.....	17-3

18 Plug-In Manager

18.1	Introduction to Plug-ins	18-1
18.1.1	What Is the Enterprise Manager 12c Extensibility Paradigm?	18-1
18.1.2	What Are Plug-Ins?	18-2
18.1.3	Are All Plug-Ins Deployed by Default?.....	18-2
18.1.4	How Often Are Plug-Ins Released?	18-2
18.2	Workflow of Plug-In Deployment.....	18-3
18.3	Introduction to Plug-In Manager.....	18-6
18.3.1	What is Plug-In Manager?	18-7
18.3.2	Accessing and Performing Operations Using Plug-In Manager	18-7
18.3.2.1	Accessing Plug-In Manager	18-7
18.3.2.2	Performing Operations Using Plug-In Manager.....	18-7
18.4	Actions that Can Be Performed Using Plug-In Manager	18-8
18.4.1	Customizing Your View	18-9
18.4.1.1	Customizing Displayed Plug-Ins	18-9
18.4.1.2	Customizing Displayed Columns	18-9
18.4.2	Checking the Availability of Plug-Ins.....	18-10
18.4.3	Viewing Information about Plug-Ins	18-10
18.4.3.1	Differentiating Plug-In Releases from Enterprise Manager Platform Releases	18-10
18.4.3.2	Identifying Plug-In ID.....	18-11
18.4.3.3	Viewing Targets and Operating Systems Certified for Deployed Plug-Ins....	18-11
18.4.3.4	Viewing Plug-In Dependencies	18-11
18.4.3.5	Verifying Deployed Plug-Ins	18-12
18.4.4	Downloading Plug-Ins	18-13
18.4.4.1	Downloading Plug-Ins in Online Mode	18-13
18.4.4.2	Downloading Plug-Ins in Offline Mode	18-14
18.4.4.3	Importing Catalog Archives	18-14

18.4.4.4	Importing Plug-In Archives.....	18-15
18.4.5	Deploying Plug-Ins to Oracle Management Service (Reduce OMS Restart time and Downtime) 18-16	
18.4.5.1	Tracking the Deployment Status of Plug-Ins on Oracle Management Service	18-18
18.4.6	Upgrading Plug-Ins Deployed to Oracle Management Service	18-19
18.4.6.1	Upgrading Across Plug-In Versions Deployed to Oracle Management Service	18-19
18.4.6.2	Upgrading Across Plug-In Revisions Within a Plug-In Version Deployed to Oracle Management Service 18-19	
18.4.7	Deploying Plug-Ins on Oracle Management Agent	18-20
18.4.7.1	Tracking the Deployment Status of Plug-Ins on Oracle Management Agent	18-20
18.4.8	Re-deploying Plug-Ins on Oracle Management Agent	18-20
18.4.9	Upgrading Plug-Ins Deployed to Oracle Management Agent	18-21
18.4.10	Deploying Plug-In Patches While Deploying or Upgrading Management Agent (Create Custom Plug-In Update) 18-22	
18.4.10.1	Creating Custom Plug-In Update Using EMCLI	18-22
18.4.10.2	Creating Custom Plug-In Update Using EDK	18-24
18.4.11	Undeploying Plug-Ins from Oracle Management Agent	18-25
18.4.12	Undeploying Plug-Ins from Oracle Management Service.....	18-25
18.5	Troubleshooting	18-26
18.5.1	Understanding Plug-In Homes.....	18-26
18.5.2	Troubleshooting OMS Plug-In Deployment and Upgrade Issues	18-27
18.5.2.1	Troubleshooting OMS Plug-In Deployment Issues.....	18-27
18.5.2.2	Rollback and Resume OMS Plug-In Upgrade.....	18-28
18.5.3	Troubleshooting Management Agent Plug-In Deployment and Upgrade Issues	18-28
18.5.3.1	Troubleshooting Management Agent Plug-In Deployment Issues	18-29
18.5.3.2	Troubleshooting Management Agent Plug-In Upgrade Issues.....	18-29

19 Maintaining Enterprise Manager

19.1	Overview: Managing the Manager	19-1
19.2	Health Overview	19-2
19.3	Repository	19-4
19.4	Controlling and Configuring Management Agents.....	19-6
19.4.1	Manage Cloud Control Agents Page	19-6
19.4.2	Agent Home Page.....	19-7
19.4.3	Controlling a Single Agent	19-8
19.4.4	Configuring Single Management Agents.....	19-9
19.4.5	Controlling Multiple Management Agents.....	19-9
19.4.6	Configuring Multiple Agents.....	19-11
19.4.7	Upgrading Multiple Management Agents.....	19-12
19.5	Management Services.....	19-12
19.6	Viewing Enterprise Manager Topology and Charts.....	19-13

20 Maintaining and Troubleshooting the Management Repository

20.1	Management Repository Deployment Guidelines	20-1
20.2	Management Repository Data Retention Policies.....	20-2

20.2.1	Management Repository Default Aggregation and Purging Policies.....	20-2
20.2.2	Management Repository Default Aggregation and Purging Policies for Other Management Data	20-3
20.2.3	Modifying the Default Aggregation and Purging Policies.....	20-3
20.2.4	Modifying Data Retention Policies When Targets Are Deleted	20-4
20.2.5	How to Modify the Retention Period of Job History.....	20-5
20.2.6	DBMS_SCHEDULER Troubleshooting.....	20-7
20.3	Dropping and Recreating the Management Repository	20-8
20.3.1	Dropping the Management Repository.....	20-8
20.3.2	Recreating the Management Repository	20-9
20.3.2.1	Using a Connect Descriptor to Identify the Management Repository Database	20-10
20.4	Troubleshooting Management Repository Creation Errors	20-10
20.4.1	Package Body Does Not Exist Error While Creating the Management Repository.....	20-10
20.4.2	Server Connection Hung Error While Creating the Management Repository.....	20-11
20.4.3	General Troubleshooting Techniques for Creating the Management Repository	20-11
20.5	Cross Platform Enterprise Manager Repository Migration.....	20-12
20.5.1	Common Prerequisites.....	20-12
20.5.2	Methodologies.....	20-13
20.5.2.1	Cross Platform Transportable Database.....	20-13
20.5.2.2	Migration Using Physical Standby	20-18
20.5.3	Post Migration Verification	20-19

21 Updating Cloud Control

21.1	Using Self Update	21-1
21.1.1	What Can Be Updated?	21-1
21.2	Setting Up Self Update	21-2
21.2.1	Setting Up Enterprise Manager Self Update Mode	21-2
21.2.2	Assigning Self Update Privileges to Users.....	21-2
21.2.3	Setting Up the Software Library	21-3
21.2.4	Setting Up the EM CLI Utility (Optional)	21-3
21.3	Applying an Update	21-3
21.3.1	Applying an Update in Online Mode.....	21-4
21.3.2	Applying an Update in Offline Mode.....	21-5
21.4	Accessing Informational Updates	21-5
21.5	Acquiring or Updating Management Agent Software.....	21-6

22 Patching Enterprise Manager Agents

22.1	Overview	22-1
22.2	Automated Agent Patching (Recommended Option).....	22-2
22.2.1	Accessing Patch Plans	22-2
22.2.2	Viewing Patch Recommendations	22-3
22.2.3	Searching Patches	22-3
22.2.4	Applying Management Agent Patches.....	22-5
22.2.5	Verifying the Applied Agent Patches	22-10
22.2.6	Validating Agent Patch Errors.....	22-10

22.3	Manual Agent Patching	22-13
------	-----------------------------	-------

23 Patching the OMS and Repository

23.1	OPatch Automation	23-1
23.1.1	Supported OMS Configurations and OPatchauto Patchability	23-1
23.1.2	OUI Inventory Configurations	23-2
23.1.3	Supported Patch Format	23-3
23.1.4	Supported Patching Methodologies	23-3
23.2	Required OPatchauto Parameters	23-3
23.2.1	Creating a Property File	23-4
23.3	Prerequisites for Running OPatchauto	23-4
23.4	Using OPatchauto	23-7
23.4.1	My Oracle Support: Searching for Patches	23-8
23.4.2	Running <code>opatchauto apply</code>	23-12
23.4.3	Running <code>opatchauto rollback</code>	23-13
23.4.4	Running <code>opatchauto lspatches</code>	23-13
23.4.5	Running <code>opatchauto version</code>	23-14
23.4.6	Patching a Standby OMS System	23-14
	OPatchauto Command Syntax	23-15
	Apply	23-16
	Rollback	23-18
	lspatches	23-20
	version	23-21
	Standby OMS Patching	23-22
	Troubleshooting	23-23

24 Starting and Stopping Enterprise Manager Components

24.1	Controlling the Oracle Management Agent	24-1
24.1.1	Starting, Stopping, and Checking the Status of the Management Agent on UNIX	24-1
24.1.2	Starting and Stopping the Management Agent on Windows	24-3
24.1.3	Checking the Status of the Management Agent on Windows	24-3
24.1.4	Troubleshooting Management Agent Startup Errors	24-4
24.1.4.1	Management Agent starts up but is not ready	24-4
24.1.4.2	Management Agent fails to start because of time zone mismatch between agent and OMS	24-4
24.1.4.3	Agent fails to start due to possible port conflict	24-4
24.1.4.4	Agent secure/unsecure fails	24-5
24.2	Controlling the Oracle Management Service	24-5
24.2.1	Controlling the Management Service on UNIX	24-5
24.2.2	Controlling the Management Service on Windows	24-6
24.2.3	Troubleshooting Oracle Management Service Startup Errors	24-6
24.3	Controlling JVM and ADP Engines	24-7
24.3.1	Controlling JVM Engines	24-7
24.3.2	Controlling ADP Engines	24-8
24.4	Guidelines for Starting Multiple Enterprise Manager Components on a Single Host	24-8
24.5	Starting and Stopping Oracle Enterprise Manager 12c Cloud Control	24-9

24.5.1	Starting Cloud Control and All Its Components	24-9
24.5.2	Stopping Cloud Control and All Its Components	24-10
24.6	Additional Management Agent Commands	24-11
24.6.1	Uploading and Reloading Data to the Management Repository	24-11
24.6.2	Specifying New Target Monitoring Credentials	24-11
24.6.3	Listing the Targets on a Managed Host	24-12
24.6.4	Controlling Blackouts.....	24-13
24.6.5	Changing the Management Agent Time Zone.....	24-15
24.6.6	Reevaluating Metric Collections.....	24-15
24.7	emctl Commands	24-17
24.8	Using emctl.log File	24-24

25 Locating and Configuring Enterprise Manager Log Files

25.1	Managing Log Files	25-1
25.1.1	Viewing Log Files and Their Messages	25-3
25.1.1.1	Restricting Access to the View Log Messages Menu Item and Functionality ..	25-4
25.1.2	Searching Log Files.....	25-5
25.1.2.1	Searching Log Files: Basic Searches	25-5
25.1.2.2	Searching Log Files: Advanced Searches	25-6
25.1.3	Downloading Log Files.....	25-7
25.2	Managing Saved Searches	25-8
25.2.1	Saving Searches	25-8
25.2.2	Retrieving Saved Searches.....	25-8
25.2.3	Managing Saved Searches	25-9
25.3	Locating Management Agent Log and Trace Files	25-9
25.3.1	About the Management Agent Log and Trace Files.....	25-10
25.3.1.1	Structure of Agent Log Files	25-10
25.3.2	Locating the Management Agent Log and Trace Files.....	25-11
25.3.3	Setting Oracle Management Agent Log Levels	25-11
25.3.3.1	Modifying the Default Logging Level.....	25-12
25.3.3.2	Setting gcagent.log	25-12
25.3.3.3	Setting gcagent_error.log.....	25-13
25.3.3.4	Setting the Log Level for Individual Classes and Packages.....	25-13
25.3.3.5	Setting gcagent_mdu.log.....	25-13
25.3.3.6	Setting the TRACE Level.....	25-15
25.4	Locating and Configuring Oracle Management Service Log and Trace Files	25-15
25.4.1	About the Oracle Management Service Log and Trace Files	25-16
25.4.2	Locating Oracle Management Service Log and Trace Files.....	25-16
25.4.3	Controlling the Size and Number of Oracle Management Service Log and Trace Files...	25-17
25.4.4	Controlling the Contents of the Oracle Management Service Trace File	25-18
25.4.5	Controlling the Oracle WebLogic Server and Oracle HTTP Server Log Files	25-19
25.5	Monitoring Log Files	25-21
25.5.1	About Log Viewer	25-21
25.5.2	Overview of WebLogic Server and Application Deployment Log File Monitoring	25-22
25.5.3	Enabling Log File Monitoring.....	25-23

25.5.4	Configuring Log File Monitoring.....	25-23
25.5.5	Viewing Alerts from Log File Monitoring	25-25
25.6	Configuring Log Archive Locations.....	25-25

Part VII Configuring Enterprise Manager for High Availability

26 High Availability Solutions

26.1	Latest High Availability Information.....	26-1
26.2	Defining High Availability	26-2
26.2.1	Levels of High Availability	26-2
26.3	Comparing Availability Levels.....	26-3
26.4	Implementing High Availability Levels.....	26-4

27 Enterprise Manager High Availability

27.1	Agent High Availability.....	27-1
27.1.1	Configuring the Management Agent to Automatically Start on Boot and Restart on Failure 27-1	
27.1.2	Configuring Restart for the Management Agent	27-1
27.1.3	Installing the Management Agent Software on Redundant Storage	27-2
27.2	Repository High Availability	27-2
27.2.1	General Best Practice for Repository High Availability.....	27-2
27.2.2	Configuring RAC for the Management Repository.....	27-2
27.3	Oracle Management Service High Availability	27-3
27.3.1	Configuring the Cloud Control OMS in an Active/Passive Environment for High Availability Failover Using Virtual Host Names 27-3	
27.3.1.1	Overview and Requirements	27-3
27.3.1.2	Installation and Configuration	27-3
27.3.1.3	Setting Up the Virtual Host Name/Virtual IP Address	27-4
27.3.1.4	Setting Up Shared Storage.....	27-4
27.3.1.5	Setting Up the Environment	27-4
27.3.1.6	Synchronizing Operating System IDs	27-5
27.3.1.7	Setting Up Shared Inventory	27-5
27.3.1.8	Installing the Software	27-5
27.3.1.9	Starting Up Services	27-6
27.3.1.10	Repository Host Name Change.....	27-6
27.3.1.11	Multiple OMS, SLB, First OMS Host Name Change.....	27-7
27.3.1.12	Multiple OMS, SLB, Additional OMS Host Name Change	27-9
27.3.2	Configuring the Cloud Control OMS for High Availability Failover.....	27-10
27.3.2.1	Configuring the Software Library	27-10
27.3.2.2	Configuring a Load Balancer	27-10
27.3.2.3	Installing Additional Management Services	27-13

28 Enterprise Manager Disaster Recovery

28.1	Disaster Recover Overview and Topology	28-1
28.2	Design Considerations	28-3
28.2.1	Network Considerations.....	28-3

28.2.1.1	Planning Host Names	28-4
28.2.1.2	Load Balancers Consideration	28-5
28.2.1.3	Application Virtual Host Name Consideration	28-5
28.2.2	Storage Considerations	28-6
28.2.3	Database Considerations	28-7
28.2.4	Starting Points	28-8
28.2.4.1	The primary site is already created, standby site is being planned	28-8
28.2.4.2	The primary site is already created, standby site is already created using the deprecated "Standby WLS Domain" method. 28-9	
28.2.4.3	No installation exists, both primary and standby sites are being planned.....	28-9
28.3	Setting Up Management Repository Disaster Recovery	28-9
28.3.1	Configuring a Standby Database for the Management Repository	28-9
28.4	Setting Up Management Service and Software Library Disaster Recovery.....	28-10
28.4.1	Management Service Disaster Recovery	28-11
28.4.2	Monitoring Standby OMS Hosts	28-13
28.4.3	Software Library Disaster Recovery	28-14
28.4.4	Migrating an Existing Site to Shared Storage	28-14
28.5	Performing Switchover and Failover Operations	28-14
28.5.1	Switchover Procedure	28-15
28.5.2	Failover Procedure.....	28-16
28.6	Keeping the Standby Site in Sync with the Primary	28-17

29 Backing Up and Recovering Enterprise Manager

29.1	Backing Up Your Deployment.....	29-1
29.2	Software Library Backup	29-1
29.3	Management Repository Backup.....	29-2
29.4	Oracle Management Service Backup.....	29-2
29.5	Management Agent Backup	29-4
29.6	Recovery of Failed Enterprise Manager Components.....	29-4
29.6.1	Repository Recovery	29-4
29.6.2	Recovery Scenarios	29-5
29.6.2.1	Full Recovery on the Same Host	29-6
29.6.2.2	Incomplete Recovery on the Same Host.....	29-6
29.6.2.3	Full Recovery on a Different Host.....	29-6
29.6.2.4	Incomplete Recovery on a Different Host.....	29-7
29.6.3	Recovering the OMS.....	29-8
29.6.3.1	Recovering the Software Homes	29-8
29.6.3.2	Recreating the OMS.....	29-9
29.6.4	OMS Recovery Scenarios.....	29-9
29.6.4.1	Single OMS, No Server Load Balancer (SLB), OMS Restored on the same Host..... 29-9	
29.6.4.2	Single OMS, No SLB, OMS Restored on a Different Host.....	29-11
29.6.4.3	Single OMS, No SLB, OMS Restored on a Different Host using the Original Hostname 29-12	
29.6.4.4	Multiple OMS, Server Load Balancer, Primary OMS Recovered on the Same Host.. 29-13	

29.6.4.5	Multiple OMS, Server Load Balancer Configured, Primary OMS Recovered on a Different Host	29-15
29.6.4.6	Multiple OMS, SLB configured, additional OMS recovered on same or different host	29-16
29.6.5	Recovering the Software Library	29-17
29.6.6	Recovering Management Agents	29-17
29.6.7	Management Agent Recovery Scenarios.....	29-18
29.6.7.1	Management Agent Reinstall Using the Same Port	29-18
29.6.7.2	Management Agent Restore from Filesystem Backup	29-18
29.7	Recovering from a Simultaneous OMS-Management Repository Failure	29-19
29.7.1	Collapsed Configuration: Incomplete Management Repository Recovery, Primary OMS on the Same Host	29-19
29.7.2	Distributed Configuration: Incomplete Management Repository Recovery, Primary OMS and additional OMS on Different Hosts, SLB Configured	29-20
n	oraEMNGEvent.....	A-1
Table A-1	oraEMNGEventIndex.....	A-3
Table A-1	oraEMNGEventNotifType	A-4
n	oraEMNGEventMessage	A-4
n	oraEMNGEventMessageURL	A-4
n	oraEMNGEventSeverity	A-5
n	oraEMNGEventSeverityCode	A-5
n	oraEMNGEventRepeatCount	A-5
n	oraEMNGEventActionMsg	A-6
n	oraEMNGEventOccurrenceTime	A-6
n	oraEMNGEventReportedTime	A-6
n	oraEMNGEventCategories.....	A-7
n	oraEMNGEventCategoryCodes	A-7
n	oraEMNGEventType.....	A-7
n	oraEMNGEventName.....	A-8
n	oraEMNGAssocIncidentId	A-8
n	oraEMNGAssocIncidentOwner.....	A-9
n	oraEMNGAssocIncidentAcked	A-9
n	oraEMNGAssocIncidentStatus	A-9
n	oraEMNGAssocIncidentPriority	A-10
n	oraEMNGAssocIncidentEscLevel	A-10
n	oraEMNGEventTargetName	A-10
n	oraEMNGEventTargetNameURL	A-11
n	oraEMNGEventTargetType	A-11
n	oraEMNGEventHostName	A-11
n	oraEMNGEventTargetOwner	A-12
n	oraEMNGEventTgtLifeCycleStatus	A-12
n	oraEMNGEventTargetVersion	A-12
n	oraEMNGEventUserDefinedTgtProp.....	A-12
n	oraEMNGEventSourceObjName.....	A-13
n	oraEMNGEventSourceObjNameURL	A-13
n	oraEMNGEventSourceObjType	A-13
n	oraEMNGEventSourceObjSubType.....	A-14
n	oraEMNGEventSourceObjOwner	A-14

n	oraEMNGEventCAJobName	A-14
n	oraEMNGEventCAJobStatus	A-15
n	oraEMNGEventCAJobOwner	A-15
n	oraEMNGEventCAJobStepOutput	A-15
n	oraEMNGEventCAJobType	A-16
n	oraEMNGEventRuleSetName	A-16
n	oraEMNGEventRuleName	A-16
n	oraEMNGEventRuleOwner	A-17
n	oraEMNGEventSequenceId	A-17
n	oraEMNGEventRCADetails	A-17
n	oraEMNGEventContextAttrs	A-18
n	oraEMNGEventUserComments	A-18
n	oraEMNGEventUpdates	A-18
n	oraEMNGEventTypeAttr(1-71)	A-19
	Table A-11 oraEM4AlertTable	A-22
	Table A-12 oraEM4AlertTargetName	A-23
	Table A-12 oraEM4AlertTargetType	A-24
	Table A-12 oraEM4AlertHostName	A-24
	Table A-12 oraEM4AlertMetricName	A-24
	Table A-12 oraEM4AlertKeyName	A-25
	Table A-12 oraEM4AlertKeyValue	A-25
	Table A-12 oraEM4AlertTimeStamp	A-26
	Table A-12 oraEM4AlertSeverity	A-26
	Table A-12 oraEM4AlertMessage	A-27
	Table A-12 oraEM4AlertRuleName	A-27
	Table A-12 oraEM4AlertRuleOwner	A-28
	Table A-12 oraEM4AlertMetricValue	A-28
	Table A-12 oraEM4AlertContext	A-28
	Table A-12 oraEM4AlertCycleGuid	A-29
	Table A-12 oraEM4AlertRepeatCount	A-29
	Table A-12 oraEM4AlertUDTargetProperties	A-30
	Table A-12 oraEM4AlertAck	A-30
	Table A-12 oraEM4AlertAckBy	A-31
	Table A-12 oraEM4AlertNotifType	A-31
n	oraEM4AlertViolationGuid	A-32
n	oraEM4JobAlertTable	A-32
	Table A-13 oraEM4JobAlertJobName	A-33
	Table A-13 oraEM4JobAlertJobOwner	A-33
	Table A-13 oraEM4JobAlertJobType	A-34
	Table A-13 oraEM4JobAlertJobStatus	A-34
	Table A-13 oraEM4JobAlertTargets	A-34
	Table A-13 oraEM4JobAlertTimeStamp	A-35
	Table A-13 oraEM4JobAlertRuleName	A-35
	Table A-13 oraEM4JobAlertRuleOwner	A-36
	Table A-13 oraEM4JobAlertMetricName	A-36
	Table A-13 oraEM4JobAlertMetricValue	A-37
	Table A-13 oraEM4JobAlertContext	A-37

Table A–13 oraEM4JobAlertKeyName.....	A-38
Table A–13 oraEM4JobAlertKeyValue.....	A-38
Table A–13 oraEM4JobAlertSeverity	A-38
Table A–13 oraEM4JobAlertJobId	A-39
Table A–13 oraEM4JobAlertJobExecId.....	A-39
B.1 MIB Definition.....	B-1
C.1 Pre-12c Enterprise Manager Metric Alerts	C-1
C.2 Pre-12C Target Availability Alerts	C-2
C.3 Pre-12C Corrective Action Results for Metric Alerts.....	C-2
C.4 Corrective Action Results for Target Availability.....	C-3
C.5 Job Status Change	C-4
D.1 Configuring Standby Management Services on a Standby Site.....	D-1
D.1.1 Installing the First Standby Management Service	D-2
D.1.2 Installing the First Standby Management Service (Oracle Installer, Software Only Mode) D-3	
D.1.3 Installing Additional Standby Management Services.....	D-6
D.1.4 Installing Additional Standby Management Services (Oracle Installer, Software Only Mode) D-7	
D.1.5 Validating Your Installation and Complete the Setup	D-10
D.1.5.1 Keeping the Standby Site in Sync.....	D-10

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 at

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>

Access to Oracle Support

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

Related Documents

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

Part I

Monitoring and Managing Targets

This section contains the following chapters:

- [Overview of Enterprise Monitoring](#)
- [Discovering and Monitoring Targets](#)
- [Using Incident Management](#)
- [Using Notifications](#)
- [Using Metric Extensions](#)
- [Using Blackouts](#)
- [Managing Groups](#)
- [Using Administration Groups](#)
- [Using Monitoring Templates](#)
- [Configuring and Using Services](#)
- [Utilizing the Job System and Corrective Actions](#)
- [Configuring a Software Library](#)
- [Introducing Enterprise Manager Support for SNMP](#)

Overview of Enterprise Monitoring

This chapter covers the following topics:

- [Monitoring Overview](#)
- [Monitoring: Basics](#)
- [Monitoring: Advanced Setup](#)
- [Notifications](#)
- [Managing Events, Incidents, and Problems](#)
- [Accessing Monitoring Information](#)

1.1 Monitoring Overview

Enterprise Manager Cloud Control 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 (targets) on that host. If a target goes down, or if a performance metric crosses a warning or critical threshold, an event is triggered and sent to Enterprise Manager. Administrators or any interested party can be notified of the triggered event through the Enterprise Manager notification system.

Adding targets to monitor is simple. Enterprise Manager provides you with the option of either adding targets manually or automatically discovering all targets on a host. Enterprise Manager can also automatically and intelligently apply monitoring settings for newly added targets. For more information, see [Section 1.4.2, "Administration Groups and Template Collections"](#)). While Enterprise Manager provides a comprehensive set of metrics used for monitoring, you can also use metric extensions (see [Section 1.3.5, "Metric Extensions: Customizing Monitoring"](#)) to monitor conditions that are specific to your environment. As your data center grows, it will become more challenging to manage individual targets separately, thus you can use Enterprise Manager's group management functionality to organize large sets of targets into groups, allowing you to monitor and manage many targets as one.

1.2 Comprehensive Out-of-Box Monitoring

Monitoring begins as soon as you install Enterprise Manager Cloud Control 12c. 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. 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.
- Monitoring 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)

Monitoring Without Management Agents

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. Enterprise Manager monitoring concepts and the underlying subsystems that support this functionality are discussed in the following sections.

1.3 Monitoring: Basics

Enterprise Manager Cloud Control 12c comes with a comprehensive set of predefined performance and health metrics that enables automated 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. While Enterprise Manager can monitor for many types of conditions (events), the most common use of its monitoring capability centers around the basics of monitoring for violation of acceptable performance boundaries defined by metric

values. The following sections discuss the basic concepts and Enterprise Manager functionality that supports monitoring of targets.

1.3.1 Metric Thresholds: Determining When a Monitored Condition is an Issue

Some metrics have associated predefined limiting parameters called thresholds that cause metric alerts (specific type of event) 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%.

While the out-of-box predefined metric threshold values will work for most monitoring conditions, your environment may require that you customize threshold values to more accurately reflect the operational norms of your environment. Setting accurate threshold values, however, may be more challenging for certain categories of metrics such as performance metrics.

For example, what are appropriate warning and critical thresholds for the *Response Time Per Transaction* database metric? For such metrics, it might make more sense to be alerted when the monitored values for the performance metric deviates from normal behavior. Enterprise Manager provides features to enable you to capture normal performance behavior for a target and determine thresholds that are deviations from that performance norm.

Note: Enterprise Manager administrators must be granted *Manage Target Metrics* or greater privilege on a target in order to perform any metric threshold changes.

1.3.2 Metric Baselines: Determining Valid Metric Thresholds

Determining what metric threshold values accurately reflect the performance monitoring needs of your environment is not trivial. Rather than relying on trial and error to determine the correct values, Enterprise Manager provides metric baselines. Metric baselines are well-defined time intervals (baseline periods) over which Enterprise Manager has captured system performance metrics, creating statistical characterizations of system performance over specific time periods. This historical data greatly simplifies the task of determining valid metric threshold values by providing normalized views of system performance. Baseline normalized views of metric behavior help administrators explain and understand event occurrences.

The underlying assumption of metric baselines is that systems with relatively stable performance should exhibit similar metric observations (values) over times of comparable workload. Two types of baseline periods are supported:

- **Moving Window Baseline Periods:** Moving window baseline periods are defined as some number of days prior to the current date (Example: Last 7 days). This allows comparison of current metric values with recently observed history. Moving window baselines are useful for operational systems with predictable workload cycles (Example: OLTP days and batch nights).

- **Static Baseline Periods:** Static baselines are periods of time that you define that are of particular interest to you (Example: End of the fiscal year). These baselines can be used to characterize workload periods for comparison against future occurrences of that workload (Example: compare end of the fiscal year from one calendar year to the next).

1.3.3 Events: Defining What Conditions are of Interest

When a metric threshold value is reached, a metric alert is raised. A metric alert is a type of event. An event is a significant occurrence that indicates a potential problem; for example, either a warning or critical threshold for a monitored metric has been crossed. Other examples of events include: database instance is down, a configuration file has been changed, job executions ended in failure, or a host exceeded a specified percentage CPU utilization. Two of the most important event types used in enterprise monitoring are:

- Metric Alert
- Target Availability

For more information on events and available event types for which you can monitor, see [Chapter 3, "Using Incident Management"](#).

1.3.4 Corrective Actions: Resolving Issues Automatically

Corrective actions allow you to specify automated responses to metric alerts, saving administrator time and ensuring issues are dealt with before they noticeably impact users. For example, if Enterprise Manager detects that a component, such as the SQL*Net listener is down, a corrective action can be specified to automatically start it back up. A corrective action is, therefore, any task you specify that will be executed when a metric triggers a warning or critical alert severity. By default, the corrective action runs on the target on which the event has been raised. A corrective action can also consist of multiple tasks, with each task running on a different target.

Administrators can also receive notifications for the success or failure of corrective actions. A corrective action can also consist of multiple tasks, with each task running on a different target.

Corrective actions for a target can be defined by all Enterprise Manager administrators who have been granted *Manage Target Metrics* or greater privilege on the target. For any metric, you can define different corrective actions when the metric triggers at warning severity or at critical severity.

Corrective actions must run using the credentials of a specific Enterprise Manager administrator. For this reason, whenever a corrective action is created or modified, the credentials that the modified action will run with must be specified.

1.3.5 Metric Extensions: Customizing Monitoring

Metric Extensions let you extend Enterprise Manager's monitoring capabilities to cover conditions specific to your IT environment, thus providing you with a complete and comprehensive view of your monitored environment.

Metric extensions allow you to define new metrics on any target type that utilize the same full set of data collection mechanisms used by Oracle provided metrics. For example, some target types you can create metrics on are:

- Hosts
- Databases

- IBM Websphere
- Oracle Exadata Databases and Storage Servers
- Oracle Business Intelligence Components

Once these new metrics are defined, they are used like any other Enterprise Manager metric. For more information about metric extensions, see [Chapter 8, "Using Metric Extensions"](#).

User-Defined Metrics (Pre-12c)

If you upgraded your Enterprise Manager 12c site from an older version of Enterprise Manager, then all user-defined metrics defined in the older version will also be migrated to Enterprise Manager 12c. These user-defined metrics will continue to work, however they will no longer be supported a future release. 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. For information about the migration process, see *Converting User-defined Metrics to Metric Extensions* in [Chapter 8, "Using Metric Extensions"](#)

1.3.6 Blackouts

Blackouts allow you to support planned outage periods to perform emergency or scheduled maintenance. When a target is put under blackout, monitoring is suspended, thus preventing unnecessary alerts from being sent when you bring down a target for scheduled maintenance operations such as database backup or hardware upgrade. Blackout periods are automatically excluded when calculating a target's overall availability.

A blackout period can be defined for individual targets, a group of targets or for all targets on a host. The blackout can be scheduled to run immediately or in the future, and to run indefinitely or stop after a specific duration. Blackouts can be created on an as-needed basis, or scheduled to run at regular intervals. If, during the maintenance period, you discover that you need more (or less) time to complete maintenance tasks, you can easily extend (or stop) the blackout that is currently in effect. Blackout functionality is available from both the Enterprise Manager console as well as via the Enterprise Manager command-line interface (EM CLI). EM CLI is often useful for administrators who would like to incorporate the blacking out of a target within their maintenance scripts. When a blackout ends, the Management Agent automatically re-evaluates all metrics for the target to provide current status of the target post-blackout.

If an administrator inadvertently performs scheduled maintenance on a target without first putting the target under blackout, these periods would be reflected as target downtime instead of planned blackout periods. This has an adverse impact on the target's availability records. In such cases, Enterprise Manager allows Super Administrators to go back and define the blackout period that should have happened at that time. The ability to create these retroactive blackouts provides Super Administrators with the flexibility to define a more accurate picture of target availability.

1.4 Monitoring: Advanced Setup

Enterprise Manager greatly simplifies managing your monitored environment and also allows you to customize and extend Enterprise Manager monitoring capabilities. However, the primary advantage Enterprise Manager monitoring provides is the ability to monitor and manage large-scale, heterogeneous environments. Whether you

are monitoring an environment with 10 targets or 10,000 targets, the following Enterprise Manager advanced features allow you to implement and maintain your monitored environment with the equal levels of convenience and simplicity.

1.4.1 Monitoring Templates

Monitoring Templates simplify the task of standardizing monitoring settings across your enterprise by allowing you to specify your standards for monitoring in a template once and apply them to monitored targets across your organization. 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. The apply operation can be automated using Administration Groups and Template Collections. For any target, you can preserve custom monitoring settings by specifying metric settings that can never be overwritten by a template.

Enterprise Manager comes with an array of Oracle-certified templates that provide recommended metric settings for various Oracle target types.

For more information about monitoring templates, see [Chapter 7, "Using Monitoring Templates"](#).

1.4.2 Administration Groups and Template Collections

Monitored environments are rarely static—new targets are constantly being added from across your ecosystem. Enterprise Manager allows you to maintain control of this dynamic environment through administration groups. Administration groups automate the process of setting up targets for management in Enterprise Manager by automatically applying management settings such as monitoring settings or compliance standards. Typically, these settings are manually applied to individual targets, or perhaps semi-automatically using monitoring templates (see [Section 1.4.1, "Monitoring Templates"](#)) or custom scripts. Administration groups combine the convenience of applying monitoring settings using monitoring templates with the power of automation.

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.

Every target added to Enterprise Manager possesses innate attributes called *target properties*. Enterprise Manager uses these target properties to add targets to the correct administration group. Administration group membership is based on target properties as membership criteria so target membership is dynamic. Once added to the administration group, Enterprise Manager automatically applies the requisite

monitoring settings using monitoring templates that are part of the associated template collection .

Administration groups use the following target properties to define membership criteria:

- Contact
- Cost Center
- Customer Support Identifier
- Department
- Lifecycle Status
- Line of Business
- Location
- Target Version
- Target Type

1.5 Notifications

For a typical monitoring scenario, when a target becomes unavailable or if thresholds for performance are crossed, events are raised and notifications are sent to the appropriate administrators. Enterprise Manager supports notifications via e-mail, pager, SNMP traps, or by running custom scripts and allows administrators to control these notification mechanisms through:

- Notification Methods
- Rules and Rule Sets

Notification Methods

A notification method represents a specific way to send notifications. Besides e-mail, there are three types of notification methods: OS Command, PL/SQL, SNMP Traps. When configuring a notification method, you need to specify the particulars associated with a specific notification mechanism such as which SMTP gateway(s) to use for e-mail or which custom OS script to run. Super Administrators perform a one-time setup of the various types of notification methods available for use.

Rules

A rule instructs Enterprise Manager to take specific action when events or incidents (entity containing one important event or related events) occur, such as notifying an administrator or opening a helpdesk ticket (see [Section 1.6, "Managing Events, Incidents, and Problems"](#)). For example, you can define a rule that specifies e-mail should be sent to you when CPU Utilization on any host target is at critical severity, or another rule that notifies an administrator's supervisor if an incident is not acknowledged within 24 hours.

1.5.1 Customizing Notifications

Notifications that are sent to Administrators can be customized based on message type and on-call schedule. Message customization is useful for administrators who rely on both e-mail and paging systems as a means for receiving notifications. The message formats for these systems typically vary—messages sent to e-mail can be lengthy and can contain URLs, and messages sent to a pager are brief and limited to a finite

number of characters. To support these types of mechanisms, Enterprise Manager allows administrators to associate a long or short message format with each e-mail address. E-mail addresses that are used to send regular e-mails can be associated with the *long* format; pages can be associated with the *short* format. The *long* format contains full details about the event/incident; the *short* format contains the most critical pieces of information.

Notifications can also be customized based on an administrator's on-call schedule. An administrator who is on-call might want to be contacted by both his pager and work e-mail address during business hours and only by his pager address during off hours. Enterprise Manager offers a flexible notification schedule to support the wide variety of on-call schedules. Using this schedule, an administrator defines his on-call schedule by specifying the e-mail addresses by which they should be contacted when they are on-call. For periods where they are not on-call, or do not wish to receive notifications for incidents, they simply leave that part of the schedule blank. All alerts that are sent to an administrator automatically adhere to his specified schedule.

1.6 Managing Events, Incidents, and Problems

Enterprise Manager's monitoring functionality is built upon the precept of monitoring by exception. This means it monitors and raises events when exception conditions exist in your IT environment and allowing administrators to address them in a timely manner. As discussed earlier, the two most commonly used event types to monitor for are metric alert and target availability. Although these are the most common event types for which Enterprise Manager monitors, there are many others. Available event types include:

- Target Availability
- Metric Alert
- Metric Evaluation Errors
- Job Status Changes
- Compliance Standard Rule Violations
- Compliance Standard Score Violations
- High Availability
- Service Level Agreement Alerts
- User-reported
- Application Dependency and Performance Alert
- JVM Diagnostics Threshold Violation

By definition, an incident is a unit containing a single, or closely correlated set of events that identify an issue that needs administrator attention within your managed environment. So an incident might be as simple as a single event indicating available space in a tablespace has fallen below a specified limit, or more complex such as an incident consisting of multiple events relating to potential performance issue when a server is running out of resources. Such an incident would contain events relating to the usage of CPU, I/O, and memory resources. Managing by incident gives you the ability to address issues that may consist of any number of causal factors. For an in-depth discussion on incidents and events, see [Chapter 3, "Using Incident Management"](#).

Although incidents can correspond to a single events, incidents more commonly correspond to groups of related events. A large number of discrete events can quickly

become unmanageable, but handled as an assemblage of related events, incidents allow you to manage large numbers of event occurrences more effectively.

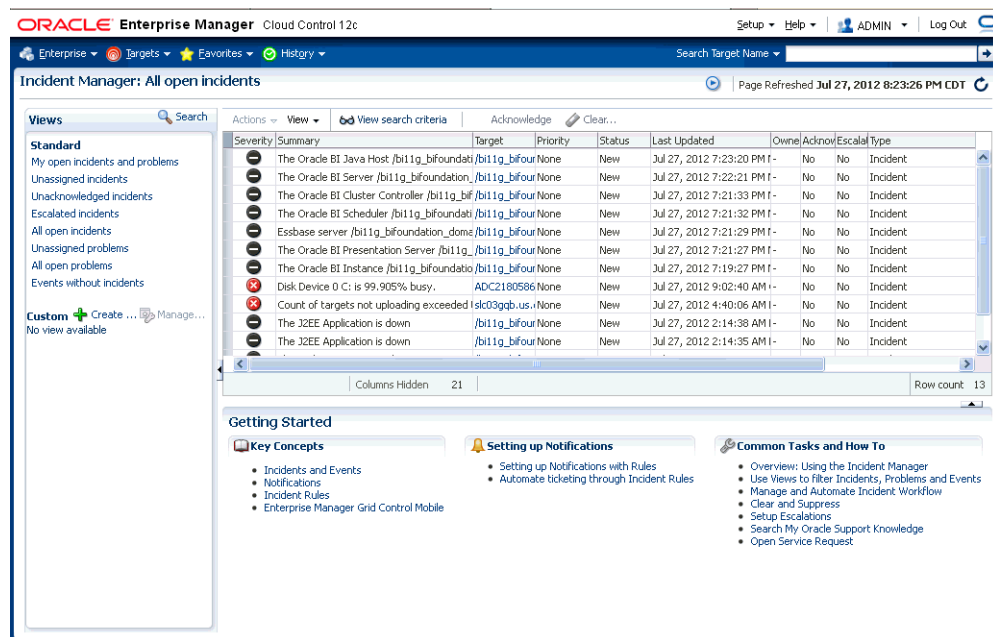
Once an incident is created, Enterprise Manager makes available a rich set of incident management workflow features that let you to manage and track the incident through its complete lifecycle. Incident management features include:

- Assign incident ownership.
- Track the incident resolution status.
- Set incident priority.
- Set incident escalation level.
- Ability to provide a manual summary.
- Ability to add user comments.
- Ability to suppress/unsuppress
- Ability to manually clear the incident.
- Ability to create a ticket manually.

Problems pertain to the diagnostic incidents and problems stored in Automatic Diagnostic Repository (ADR), which are automatically raised by Oracle software when it encounters critical errors in the software. When problems are 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. A problem represents the underlying root cause of a set of incidents. Enterprise Manager provides features to track and manage the lifecycle of a problem.

1.6.1 Incident Manager

Enterprise Manager Cloud Control simplifies managing incidents through an intuitive UI called Incident Manager. Incident Manager provides an easy-to-use interface that allows you to search, view, manage, and resolve incidents and problems impacting your environment. To access Incident Manager, from the **Enterprise** menu, select **Monitoring**, and then **Incident Manager**.

Figure 1–1 Incident Manager

From the Incident Manager UI, you can:

- Filter incidents, problems, and events by using custom views
- Respond and work on an incident
- Manage incident lifecycle including assigning, acknowledging, tracking its status, prioritization, and escalation
- 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.

1.6.2 Incident Rules and Rule Sets

An *incident rule* specifies criteria and actions that determine when a notification should be sent and how it should be sent whenever an event or incident is raised. The criteria defined within a rule can apply to attributes such as the target type, events and severity states (clear, warning or critical) and the notification method that should be used when an incident is raised that matches the rule criteria. Rule actions can be conditional in nature. For example, a rule action can be defined to page a user when an incident severity is critical or just send e-mail if it is warning.

A *rule set* is a collection of rules that apply to a common set of targets such as hosts, databases, groups, jobs, metric extensions, or self updates and take appropriate actions to automate the business processes underlying incident. Incident rule sets can be made public for sharing across administrators. For example, administrators can subscribe to the same rule set if they are interested in receiving notifications for the same criteria defined in the rule. Alternatively, an Enterprise Manager Super Administrator can assign incident rule sets to other administrators so that they receive notifications for incidents as defined in the rule.

In addition to being used by the notification system (see *Rules* in [Section 1.5, "Notifications"](#)), rule sets can also instruct Enterprise Manager to perform other

actions, such as creating incidents, updating incidents, or call into a trouble ticketing system as discussed in [Section 1.6.3, "Connectors"](#).

1.6.3 Connectors

An Oracle Management Connector integrates third-party management systems with Enterprise Manager. There are two types of connectors: Event connectors and helpdesk connectors.

Using the event connector, you can configure Enterprise Manager to share events with non-Oracle management systems. The connector monitors all events sent from Oracle Enterprise Manager and automatically updates alert information in the third-party management system. Event connectors support the following functions:

- Sharing of event information from Oracle Enterprise Manager to the third-party management system.
- Customization of event to alert mappings between Oracle Enterprise Manager and the third-party management system.
- Synchronization of event changes in Oracle Enterprise Manager with the alerts in the third-party management system.

Using the helpdesk connector, you can configure Enterprise Manager to create, update, or close a ticket for any event created in Enterprise Manager. The ticket generated by the connector contains the relevant information about the Enterprise Manager incident, including a link to the Enterprise Manager console to enable helpdesk analysts leverage Enterprise Manager's diagnostic and resolution features to resolve the incident. In Enterprise Manager, the ticket ID, ticket status, and link to the third-party ticketing system is shown in the context of the incident. This provides Enterprise Manager administrators with ticket status information and an easy way to quickly access the ticket.

Available connectors include:

- BMC Remedy Service Desk Connector
- HP Service Manager Connector
- CA Service Desk Connector
- HP Operations Manager Connector
- Microsoft Systems Center Operations Manager Connector
- IBM Tivoli Enterprise Console Connector
- IBM Tivoli Netcool/OMNIbus Connector

For more information about Oracle-built connectors, see the Enterprise Manager Extensions Exchange.


<http://www.oracle.com/technetwork/oem/extensions/index.html>

1.7 Accessing Monitoring Information

Enterprise Manager provides multiple ways to access monitoring information. The primary focal point for incident management is the Incident Manager console, however Enterprise Manager also provides other ways to access monitoring information. The following figures show the various locations within Enterprise Manager that display target monitoring information. The following figure shows the

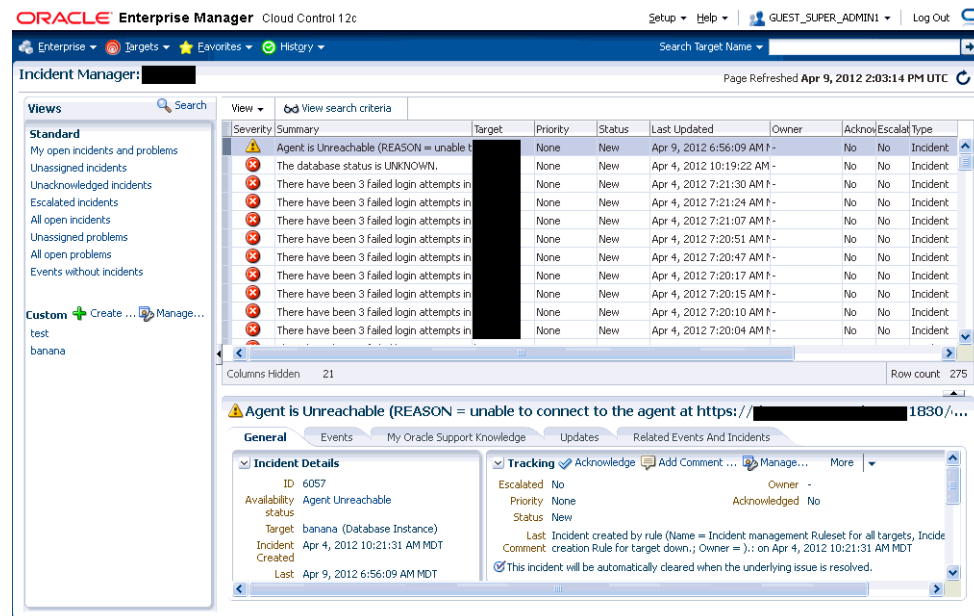
Enterprise Manager Overview page that conveniently displays target status rollup and rollup of incidents.

Figure 1–2 Enterprise Manager Console



The next figure shows the Incident Manager home page which displays incidents for a system or target.

Figure 1–3 Incident Manager (in context of a system or target)



Monitoring information is also display on target home pages. In the following figure, you can see target status as well as a rollup of incidents.

Figure 1–4 Target Home Pages

The screenshot displays the Oracle Enterprise Manager Cloud Control 12c interface. The top navigation bar includes 'Enterprise', 'Targets', 'Favorites', and 'History'. The main content area is divided into several sections:

- Summary:** Displays target status including Up Time (17 days, 17 hrs), Version (11.2.0.3.0), Load (0.00 average active sessions), Total Sessions (71), Last Backup (Mar 5, 2012 6:14:20 AM), Available Space (0.13 GB), and Total SGA (999.51 MB).
- Diagnostics:** Shows ADDM Findings (0) and Incidents (0).
- Incidents and Problems:** A table listing incidents and problems. The table has columns: Summary, Target, SeveStatus, Escalation level, Type, and Time since last update. The table shows three incidents: 'Problem: ORA 600 [4]', 'Problem: ORA 603', and 'Compliance score 51%'. The 'SeveStatus' column shows 'New' for all incidents. The 'Escalation level' column shows '-' for all incidents. The 'Type' column shows 'Problem' for the first two and 'Incident' for the last one. The 'Time since last update' column shows '7 days 7 hours' for the first two and '21 days 4 hours' for the last one.
- Compliance Summary:** Displays Compliance Standards and a table with columns: Name, Average Score, and No data to display.

The bottom navigation bar includes 'Performance', 'SQL Monitor - Last Hour', and 'Resources'.

Discovering and Monitoring Targets

This chapter contains the following sections:

- [About Discovering, Promoting, and Monitoring Hosts and Targets](#)
- [Discovery and Monitoring in Enterprise Manager Lifecycle](#)
- [Overview of the Discovery and Monitoring Process](#)
- [Automatically Discovering and Promoting Targets](#)
- [Manually Adding Targets](#)

2.1 About Discovering, Promoting, and Monitoring Hosts and Targets

Enterprise Manager Cloud Control enables you to manage and monitor Oracle components such as Oracle Database or an Oracle Weblogic Server domain running on your system using a single console. These Oracle components are also known as *targets*. Targets that are managed and monitored by Enterprise Manager Cloud Control are known as *managed targets*.

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

Enterprise Manager Cloud Control enables you to either add a target manually without any discovery, or discover the target automatically and then promote it.

In the manual process, you explicitly add a specific host or Oracle component as a target to bring under management. The benefit of using this process is that you eliminate the need to consume resources on the Management Agent to perform discovery when it is not needed.

In automatic discovery, you enable a Management Agent running on the host to run an Enterprise Manager job that scans for unmanaged hosts. You then convert these unmanaged hosts to managed hosts by deploying Management Agents on these hosts, then you search for targets on these managed hosts, and finally you promote these targets to managed status. The benefit of using this process is that as new Oracle components are added to your infrastructure, they can be found and brought under management on a regularly-scheduled basis.

For a better understanding of this concept, refer to the following subsections:


- [What is Discovery?](#)
- [What is Promotion?](#)
- [What is Monitoring?](#)

2.1.1 What is Discovery?

Discovery refers to the process of identifying unmanaged hosts and targets in your environment. You can discover hosts and targets automatically or manually.

Figure 2–1 illustrates the discovery process.

Figure 2–1 *Discovery*



For further information on how to discover unmanaged hosts refer to [Step 1: Discovering Unmanaged Hosts Using Network Scan](#).


For further information on how to discover targets on managed hosts refer to [Step 3: Discovering Targets on Managed Hosts](#)

2.1.2 What is Promotion?

Promotion refers to the process of converting unmanaged hosts and targets to managed hosts so that they can be monitored in Enterprise Manager Cloud Control. Promotion involves installation or deployment of a Management Agent on the host, only if you are promoting host targets.

Note: Promoting a non host target such as a database target can result in pushing the plug-in to the Management Agent if the plug-in does not exist on the Management Agent.

Figure 2–2 illustrates the promotion process.

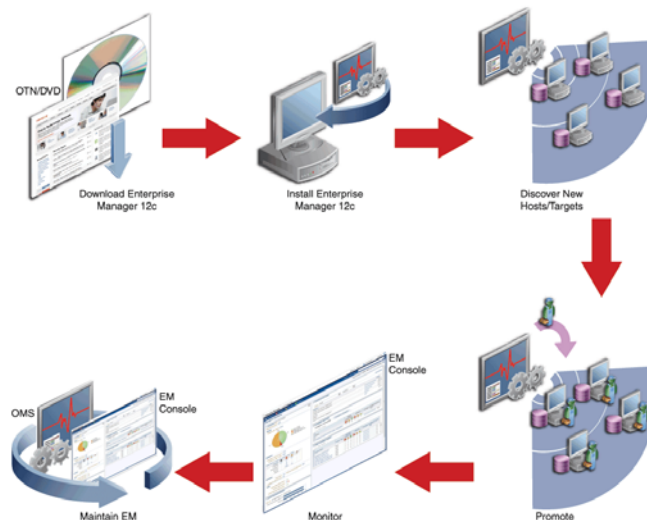
Figure 2–2 Promotion

2.1.3 What is Monitoring?

Monitoring refers to the process of gathering information and keeping track of activity, status, performance, and health of targets managed by Enterprise Manager Cloud Control on your host. A Management Agent deployed on the host in conjunction with plug-ins monitors every target in your environment.

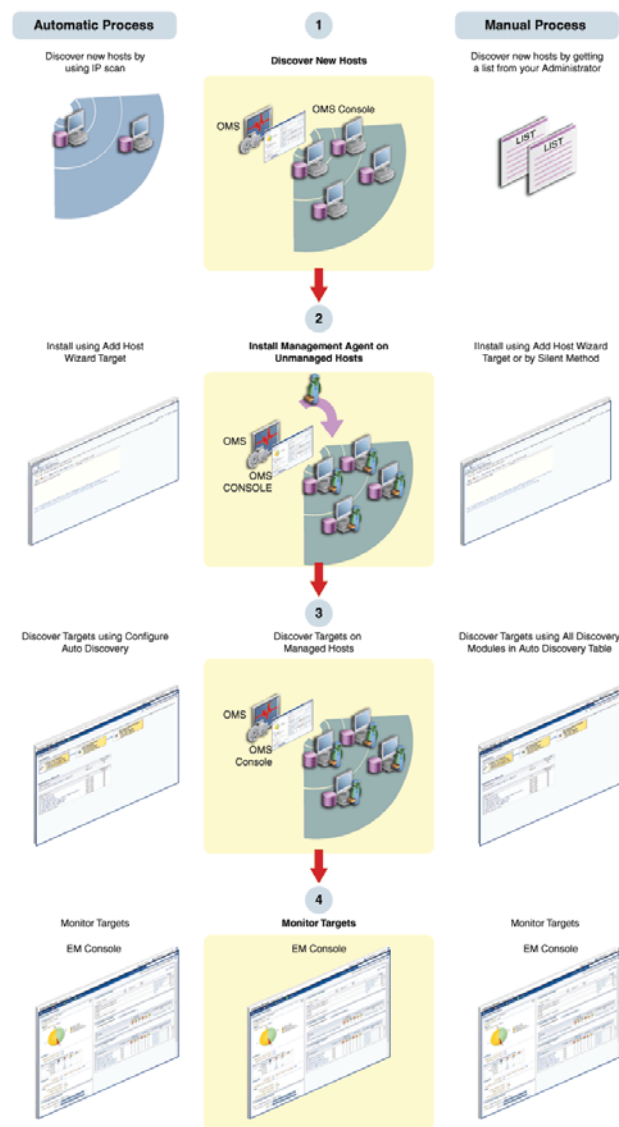
2.2 Discovery and Monitoring in Enterprise Manager Lifecycle

Figure 2–3 illustrates the lifecycle process of discovering and monitoring targets in Enterprise Manager Cloud Control.

Figure 2–3 Discovery and Monitoring in Enterprise Manager Lifecycle

2.3 Overview of the Discovery and Monitoring Process

Figure 2–4 illustrates the process of discovering and monitoring targets:

Figure 2–4 Discovery and Monitoring Process

2.4 Automatically Discovering and Promoting Targets

Automatic discovery refers to the process of scanning hosts for targets that can be managed and monitored by Enterprise Manager Cloud Control.

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

You can configure automatic discovery to set up a schedule for discovery, the target types to be discovered, and the hosts to scan for targets. Discovered hosts can then be promoted to managed target status, enabling Enterprise Manager Cloud Control to manage and monitor these targets. When new targets are added to your infrastructure, they can be found and brought under management on a regularly-scheduled basis.

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.

To automatically discover and monitor targets, follow these steps:

- [Prerequisites](#)
- [Step 1: Discovering Unmanaged Hosts Using Network Scan](#)
- [Step 2: Converting Unmanaged Hosts to Managed Hosts](#)
- [Step 3: Discovering Targets on Managed Hosts](#)
- [Step 4: Promoting Targets](#)

2.4.1 Prerequisites

- To run automatic host discovery (Nmap binary) on a Management Agent installed on a Solaris system, follow these steps:

1. Install GCC 3.4.3 or higher on the Solaris system.
2. Stop the Management Agent from the running.
3. Execute the following commands on the terminal session which you will use to start the Management Agent:

```
bash-2.03$ LD_LIBRARY_PATH=<directory path of gcc libraries>:$LD_LIBRARY_PATH
```

```
bash-2.03$ export LD_LIBRARY_PATH
```

Note: These steps ensure that Nmap refers to the GCC libraries while the Management Agent is running.

4. Start the Management Agent using the terminal session used for the previous step.

Note: Add LD_LIBRARY_PATH to your start up scripts so that this setting is retained after you reboot your system.

- To discover hosts from a Management Agent running on the following systems, follow the prerequisites stated in [Table 2-1](#).

Table 2-1 Prerequisites for discovering hosts

System	Prerequisites
Solaris 11 cluster (SPARC or x86-64)	<p>Create a static IPv4 address on interface net0:</p> <pre>ipadm create-addr -T static -a local=X.X.X.X/YY net0/ZZ</pre> <p>Here, X.X.X.X is a static IPv4 address, YY is the sub network mask, and ZZ is the specific identity for the net0 interface. The created static address will subsequently be identified by net0/ZZ.</p> <p>For example:</p> <pre>ipadm create-addr -T static -a local=10.134.108.101/24 net0/hh</pre>


Table 2–1 (Cont.) Prerequisites for discovering hosts

System	Prerequisites
SUSE Linux for System z™	<ol style="list-style-type: none"> Run the <code>QETH_OPTIONS='fake_ll=1'</code> option by adding it to the configuration file for the NIC present in the <code>/etc/sysconfig/hardware</code> directory. The name of the configuration file changes according to the NIC used. Contact your system administration for the name of the configuration file that your system uses. Restart your system for the changes to take effect.
RedHat Linux for System z	<ol style="list-style-type: none"> Run the <code>OPTIONS='fake_ll=1'</code> option by adding it to the configuration file for the NIC present in the <code>/etc/sysconfig/network-scripts</code> directory. The name of the configuration file changes according to the NIC used. Contact your system administration for the name of the configuration file that your system uses. Verify that the alias in the <code>/etc/modprobe.conf</code> file includes the following command: <code>alias eth0 geth</code> Restart the system for the changes to take effect.

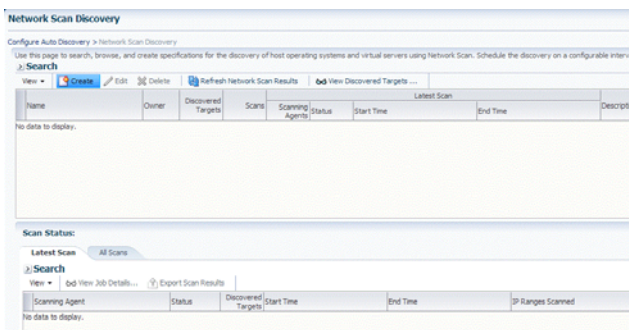
2.4.2 Step 1: Discovering Unmanaged Hosts Using Network Scan

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

- From the **Setup** menu, select **Add Target**, then select **Configure Auto Discovery**.
- On the Configure Auto Discovery page, click the **Configure** icon against Host and Oracle VM Manager in the Network Scan-based Auto Discovery table.



- On the Network Scan Discovery page, click **Create**. You will now 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.



- On the Network Scan Discovery: Create page, 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.

Network Scan Discovery: Create
Specify Network Scan Discovery to discover host operating systems and virtual servers. Schedule the discovery on a configurable interval.
☒ To perform Network Scan discovery, configure the scanning agent hosts with Sudo Privilege Delegation, and use credentials that have Run As 'root' attribute set.

* Name: Host Discovery Jun 19, 2013 3:11:54 AM PDT
Description:
Owner: SYSMAN

Network Scans
View: ☒ Add ☒ Remove ☒ Edit Nmap Scan Services and Ports
IP Ranges to Scan: No data to display. Scanning Agent:
Exclude IP Ranges:
Schedule:
Start: ☒ Immediately ☐ Later (GMT-08:00) Los Angeles - Pacific Time (PDT)
Repeat: ☐ Do not repeat
Credentials:
The discovery Network Scan is run as root. It is required that you set privileged host credentials or named credentials that use Sudo.
Credential: ☒ Preferred ☐ Named ☐ New
Preferred Credential Name: Privileged Host Credentials
Credential Details: Credentials will be determined at runtime.

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.

5. 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

6. 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.

7. 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.
8. Specify the following:
 - 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.

Click **Save and Submit Scan**.

9. 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**.

2.4.3 Step 2: Converting Unmanaged Hosts to Managed Hosts

To convert unmanaged hosts to managed hosts, follow these steps:

1. From the **Setup** menu, select **Add Target**, and then select **Auto Discovery Results**.
2. On the Auto Discovery Results page, select the **Network-scanned 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.
3. Select a host from 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.




Installing a Management Agent on an unmanaged host promotes the unmanaged host to managed target status, thereby converting the host to a managed host.


2.4.4 Step 3: Discovering Targets on Managed Hosts

To discover targets on managed hosts, follow these steps:

1. From the **Setup** menu, select **Add Target**, and then select **Configure Auto Discovery**.
2. On the Configure Auto Discovery page, click the **Configure** icon in the All Discovery Modules row in the Configure Auto Discovery table.



3. On the Target Discovery (Agent-based) page, 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.
4. To search for a specific Management Agent, click **Search**. The table lists all the Management Agents and filters the list based on what you search for.
5. Select the host in the table and click **Configure Discovery Modules**.

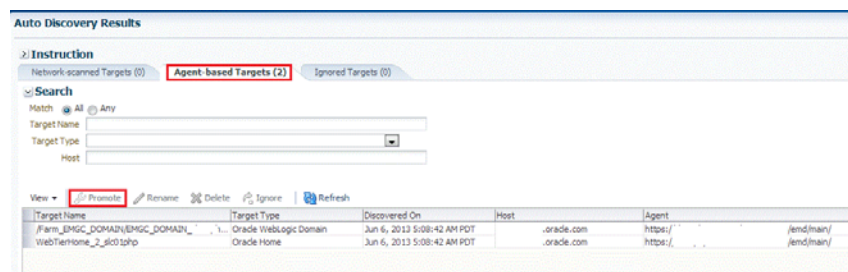


6. On the Configure Target Discovery page, set the schedule at which discovery will be run, in days. This schedule will be applied to all selected hosts. By default the discovery will run every 24 hours.
7. Select the target types you want to discover on the host. Note that you must supply search parameters 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 Domain: 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.
8. Click **OK** when finished. Target discovery has been configured on this host.
9. Repeat these steps for each additional host on which you want to configure discovery.
10. Click **Run Discovery Now** to discover targets immediately. The discovery will also run at the scheduled interval.

2.4.5 Step 4: Promoting Targets

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

1. 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**.
2. Select a target to promote, then click **Promote**. A wizard specific to the target type you are promoting opens. Supply the required values.



3. Click the **Agent-based Targets** tab. You can choose one or several targets to promote.
4. Note that you can optionally click **Ignore** for a discovered target. Ignoring a target puts it into a list of targets that you do not want to manage.

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. If you delete a target, it would be rediscovered the next time discovery runs.

5. 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:

- When you promote a discovered target to managed status, the plug-in required for the target is automatically deployed to the Management Agent, which monitors the host where the target has been discovered. For the plug-in to be deployed, the Management Agent must be secure. Therefore, before promoting the discovered targets to managed status, ensure that the Management Agent is secure. You can always unsecure it after the discovered target is promoted to managed status, that is, after the required plug-in is deployed.

To verify the secure status of a Management Agent, and to secure it if required, use any one of the following methods:

- From the **Setup** menu, select **Manage Cloud Control**, and then click **Agents**. Click the required Management Agent. Verify whether the Management Agent is secure. If it is not secure, from the **Agent** menu, click **Secure** to secure it.
- Run the following command to verify if the Management Agent is secure:

```
<EMSTATE>/bin/emctl status agent
```

If the Management Agent is secure, the Management Agent URL displayed in the output of the previous command is an HTTPS URL. However, if the Management Agent URL displayed is an HTTP URL, secure the Management Agent by running the following command:

```
<EMSTATE>/bin/emctl secure agent
```

- Enterprise Manager Cloud Control supports simultaneous promotion of multiple targets only for some target types. 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.
-

2.5 Manually Adding Targets

Enterprise Manager Cloud Control enables you to manually add hosts as well as a wide variety of Oracle software and components as managed targets. You can add targets directly by specifying monitoring properties. Adding targets in this way eliminates the need to consume resources on the Management 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.

You can manually add targets using the following three options:

- [Add Host Targets](#)
- [Add Targets Using Guided Process](#)
- [Add Targets Declaratively by Specifying Target Monitoring Properties](#)

2.5.1 Add Host Targets

To convert unmanaged hosts to managed hosts manually, you should manually install a Management Agent on each host. You can install a Management Agent on a managed host by using:

- Add Host Wizard

To convert an unmanaged host to a managed host, follow these steps:

1. From the **Setup** menu, select **Add Target**, then select **Add Targets manually**. Enterprise Manager Cloud Control displays the Add Targets Manually page.
2. Under the Add Targets Manually page, go to the Add Targets Manually subsection and choose **Add Host Targets**. Click **Add Host...**

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

- Silent Method

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

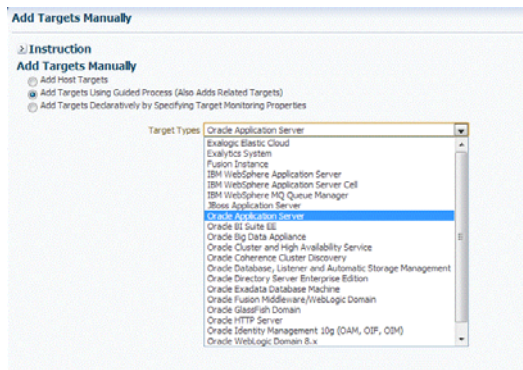
2.5.2 Add Targets Using Guided Process

To add targets manually to a specific Management Agent for monitoring, using a guided process, follow these steps:

Note: When you add a target using Guided Process, some scripts and automated processes are run that are particular for the target type that you select. You may have to input credentials in order to run the guided process.

1. From the **Setup** menu, select **Add Target**, then select **Add Targets Manually**. Enterprise Manager Cloud Control displays the Add Targets Manually page.

2. On the Add Targets Manually page, select **Add Targets Using Guided Process (Also Adds Related Targets)**.
3. Choose one of the target types to add from the Target Types list, such as **Exalogic Elastic Cloud**, **Oracle Cluster and High Availability Service**, or **Oracle WebLogic Domain**. Click **Add Using Guided Process...**



4. After you select the target type, a wizard specific to the target type guides you through the process of manually adding the target.

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

Note: When you manually add a non-host target to Enterprise Manager Cloud Control, the plug-in required for the target is automatically deployed to the Management Agent, which monitors the host where the non-host target exists. For the plug-in to be deployed, the Management Agent must be secure. Therefore, before manually adding a non-host target to Enterprise Manager Cloud Control, ensure that the Management Agent is secure. You can always unsecure it after the target is added to Enterprise Manager Cloud Control, that is, after the required plug-in is deployed.

To verify the secure status of a Management Agent, and to secure it if required, use any one of the following methods:

- From the **Setup** menu, select **Manage Cloud Control** and then, click **Agents**. Click the required Management Agent. Verify whether the Management Agent is secure. If it is not secure, from the **Agent** menu, click **Secure** to secure it.
- Run the following command to verify if the Management Agent is secure:

```
<EMSTATE>/bin/emctl status agent
```


If the Management Agent is secure, the Management Agent URL displayed in the output of the previous command is an HTTPS URL. However, if the Management Agent URL displayed is an HTTP URL, secure the Management Agent by running the following command:

```
<EMSTATE>/bin/emctl secure agent
```

2.5.3 Add Targets Declaratively by Specifying Target Monitoring Properties

To manually add a target on a managed host, follow these steps:

1. From the **Setup** menu, select **Add Target**, then select **Add Targets Manually**. Enterprise Manager Cloud Control displays the Add Targets Manually page.
2. On the Add Targets Manually page, select **Add Targets Declaratively by Specifying Target Monitoring Properties**.



3. Choose one of the target types to add from the Target Types list, such as **ADF Business Components for Java**, **Cluster Database**, or **Oracle HTTP Server**.
4. Specify the Management Agent that will be used to monitor the target, or click on the Search icon to search for and select the Management Agent. Click **Add Manually...**
5. After you select the target type, a wizard specific to the target type guides you through the process of manually adding the target.

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

Note: When you manually add a non-host target to Enterprise Manager Cloud Control, the plug-in required for the target is automatically deployed to the Management Agent, which monitors the host where the non-host target exists. For the plug-in to be deployed, the Management Agent must be secure. Therefore, before manually adding a non-host target to Enterprise Manager Cloud Control, ensure that the Management Agent is secure. You can always unsecure it after the target is added to Enterprise Manager Cloud Control, that is, after the required plug-in is deployed.

To verify the secure status of a Management Agent, and to secure it if required, use any one of the following methods:

- From the **Setup** menu, select **Manage Cloud Control**, and then, click **Agents**. Click the required Management Agent. Verify whether the Management Agent is secure. If it is not secure, from the **Agent** menu, click **Secure** to secure it.
- Run the following command to verify if the Management Agent is secure:

```
<EMSTATE>/bin/emctl status agent
```

If the Management Agent is secure, the Management Agent URL displayed in the output of the previous command is an HTTPS URL. However, if the Management Agent URL displayed is an HTTP URL, secure the Management Agent by running the following command:

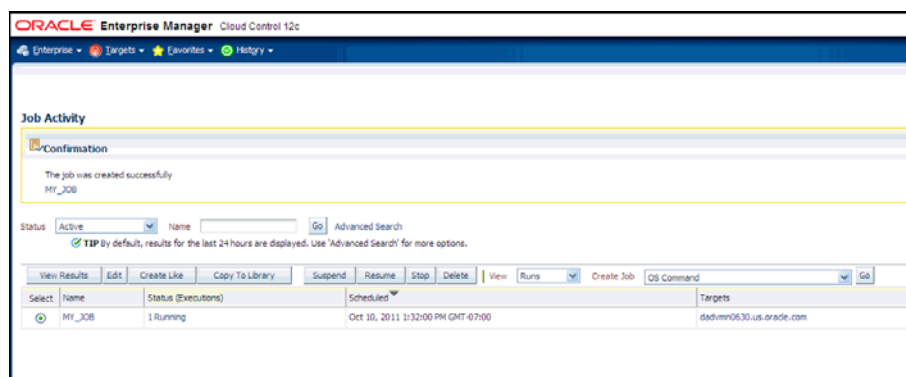
```
<EMSTATE>/bin/emctl secure agent
```

2.6 Discovering and Promoting Oracle Homes

When you deploy an Oracle software component outside of the deployment procedures provided by Enterprise Manager, the Oracle home is not automatically discovered and promoted as targets. You will have to manually discover and promote the Oracle home target.

To discover and promote an Oracle home target, follow these steps:

1. From the **Enterprise** menu, select **Job**, and then select **Library**.
2. On the Job Library page, from the drop-down list in the table, select **Discover Promote Oracle Home Target**.



Click **Go**.

- On the 'Create Discover Promote Oracle Home Target' Library Job page, in the General tab, specify the name of the discovery.

For example: OHDiscovery

You can optionally add a description for the discovery.

Click **Add**.

The screenshot shows the 'Create Discover Promote Oracle Home Target' Library Job page with the 'General' tab selected. The 'Name' field is highlighted with a red box and contains the text 'OH discovery'. Below it, the 'Description' field is empty. The 'Target Type' is set to 'Host'. In the 'Target' section, the 'Add' button is highlighted with a red box. Below the 'Add' button, there is a table with one row containing the text 'No targets are currently selected.'

- In the Search and Select: Target dialog box, select Target Type as **Host**, and then select all the host targets listed by clicking **Select All**.

Click **Select**.

- On the 'Create Discover Promote Oracle Home Target' Library Job page, the host targets that you selected are displayed in the table. Click **Save to Library**.
- On the Job Library page, you will receive a confirmation that the library job was created successfully.

From the table, select the library job that you have created, and click **Submit**.

The screenshot shows the 'Job Library' page. At the top, there is a confirmation message: 'Confirmation: The library job was created successfully. OH DISCOVERY'. Below the message, there is a table with columns 'Name', 'Owner', and 'Go'. The 'Submit' button is highlighted with a red box. Below the table, there is a table with columns 'Name', 'Job Type', and 'Go'. The row for 'OH DISCOVERY' is highlighted with a red box.

- On the Submit page, specify a name for the discovery.

For example: OHDiscovery.1

You can optionally add a description.

Select all the host targets by clicking **Select All**, and then click **Submit**.

When the discovery is successful, a confirmation is displayed on the Job Library page.

Using Incident Management

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

This chapter covers the following:

- [Events, Incidents, and Problems](#)
- [Key Concepts: Events, Incidents, and Problems](#)
- [Setting Up Your Incident Management Environment](#)
- [Working with Incidents](#)
- [Advanced Topics](#)
- [Moving from Enterprise Manager 10/11g to 12c](#)

3.1 Events, Incidents, and Problems

Enterprise Manager Cloud Control 12c greatly expands target monitoring and management capability beyond previous releases by focusing on what is important from a broader management perspective rather than the discrete events that may point to the same underlying issue. Enterprise Manager exposes three levels of management granularity that when combined, provide complete monitoring/management coverage of your environment.

3.1.1 Event Management

Intuitively, you monitor for specific events in your monitored environment. An event is a significant occurrence on a managed target that typically indicates something has occurred outside normal operating conditions--they provide a uniform way to indicate that something of interest has occurred in an environment managed by Enterprise Manager. Examples of events are:

- Metric Alerts
- Compliance Violations
- Job Events
- Availability Alerts

3.1.2 Incident Management

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). 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). However, managing by incident also allows you to address more complex situations where the subset of events you are interested in are related and may indicate a higher level issue needs to be addressed as a single issue and not as individual events: A cluster of events by themselves may indicate a minor administrative issue, but when viewed together may signify a larger problem that can potentially consist of events from multiple domains/layers of your monitored infrastructure.

For example, you are monitoring a host. If you want to monitor 'load' being placed on one or more hosts you might be interested in events such as CPU utilization, memory utilization, and swap utilization exceeding acceptable metric thresholds. Individually, these events may or may not indicate an issue with the host, but together, these events form an incident indicating extreme load is being placed on a monitored host.

Incidents represent the larger service disruptions that may impact your business instead of discrete events. Managing by incidents, therefore, allows you to monitor for complex operational issues that may affect multiple domains that may impact your business. These incidents typically need to be tracked, assigned to appropriate personnel, and resolved as quickly as possible. You can effectively implement a centralized monitoring that consolidates monitoring information and more effectively allocate resource across your ecosystem to resolve or prevent issues from occurring. The end result is better implementation of your business processes that in turn lead to better performance of your IT resources.

3.1.3 Problem Management

Problem management involves the functionality that helps track the underlying root causes of incidents. Once the immediate service disruptions represented by incidents are resolved, you can then progress to understanding and resolving the underlying root cause of the issue.

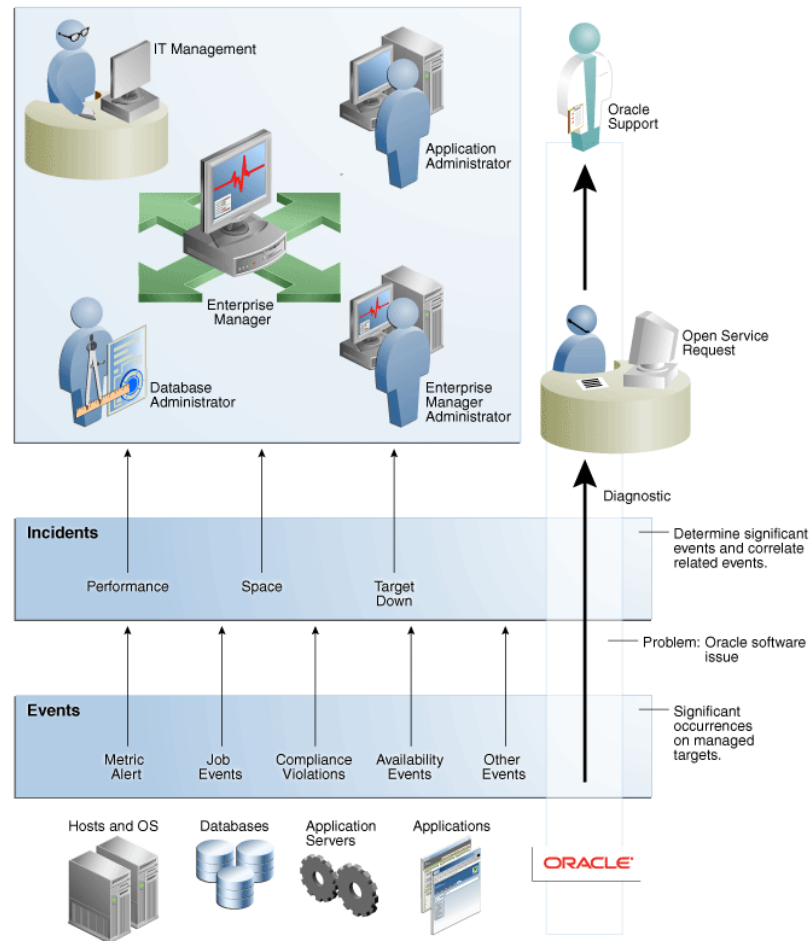
For the current release, problems pertain to the diagnostic incidents and problems stored in Automatic Diagnostic Repository (ADR), which are automatically raised by Oracle software when it encounters critical errors in the software. When problems are raised for Oracle software, Oracle has determined that the recommended recourse is to open a Service Request (SR), send support the diagnostic logs, and eventually provide a solution from Oracle. A problem represents the underlying root cause of a set of incidents. Enterprise Manager provides features to track and manage the lifecycle of a problem.

3.1.4 Summing Up

- **Event:** A significant occurrence of interest on a target that has been detected by Enterprise Manager.
Goal: Ensure that your environment is monitored.
- **Incident:** A set of significant events or combination of related events that pertain to the same issue.
Goal: Ensure that service disruptions are either avoided or resolved quickly.
- **Problems:** The underlying root cause of incidents. Currently, this represents critical errors in Oracle software that represents the underlying root cause of diagnostic incidents.
Goal: Ensure underlying root causes of issues are resolved to avoid future occurrence of issues.

Events, incidents, and problems work in concert to allow you to manage your complete IT ecosystem both effectively and efficiently. The following illustration summarizes how they work within your managed environment.

Figure 3–1 Event/Incident/Problem Flow



The following sections delve into events, incidents, and problems in more detail.

3.1.5 Working with Events, Incidents, and Problems

Effectively managing your environment via events, incidents, and problems involves planning, setup, implementation, and finally usage. Ultimately, the goal is to manage your environment through a series of incidents that address the critical IT components of your organization that affect corporate operations and ultimately profitability.

This chapter covers operational specifics about Enterprise Manager's incident management feature. However, as part of the larger Enterprise Manager core framework, other areas of the monitoring infrastructure are involved. The following topics link off to appropriate areas of this chapter as well as other areas of Enterprise Manager documentation.

- **Setting Up Your Incident Management Environment**
 - [Setting Up Your Monitoring Infrastructure](#)
 - [Setting Up Notifications](#)

- Setting Up Administrators and Privileges
- Setting Up Rule Sets
- **Working with Incidents**
 - Finding What Needs to be Worked On
 - Setting Up Custom Views
 - Responding and Working on a Simple Incident
 - Searching My Oracle Support Knowledge
 - Open Service Request (Problems-only)
 - Suppressing Incidents and Problems
 - Managing Workload Distribution of Incidents
 - Reviewing Events on a Periodic Basis
- **Advanced Topics**
 - Defining Custom Incident Statuses
 - Clearing Stateless Alerts for Metric Alert Event Types
 - User-reported Events
 - Additional Rule Applications
 - Event Prioritization
- **Moving from Enterprise Manager 10/11g to 12c**

3.2 Key Concepts: Events, Incidents, and Problems

As discussed previously, Enterprise Manager Cloud Control 12c lets you focus on what is important from a broader monitoring/management perspective via events, incidents, and problems. This section covers these concepts in greater detail.

3.2.1 Events

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

Existing Enterprise Manager customers may be familiar with metric alerts and metric collection errors. For Enterprise Manager 12c, metric alerts are a type of event, one of many different event types. The notion of an event unifies the different exception conditions that are detected by Enterprise Manager, such as monitoring issues or compliance issues, into a common concept. It is backed by a consistent and uniform set of event management capabilities that can indicate something of interest has occurred in a datacenter managed by Enterprise Manager.

All events have the following attributes:

Table 3–1 Event Attributes

Attribute	Description
Type	Type of event that is being reported. All events of a specific type share the same set of attributes that describe the exact nature of the problem. For example, Metric Alert, Compliance Standard Score Violation, or Job Status Change.
Severity	Event severity. For example, Fatal, Warning, or Critical.
Internal Name	An internal name that describes the nature of the event and can be used to search for events. For example, you can search for all <i>tablespacePctUsed</i> events.
Entity on which the event is raised.	An event can be raised on a target, a non-target source object (such as a job) or be related to a target and a non-target source object. Note: This attribute is important when determining what privileges are required to manage the event.
Message	Informational text associated with the event.
Reported Date	Time the event was reported.
Category	Functional or operational classification for an event. Available Categories: <ul style="list-style-type: none"> ■ Availability ■ Business ■ Capacity ■ Configuration ■ Diagnostics ■ Error ■ Fault ■ Jobs ■ Load ■ Performance ■ Security
Causal Analysis Update	Used for Root Cause Analysis of target down events. Possible Values: Root Cause or Symptom

Event Types







The *type* of an event defines the structure and payload of an event and provides the details of the condition it is describing. For example, a metric alert raised by threshold violation has a specific payload whereas a job state change has a different structure. As shown in the following table, the range of events types greatly expands Enterprise Manager's monitoring flexibility.

Event Type	Description
Target Availability	The Target Availability Event represents a target's availability status (Example: Up, Down, Agent Unreachable, or Blackout).

Event Type	Description
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.
Job Status Change	<p>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.</p> <p>Note: 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, from the Setup menu, select Incidents, and then Job Events.</p>
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).
Service Level Agreement Alert	These events are generated when a service level or service level objective is violated for a service. occurs for a Service Level Agreement or a Service Level Objective.
User-reported	These events are created by end-users.
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.
Application Performance Management KPI Alert	An Application Performance Management (APM) Key Performance Indicator (KPI) alert event is generated when a KPI violation alert occurs for a metric on an APM managed entity associated with a Business Application target.
JVM Diagnostics Threshold Violation	A JVM Diagnostics event is raised when a JVM metric exceeds its threshold value on a Java Virtual Machine target.

Event Severity

The severity of an event indicates the criticality of a specific issue. The following table shows the various event severity levels along with the associated icon.

Icon	Severity	Description
	Fatal	Corresponding service is no longer available. For example, a 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 can be used, for example, to report Oracle best practice violations.
	Clear	Conditions that raised the event have been resolved.
	Informational	<p>A specific condition has just occurred but does not require any remedial action.</p> <p>Events with an informational severity:</p> <ul style="list-style-type: none"> do not appear in the incident management UI. cannot create incidents. are not stored within Enterprise Manager.

3.2.2 Incidents

While events indicate issues requiring attention in your managed environment, it is more efficient to work on a collective subset of related events as a single unit of work--you can work on different events representing the same issue or you can work on one incident containing multiple space-related events. For example, you have multiple space events from various targets that indicate you are running low on space. Instead of managing numerous discrete events, you can more efficiently manage a smaller set of incidents.

An incident is a significant event or set of related significant events that need to be managed because it can potentially impact your business applications. These incidents typically need to be tracked, assigned to appropriate personnel, and resolved as quickly as possible. You perform these incident management operations through Incident Manager, an intuitive UI within Enterprise Manager.

Incident Manager 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 3.2.5, "Incident Manager"](#) for more information about this UI.

3.2.2.1 Working with Incidents

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

- Assign incident ownership.
- Track the incident resolution status.
- Set incident priority.
- Set incident escalation level.
- Ability to provide a manual summary.
- Ability to add user comments.
- Ability to suppress/unsuppress
- Ability to manually clear the incident.
- Ability to create a ticket manually.

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.

Incident Status

The lifecycle of an incident within an organization is typically determined by two pieces of information: The current resolution state of the incident (Incident Status) and how important it is to resolve the incident relative to other incidents (Priority). As key incident attributes, the following options are available:

- New
- Work in Progress
- Closed
- Resolved

You can define additional statuses if the default options are not adequate. In addition, you can change labels using the Enterprise Manager Command Line Interface (EM CLI). See **Advanced Topics** for more information.

Priority

By changing the priority, you can escalate the incident and perform operations such as assigning it to a specific IT operator or notifying upper-management. The following priority options are available:

- None
- Low
- Medium
- High
- Very High
- Urgent

Priority is often based on simple business rules determined by the business impact and the urgency of resolution.

Incident Attributes

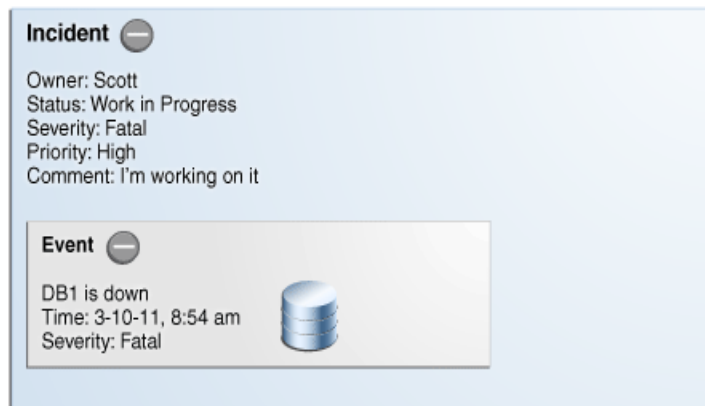
Every incident possesses attributes that provide information as identification, status for tracking, and ownership. The following table lists available incident attributes.

Incident Attribute	Definition
Escalated	<p>An escalation level signifying a escalation to raise the level of attention on the incident from your organization's IT or management hierarchy.</p> <p>Available escalation levels:</p> <ul style="list-style-type: none"> ■ None (Not escalated) ■ Level 1 through Level 5
Category	<p>Operational or organizational classification for an incident. Incidents (and events) can have multiple categories.</p> <p>Categories for all events within an incident are aggregated.</p> <p>Available Categories:</p> <ul style="list-style-type: none"> ■ Availability ■ Business ■ Capacity ■ Configuration ■ Diagnostics ■ Error ■ Fault ■ Jobs ■ Load ■ Performance ■ Security
Summary	<p>An intuitive message indicating what the incident is about. By default, the incident summary is pulled from the message of the last event of the incident, however, this message can be changed to a fixed summary by any administrator working on the incident.</p>
Incident Created	<p>Date and time the incident was created.</p>
Last Updated	<p>Date and time the incident was last updated or when the incident was closed.</p>
Severity	<p>Severity is based on the worst severity of the events in the incident. For example, Fatal, Warning, or Critical.</p>
Source	<p>Source entities of the incident.</p>
Priority	<p>Priority Values</p> <ul style="list-style-type: none"> ■ None (Default) ■ Low ■ Medium ■ High ■ Very High ■ Urgent

Incident Attribute	Definition
Status	<p>Incident Status.</p> <ul style="list-style-type: none"> ■ New (Default) ■ Work in Progress ■ Closed (Terminal state when the incident is closed. See below for more information.) ■ Resolved <p>You can define additional statuses if the default options are not adequate. In addition, you can change labels using the Enterprise Manager Command Line Interface (EM CLI).</p> <p>Closed Status: Enterprise Manager automatically sets the status to closed when an incident severity is cleared--administrators do not manually select the Closed status. The incident severity is set to Clear when all of the events contained within the incident have been cleared. Typically the Agent sets the Clear severity, as would be the case when a metric alert value falls below a severity threshold. If an event or incident supports manual clearing, then the Clear option will be shown in the Incident Manager UI. Once an incident has been cleared by an administrator or by Enterprise Manager, only then will Enterprise Manager set the status to Closed.</p> <p>If you do not see the option to clear the incident in the UI, this means Enterprise Manager will automatically set the status to Clear if it detects the monitored condition no longer holds true. For example, you want to indicate that an incident has been fixed. You can set the status to Resolved and Enterprise Manager will set the status to Closed when it clears the severity.</p>
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.
Acknowledged	<p>Indicates that a user has accepted ownership of an incident or problem. Available options: Yes or No.</p> <p>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, but can be assigned to another user. Acknowledging an incident stops any repeat notifications for that incident.</p>
Causal Analysis Update	<p>Used for Root Cause Analysis of target down incidents.</p> <p>Possible Values: Root Cause or Symptom</p>

3.2.2.2 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 all incident management functionality required to track and manage its resolution.

Figure 3–2 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. When the event is newly generated, there is no ownership or status. An incident is opened that can be updated manually or by automated rules to set owners, status, as well as other attributes. In the example, the owner/administrator Scott is currently working to resolve the issue.

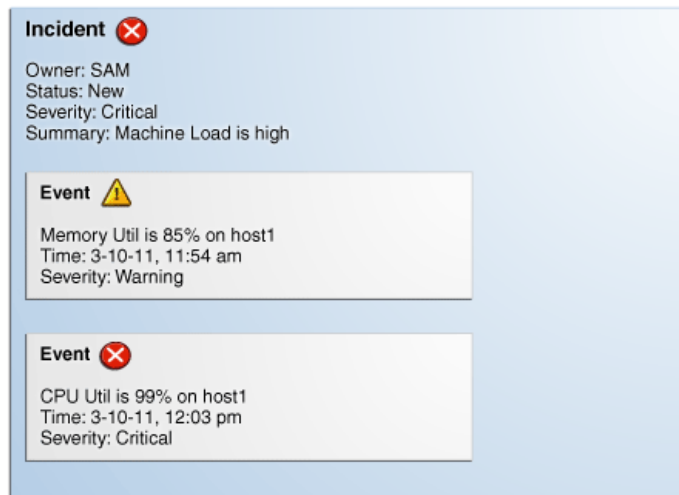
The incident severity is currently Fatal as the incident inherits the worst severity of all the events within incident. In this case there is only one event associated with the incident so the severity is Fatal.

3.2.2.3 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.

Note: Multi-event incidents are not automatically generated. An administrator must manually create them.

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. See "[Creating an Incident Manually](#)" on page 3-46 for more information. Together, these events form an incident indicating extreme load is being placed on a monitored host. The following figure illustrates this example.

Figure 3–3 Incident with Multiple Events

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 last event of the incident, however, this message can be changed by any administrator working on the incident.

Because administrators are interested in overall machine load, administrator Sam has manually 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.

3.2.2.4 How are Incidents Created?

Incidents are most commonly created automatically through rules and rule sets (user-defined instructions that tell Incident Manager how to handle specific events when they occur). As shown in the preceding examples, incidents can also be created manually. Once an incident is raised, its severity is inherited from the worst severity of all events within the incident. The latest event *Message*, by default, becomes the *Incident Summary*. Incidents can also be created manually. See ["Creating an Incident Manually"](#) on page 3-46 for more information.

3.2.3 Problems

For Enterprise Manager 12c, problems focus on the diagnostic incidents and problem diagnostic incidents/problems stored in 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 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 can use Support Workbench to package the diagnostic details gathered in ADR and open SR. Users should then manage the problems in Incident Manager. Access to Support Workbench functionality is available through Incident Manager (**Guided Resolution** area) in context of the problem.

3.2.4 Rule Sets

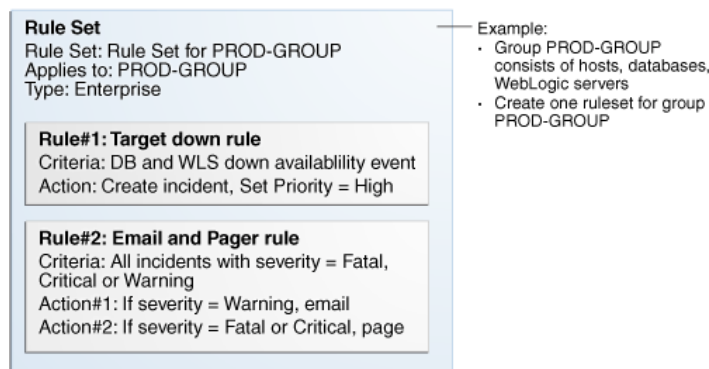
Incident rules and rule sets automate actions related to events, incidents and problems. They can automate the creation of incidents based on important events, perform notification actions such as sending e-mail or opening helpdesk tickets, or perform operations to manage the incident workflow lifecycle such as changing incident ownership, priority, or escalation level.

With previous versions of Enterprise Manager, you used notification rules to choose the individual targets and conditions for which you want to perform actions or receive notifications (send e-mail, page, open a helpdesk ticket) from Enterprise Manager. For Enterprise Manager 12c, the concept and function of notification rules has been replaced with incident 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 creating incidents, updating incidents and problems. The actions can also be conditional in nature. For example, a rule action can be defined to page a user when an incident severity is critical or just send e-mail if it is warning.
- **Rule Set:** An incident rule set is a collection of rules that apply to a common set of objects such as targets (hosts, databases, groups), jobs, metric extensions, or self updates and take appropriate actions to automate the business processes underlying event, incident and problem management.

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 they can be reordered from the Incident Rules 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 (this is also a best practice). 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. See ["Rule Set Types"](#) on page 3-15

3.2.4.1 Out-of-Box Rule Sets

Enterprise Manager provides out-of-box rule sets for incident creation and event clearing based on typical scenarios. Out-of-box rule sets cannot be edited or deleted, however, they can be disabled. As a best practice, you should create your own copies of out-of-box rule sets and then subscribe to the rule set copies rather than subscribing directly to the out-of-box rule sets. Effectively, you are making a copy of the rule set and changing the target criteria to fit your enterprise needs by selecting an appropriate group of targets (preferably an administration group).

Please note that out-of-box rule set definitions and actions they perform can be changed by Oracle at any time and will be applied during patching or software upgrade.

Regular Enterprise Manager administrators are allowed to perform the following operations on rule sets:

- Subscribe
- Subscribe for E-Mail notifications
- Unsubscribe
- Unsubscribe from E-Mail notifications

- Enable
- Disable

Note: Even though administrators can subscribe to a rule set, they will only receive notification from the targets for which they have at least the View Target privilege.

Enterprise Manager Super Administrators have the added ability to reorder the rule sets.

Enterprise rule sets are evaluated sequentially and may go through multiple passes as needed. When there is a change to the entity being processed - such as an incident being created for an event or an incident priority changing due to a rule - we rerun through all the rules from the beginning again until there are no matches. Any rule that is matched in a prior pass will not match again (to prevent infinite loops).

For example, when a new event, incident, or problem arises, the first rule set in the list is checked to see if any of its member rules apply and appropriate actions specified in those rules are taken. The second rule is then checked to see if its rules apply and so on. Private rule sets are only evaluated once all enterprise rule set evaluations are complete and in no particular order.

Important: Use caution when reordering rule sets as their order defines the event, incident, and problem handling workflow. Reordering rule sets without fully understanding the impact on your system can result in unintended actions being taken on incoming events, incidents, and problems.

3.2.4.2 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. However, if the rule set owner loses the *Create Enterprise Rule Set* system privilege at some future time, he can still edit or delete the rule set. Super Administrators can edit or delete any rule set.

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 by adding co-authors. Enterprise rule sets are visible to all administrators.

- **Private:** Used when an administrator wants to be notified about something he is monitoring but not as a standard business practice. 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 regardless of whether they have been granted the *Create Enterprise Rule Set* resource privilege. Oracle recommends that private rule sets be used only in rare or exceptional situations.

When a rule set performs actions, the privileges of the rule set creator are used. For example, a rule set owner/creator must have at least View Target privilege in order to

receive notifications and at least Manage Target Events privilege in order to update the incident. The exception is when a rule set sends a notification. In this case, the privileges of the user it is sent to is used.

3.2.4.3 Rules

Rules are instructions within a rule set that automate actions on incoming events or incidents or problems. 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 the 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.

The following table shows how rule criteria and actions determine rule application. In this rule operation example there are three rules which take actions on selected events and incidents. Within a rule set, 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.

Table 3–2 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 are unacknowledged for more than six hours		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 six hours, *Rule 3* applies and the incident escalation level is increased from None to Level 1. At this point, Enterprise Manager runs through all the rule sets and their rules from the beginning again.

3.2.4.3.1 Rule Application Each rule within a rule set applies to an event, incident OR problem. For each of these, you can choose rule application criteria such as:

- Apply the rule to incoming events or updated events only

- Apply the rule to critical events only.

Rules are applied to events, incidents, and problems according to criteria selected at the time of rule creation (or update). The following situations illustrate the methodology used to apply rules.

- If one of the rules creates a new incident in response to an incoming event, Enterprise Manager finishes matching the event to any further rules/rule sets. Once completed, Enterprise Manager then matches the newly created incident to all the rule sets from the beginning to see if any incident-specific rules match.
- If an incoming event is already associated with an incident (for example, a *Warning* event creates an incident and then a *Critical* event is generated for the same issue), Enterprise Manager applies all the matching rules to the event and then matches all rules to the incident.
- If, while applying a rule to an incident, changes are made to the incident (change priority, for example), Enterprise Manager stops rule application at that point and then re-applies the rules to the incident from the beginning. The conditional action that updated the incident will not be matched again in the same rule application cycle.

3.2.4.3.2 Rule Criteria

The following tables list selectable criteria for each type.

Table 3–3 Rule Criteria: Events

Criteria	Description
Type	Rule applies to a specific event type.
Severity	Rule applies to a specific event severity.
Category	Rule applies to a specific event category.
Target type	Rule applies to a specific target type.
Target Lifecycle Status	Rule applies to a specific lifecycle status for a target. Lifecycle status is a target property that specifies a target's operational status.
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.
Causal analysis update	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 applied 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 3–4 Rule Criteria: Incidents

Criteria	Description
Rules that created the incident	Rule applies to incidents raised by a specific rule.
Category	Rule applies to a specific incident category.
Target Type	Rule applies to a specific target type.
Target Lifecycle Status	Rule applies to a specific lifecycle status for a target. Lifecycle status is a target property that specifies a target's operational status.
Severity	Rule applies to a specific incident severity.
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.
Status	Rule applies when the incident status matches a selected incident status.
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 3–5 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.
Target Type	Rule applies to a specific target type.
Target Lifecycle Status	Rule applies to a specific lifecycle status for a target. Lifecycle status is a target property that specifies a target's operational status.
Acknowledged	Rule applies when the problem is acknowledged.
Owner	Rule applies for a specified problem owner.
Priority	Rule applies when problem priority matches a selected priority.

Table 3–5 (Cont.) Rule Criteria: Problems

Criteria	Description
Status	Rule applies when the problems matches a specific status.
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.

3.2.4.3.3 Rule Actions For each rule, Enterprise Manager allows you to define specific actions.

Some examples of the types of actions that a rule set 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 a 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.

The following table summarizes available actions for each rule application.

Table 3–6 Available Rule Actions

Action	Event	Incident	Problem
E-mail	Yes	Yes	Yes
Page	Yes	Yes	Yes
Advanced Notifications			
Send SNMP Trap	Yes	No	No
Run OS Command	Yes	Yes	Yes
Run PL/SQL Procedure	Yes	Yes	Yes
Create an Incident	Yes	No	No
Set Workflow Attributes	Yes	Yes	Yes

Note: Within an event rule, the workflow attributes of the associated incident can also be updated.

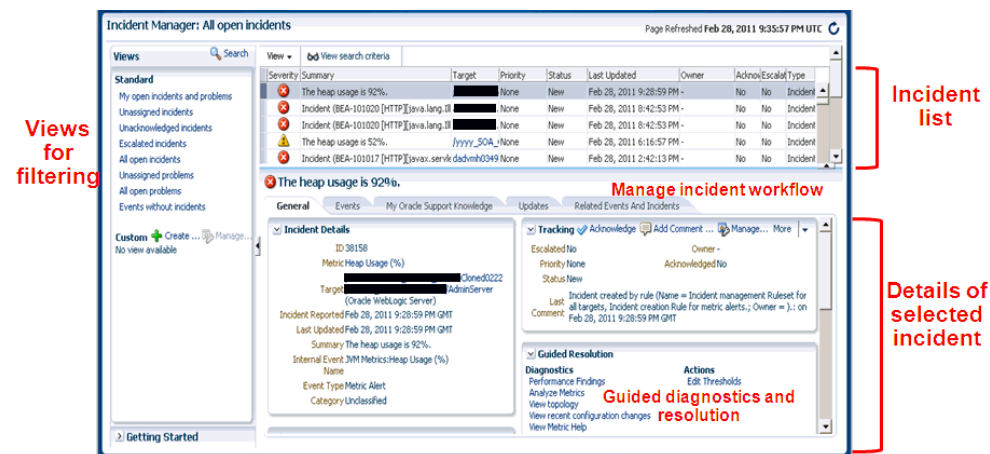
Table 3–6 (Cont.) Available Rule Actions

Action	Event	Incident	Problem
Create a Helpdesk Ticket	Yes	Yes	No
Note: Action performed indirectly by first creating an incident and then creating a ticket for the incident.			

3.2.5 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 custom views
- Search for specific incidents by properties such as target name, summary, status, or target lifecycle status
- Respond and work on an incident
- Manage incident lifecycle including assigning, acknowledging, tracking its status, prioritization, and escalation
- 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.

Figure 3–5 incident Manager

For example, you have 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 issues quickly. By drilling down on an open incident, you can access this information and modify it accordingly.

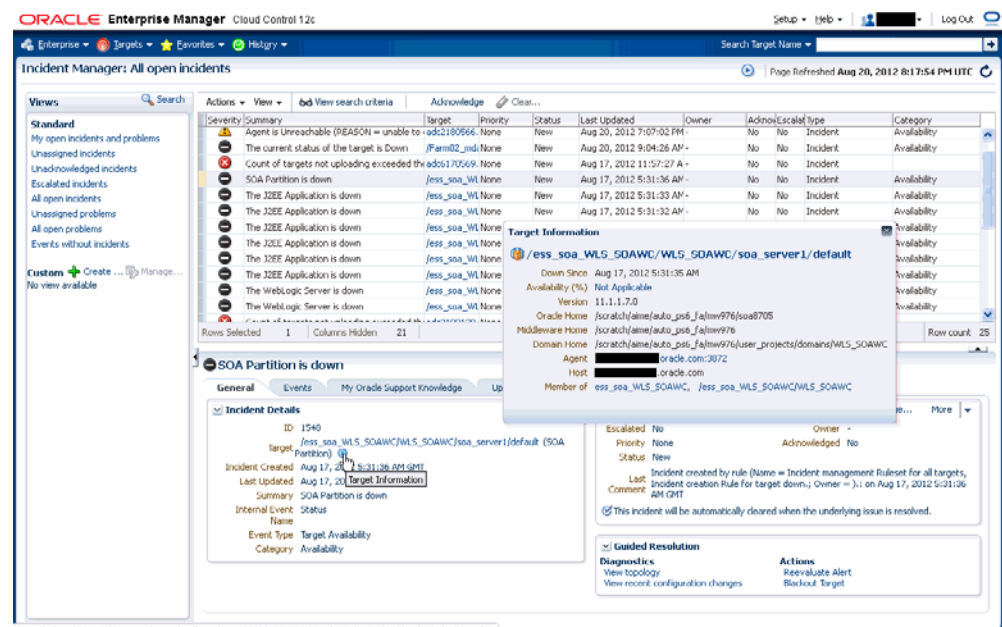
Displaying Target Information in the Context of an Incident

You can directly view information about a target for which an incident or event has been raised. The type of information shown varies depending on the target type.

To display in-context target information:

1. From the **Enterprise** menu, select **Monitoring** and then **Incident Manager**.
2. From the Incident Manager UI, choose an incident. Information pertaining to the incident displays.
3. From the Incident Details area of the General tab, click on the information icon "i" next to the *target*. Target information as it pertains to the incident displays. See [Figure 3–6](#)

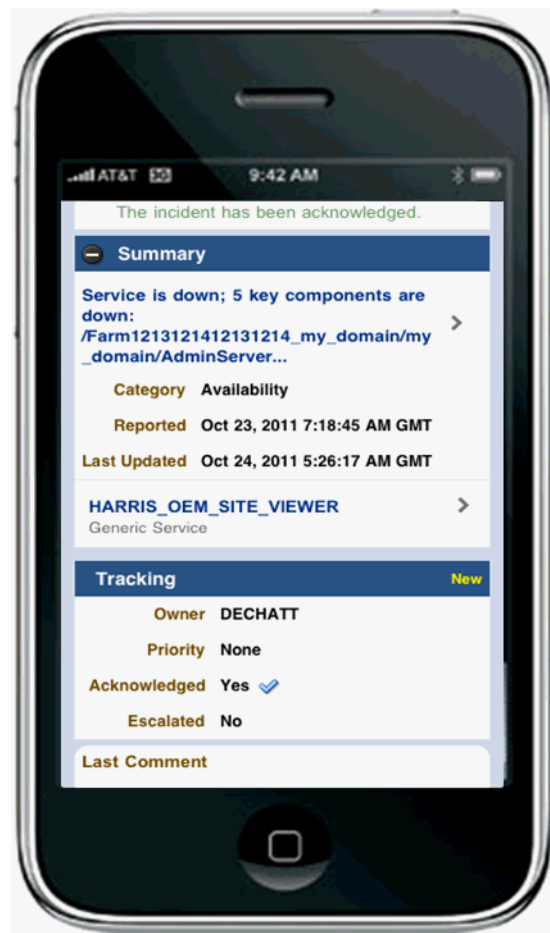
Figure 3–6 Target Information in Context of an Incident



Being able to display target information in this way provides you with more operational context about the targets on which the events and incidents are raised. This in turn helps you manage the lifecycle of the incident more efficiently.

Cloud Control Mobile

Also available is the mobile application Cloud Control Mobile, which lets you manage incidents and problems on the go using any iDevice to remotely connect to Enterprise Manager.

Figure 3–7 Cloud Control Mobile

For more information about this mobile application, see [Chapter 15, "Remote Access To Enterprise Manager"](#)

3.2.5.1 Views

Views let you work efficiently with incidents by allowing you to categorize and focus on only those incidents of interest. A view is a set of search criteria for filtering incidents and problems in the system. Incident Manager provides a set of predefined *standard* views that cover the most common event, incident, and problem search scenarios. In addition, Incident Manager also allows you to create your own custom views. Custom views cannot be shared with other users. For instructions on creating custom views, see [Setting Up Custom Views](#).

3.3 Setting Up Your Incident Management Environment

Before you can monitor and manage your environment using incidents, you must ensure that your monitoring environment is properly configured. Proper configuration consists of the following:

- [Setting Up Your Monitoring Infrastructure](#)
- [Setting Up Notifications](#)
- [Setting Up Administrators and Privileges](#)

3.3.1 Setting Up Your Monitoring Infrastructure

The first step in setting up your monitoring infrastructure is to determine which conditions need to be monitored and hence are the source of events. To prevent an inordinate number of extraneous events from being generated, thus reducing system and administrator overhead, you need to determine what is of interest to you and enable monitoring based on your requirements. You can leverage Enterprise Manager features such as Administrations Groups to automatically apply management settings such as monitoring settings or compliance standards when new targets are added to your monitored environment. This greatly simplifies the task of ensuring that events are raised only for those conditions in which you are interested. For more information, see [Chapter 6, "Using Administration Groups"](#).

Example: You want to ensure that the database containing your human resource information is available round the clock. One condition you are monitoring for is whether that database target is up or down. If it goes down, you want the appropriate person to be notified and have them resolve the problem as quickly as possible. Other conditions that you may want to monitor include performance threshold violations, any changes in application configuration files, or job failures. Working with events, you are monitoring and managing individual targets and issues directly related to those targets. For example, you monitor for individual database availability, individual host threshold violations such as CPU and I/O load, or perhaps the performance of a Web service.

In general, if you are primarily interested in availability and some key performance related metrics, you should use default monitoring templates and other template features to ensure the only those specific metrics are collected and events are raised only for those metrics.

Job Events: The status of a job can change throughout its lifecycle - from the time it is submitted to the time it has executed. For each of these job statuses, events can be raised to notify administrators of the status of the job.

As a general rule, events should be generated only for job status values that require administration attention. These job status values include Action Required and Problem status values such as Failed or Stopped. However, in order to avoid overloading the system with unnecessary events, job events are not enabled for any target by default. Hence, if you would like to generate events for jobs, you must:

1. Set the appropriate job status. You can use the default settings or modify them as required.
2. Specify the set of targets for which you would like job-related events to be generated.

You can perform these operations from the *Job Event Generation Criteria* page. From the **Setup** menu, choose **Incidents** and then **Job Events**.

3.3.1.1 Rule Set Development

Before creating incident rules/rule sets, the first step is to strategically determine when incidents should be created based on the business requirements of your organization. Important questions to consider are:

1. What events should create incidents? Which service disruptions need to be tracked and resolved by IT administrators?
2. Which administrators should be notified for incoming events or incidents?
3. Are any of the events or incidents being forwarded to external systems (such as a helpdesk ticketing system) ?

Once the exact business requirements are understood, you translate those into enterprise rule sets. Adhering to the following guidelines 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 targets into a smaller number of monitoring entities for the rule set. Groups should be composed of targets that have similar monitoring requirements including incident management and response.
- All the rules that apply to the same groups of targets should be consolidated into one rule set. You can create multiple rules that apply to the targets in the rule set. You can create rules for events specific to an event class, rules that apply to events of a specific event class and target type, or rules that apply to incidents on these targets.
- Leverage the execution order of rules within the rule set. Rule sets and rules within a rule set are executed in sequential order. Therefore, ensure that rules and rule sets are sequenced with that in mind.

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 3–7 Rule Usage Guidelines

Rule Usage	Application
Rules on Event	<p>To create incidents for the events managed in Enterprise Manager.</p> <p>To send notifications on events.</p> <p>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.</p> <p>Send events to third-party management systems.</p>
Rules on Incidents	<p>Automate management of incident workflow operations (assign owner, set priority, escalation levels..) and send notifications</p> <p>Create tickets based on incident conditions. For example, create a ticket if the incident is escalated to level 2.</p>
Rules on Problems	<p>Automate management of problem workflow operations (assign owner, set priority, escalation levels..) and send notifications</p>

Rule Set Example

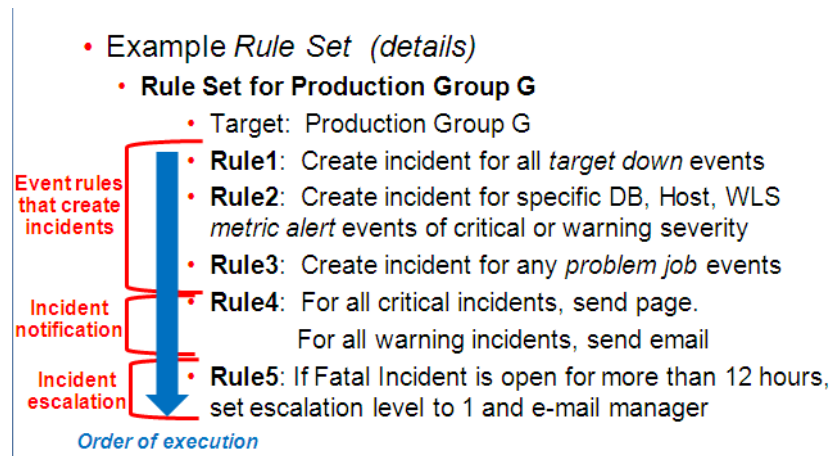
The following example illustrates many of the implementation guidelines just discussed. All targets have been consolidated into a single group, all rules that apply to group members are part of the same rule set, and the execution order of the rules has been set. In this example, the rule set applies 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, execution efficiency is achieved. As shown in this example, by using conditional actions that take different actions for the same set of events based on severity, it is easier to change the event selection criteria in the future without having to change multiple rules. **Note:** This assumes that the action requirements for all incidents (from rules 1 - 3) are the same.

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

Table 3–8 Example Rule Set for Production Group G

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

Table 3–8 (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 (%) Note: The warning and critical thresholds are defined in Metric and Policy settings, not from the rules UI. Host1 CPU Utilization (%) WLS1 Heap Usage (%)	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.

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

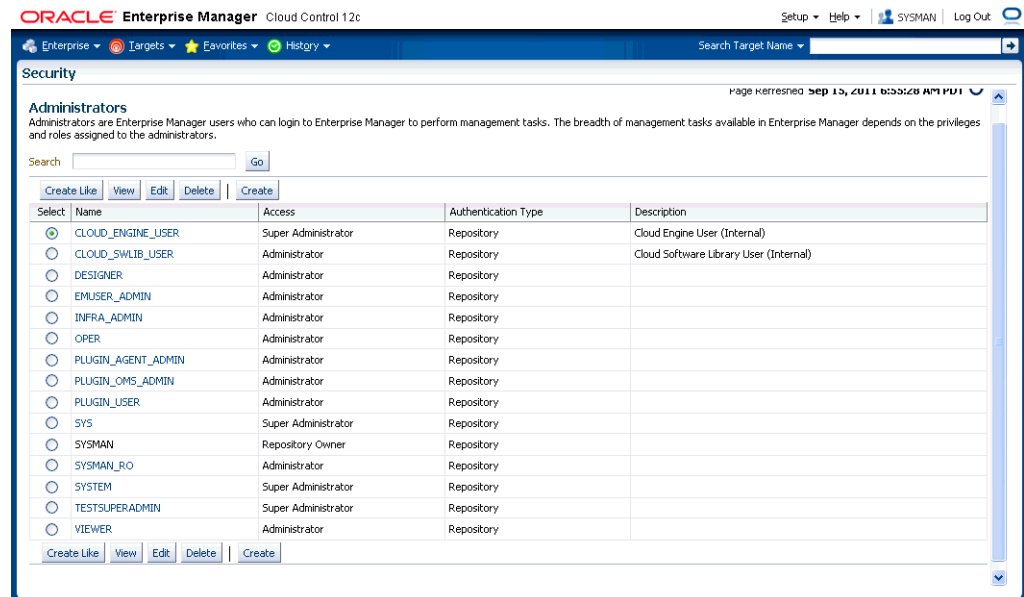
- User's Enterprise Manager account has notification preferences (e-mail and schedule). This is required not just for the administrator who is creating/editing a rule, but also for any user who is being notified as a result of the rule action.
- If you decide to use connectors, tickets, or advanced notifications, 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.

3.3.1.1.2 Setting Up Notifications After determining which events should be raised for your monitoring environment, you need to establish a comprehensive notification infrastructure for your enterprise by configuring Enterprise Manager to send out e-mail and or pages, setting up e-mail addresses for administrators and tagging them as e-mail/paging. In addition, depending on the needs of your organization, notification setup may involve configuring advanced notification methods such as OS scripts, PL/SQL procedures, or SNMP traps. For detailed information and setup instructions for Enterprise Manager notifications, see [Chapter 4, "Using Notifications"](#).

3.3.2 Setting Up Administrators and Privileges

This step involves defining the appropriate administrators (which includes assigning the proper privileges for security) and then setting up notification assignments based on their defined roles and domain ownership within your organization.

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



There are two types of administrators typically involved in incident management.

- **Business Rules Architect/Analyst:** Administrator who has a deep understanding of how the business works and translates this knowledge to operational rules. Once these rules have been deployed, the business architect uses their knowledge of the dynamic organization to keep these rules up-to-date.

In order to create or edit an enterprise rule set, the business architect/analyst must have been granted the *Create Enterprise Rule Set* privilege on the *Enterprise Rule Set* resource. The architect/analyst can share ownership of the rule sets with other administrators who may or may not have the *Create Enterprise Rule Set* privilege but are responsible for managing a specific rule set.

- **IT Operator/Manager:** The IT manager is responsible for day-to-day management of incident assignment. The IT operator is assigned the incidents and is responsible for their resolution.

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 grant privileges, from the **Setup** menu 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.

Resource Type	Description	Privilege Grants Applicable to all Resources	Number of Resources with Privilege Grants	Manage Privilege Grants
Access	Defines the access to different application in Enterprise Manager Cloud Control	-	NA	
Application Replay Entities	Application Replay Entities include captures, replay tasks, and replays.	-	NA	
Backup Configurations	Security Class for System Backup/Recovery Manager.	-	-	
Change Plan Security Class	Security behavior for Change Plans	-	-	
Chargeback and Consolidation	Extends Enterprise Manager feature to allow Chargeback and Consolidation of Targets based on configuration and resource usage	-	-	
Cloud Policy	Defines access privileges for Cloud Policies	-	-	
Cloud Policy Group	Defines access privileges for Cloud Policy Groups	-	-	
Cloud Self Service Portal	Defines the access privileges and roles for Cloud Self Service Portal.	-	NA	
Compliance Framework	Compliance Framework provides capability to define/customize/manage compliance frameworks, and compliance standards/rules and evaluate compliance of targets/systems with regards to business best practices for configuration/security/storage etc.	-	NA	
Custom Configurations	Custom Configurations allow extending target configuration collections	-	NA	
Deployment Procedure	Deployment procedures are customizable orchestration routines for various Provisioning and Patching tasks	-	-	
Discovery Security Class	Discovery Security Class	-	NA	
EM Plug-in	Manage the access control for Enterprise Manager plug-ins	-	NA	
Enterprise Manager High Availability	Enterprise Manager High Availability Administration allows to add additional management service through deployment procedure.	-	NA	
Enterprise Rule Set	Collection of rules that apply to Enterprise Manager elements, for example, targets and job. Individual rules can be used to send notifications, create incidents, update incidents, and other incident-management related actions.	-	-	

Granting User Privileges for Events, Incidents and Problems

In order to work with incidents, all relevant Enterprise Manager administrator accounts must be granted the appropriate privileges to manage incidents. Privileges for events, incidents, and problems are determined according to the following rules:

- 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 the events in the incident.
- Similarly, problem privileges are calculated based on privileges on underlying incidents.

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

For events, two privileges are defined in the system:

- 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 an event, creating an incident for an event, and creating a ticket for an event. You can also associate an event with an incident.

Important: Incident privilege is inherited from the underlying events.

If an event is raised on a target alone (the majority of event types are raised on targets such as metric alerts, availability events or service level agreement), you will need the following privileges:

- *View on target* to view the event.
- *Manage Target Events* to manage the event.

Note: This is a sub-privilege of Operator.

If an event is raised on both a target and a job, you will need the following privileges:

- *View on target* and *View on the job* to view the event.
- *View on target* and *Full on the job* to manage the event.

If the event is raised on a job alone, you will need the following privileges:

- *View on the job* to view the event.
- *Full on the job* to manage the event.

If an event is raised on a metric extension, you will need *View* privilege on the metric extension to view the event. Because events raised on metric extensions are informational (and do not appear in Incident Manager) event management privileges do not apply in this situation.

If an event is raised on a Self-update, only system privilege is required. Self-update events are strictly informational.

For incidents, two privileges are defined in the system:

- 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 events, and creating tickets for incidents.

If an incident consists of a single event, you can view the incident if you can view the event and manage the incident if you can manage the event.

If an incident consists of more than one event, you can view the incident if you can view at least one event and manage incident if you can manage at least one of the events.

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, and manually closing the problem.

In Enterprise Manager 12c, problems are always related to a single target. So the View Problem privilege, if an administrator has View privilege on the target, and the Manage Problem privilege, if an administrator has *manage_target_events* privilege on the target, implicitly grants management privileges on the associated event. This, in turn, grants management privileges on the incident within the problem.

3.3.3 Setting Up Rule Sets

Rule sets automate actions in response to incoming events, incidents and problems or updates to them. This section covers the most common tasks and examples.

- [Creating a Rule Set](#)
- [Creating a Rule to Create an Incident](#)
- [Creating a Rule to Manage Escalation of Incidents](#)
- [Creating a Rule to Escalate a Problem](#)
- [Subscribing to a Rule](#)
- [Receiving E-mail for Private Rules](#)

3.3.3.1 Creating a Rule Set

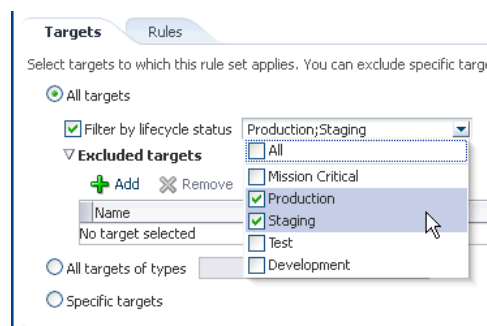
In general, to create a rule set, perform the following steps:

1. From the **Setup** menu, select **Incidents** then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, edit the existing rule set or create a new rule set. For new rule sets, you will need to first select the targets to which the rules apply. Rules are created in the context of a 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.

Narrowing Rule Set Scope Based on Target Lifecycle Status

When creating a new rule set, you can choose to have the rule set apply to a narrower set of targets based on the target's *Lifecycle Status* value. For example, you can create one rule set that only applies only to targets that have a *Lifecycle Status* of *Staging* and *Production*. As shown in the following graphic, you determine rule set scope by setting the *Lifecycle Status* filter.



Using this filter allows you to create rules for targets based on their *Lifecycle Status* without having to first create a group containing only such targets.

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 pop-up dialog. 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.

3.3.3.2 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 **Incident Rules**.
2. Determine whether there is an existing rule set that contains a rule that manages the event. In the **Incident Rules** page, use the Search option to find the rule/rule set name, description, target name, or target type for the target and the associated rule set. You can search by target name or the group target name to which this target belongs to locate the rule sets that manage the targets.

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.) **Note:** Only one event type can be selected in a single rule and, once selected, it cannot be changed when editing a rule.

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.

3.3.3.3 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 hours. If the fatal incident is still open after 96 hours, then it is escalated to level 3 and the operations VP is notified.

Provide the rule details using the Create New Rule wizard.

- a. To set up the rule to apply to all newly created incidents or when the incident is updated with *fatal* severity, select the **Specific Incidents** option and add the condition *Severity is Fatal*.
 - b. In the **Conditions for Actions** region located on the Add Actions page, select **Only execute the actions if specified conditions match**.
Select **Incident has been open for some time and is in a particular state (select time and optional expressions)**.
Select the time to be 48 hours and Status is not resolved or closed.
 - c. In the **Notification** region, type the name of the administrator to be notified by e-mail or page. Click **Continue** to save the current set of conditions and actions.

- d. Repeat steps b and c to page the DBA director (Time in this state is 72 hours, Status is Not Resolved or Closed). If open for more than 96 hours, set escalation level to 3, page Operations VP.
 - e. After reviewing added actions sets, click **Next**.
 - f. Click **Next** to go to the Summary screen. Review the summary information and click **Continue** to save the rule.
5. Review the sequence of existing enterprise rules and position the newly created rule in the sequence.

In Edit Rule Set page, click on the **desired rule from the Rules** table and select **Reorder Rules** from the **Actions** menu to reorder rules within the rule set, then click **Save** to save the rule sequence changes.

Example Scenario

To facilitate the incident escalation 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).

3.3.3.4 Creating a Rule to Escalate a Problem

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 e-mail.

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.

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. In the Rules section of the Edit Rule Set page, select **Create...**
4. From the **Select Type of Rule to Create** dialog, select **Newly created problems or updates to problems** and click **Continue**.
5. 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**.

6. 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**.

7. Review the rules summary. Make corrections as needed. Click **Continue** to return to Edit Rule Set page and then click **Save** to save the rule set.

3.3.3.5 Subscribing to a Rule

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 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**, and 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.

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. 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.

Alternate Rule Set Subscription Method: From the Incident Rules - All Enterprise Rules page, select the rule in incident rules table. From the **Actions** menu, select **E-mail** and then **Subscribe me** (or **Subscribe administrator....**).

3.3.3.6 Receiving E-mail for Private Rules

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. See [Chapter 9, "Utilizing the Job System and Corrective Actions"](#) for job privilege requirements.

Perform the following steps:

1. Navigate to the Rules page.
From the **Setup** menu, select **Incidents**, then select **Incident 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 **Incoming events and updates to events**. 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 is considered 'matching' against the incoming job events and e-mail will be sent 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.

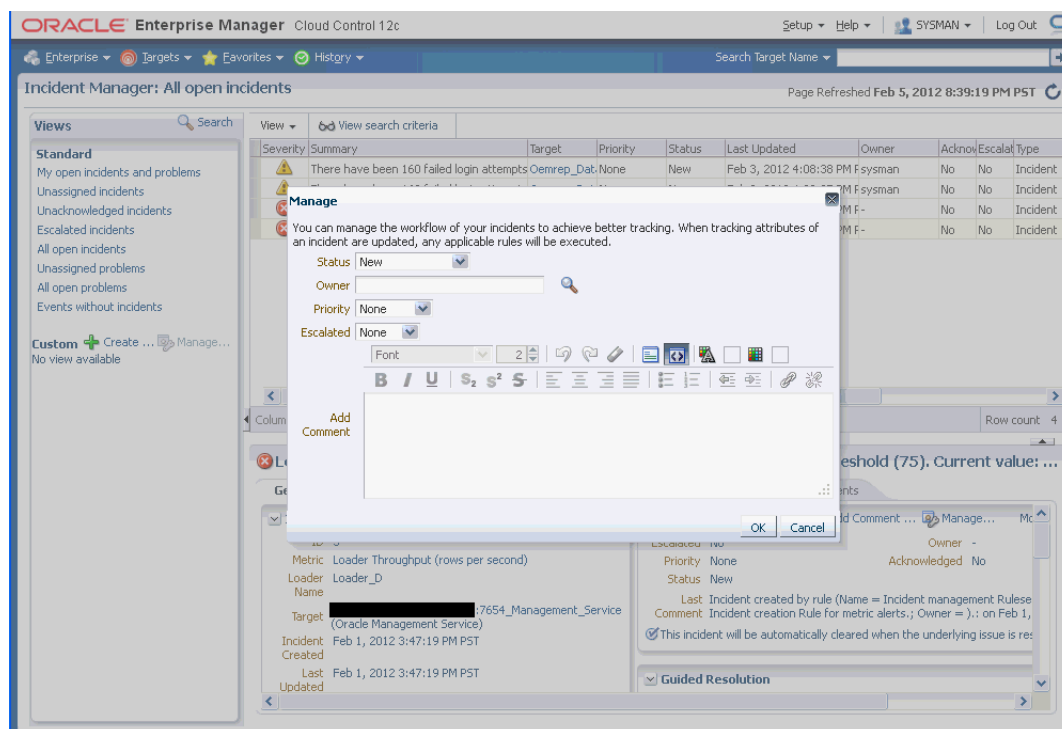
3.4 Working with Incidents

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:

- Send notifications to the appropriate administrators.
- Create incidents and rules.
- 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.
- Accessing My Oracle Support knowledge articles.
- Opening Oracle Service Requests to request assistance with issues with Oracle software (Problems).

You can update resolution information for 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 IT operator.

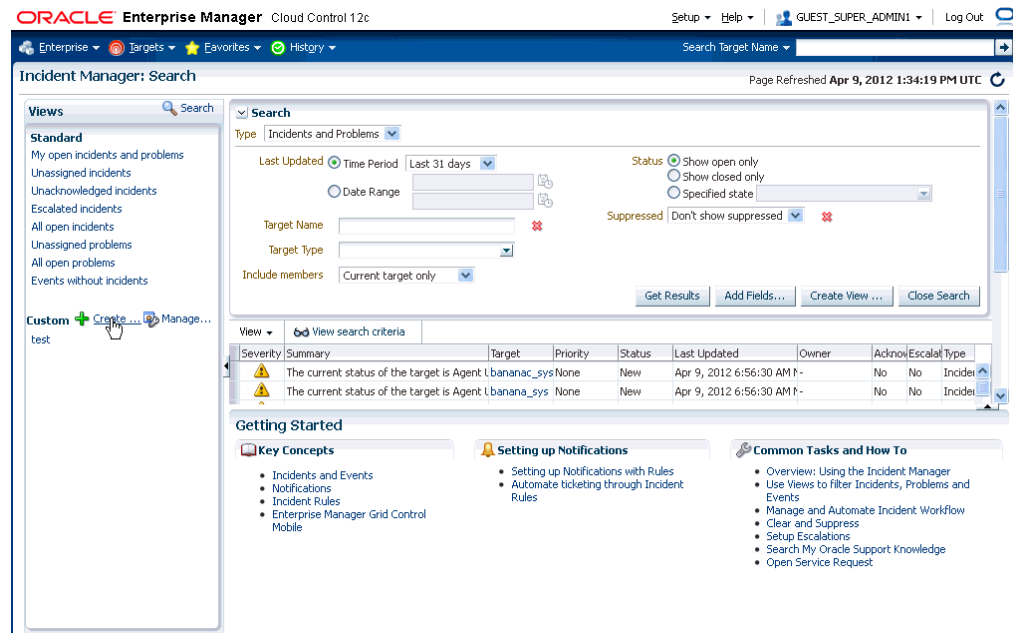
Working with incidents involves the following stages:

1. [Finding What Needs to be Worked On](#)
2. [Searching for Incidents](#)
3. [Setting Up Custom Views](#)
4. [Responding and Working on a Simple Incident](#)
5. [Responding to and Managing Multiple Incidents, Events and Problems in Bulk](#)
6. [Managing Workload Distribution of Incidents](#)
7. [Creating an Incident Manually](#)

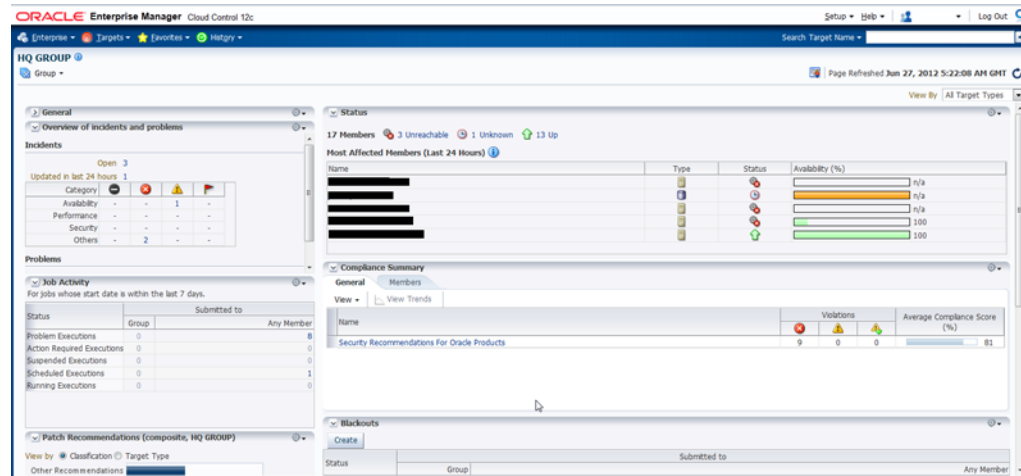
3.4.1 Finding What Needs to be Worked On

Enterprise Manager provides multiple access points that allow you to find out what needs to be worked on. The primary focal point for incident management is the Incident Manager console, however Enterprise Manager also provides other methods of notification. The most common way to be notified that you have an issue that needs to be addressed is by e-mail. However, incident information can also be found in the following areas:

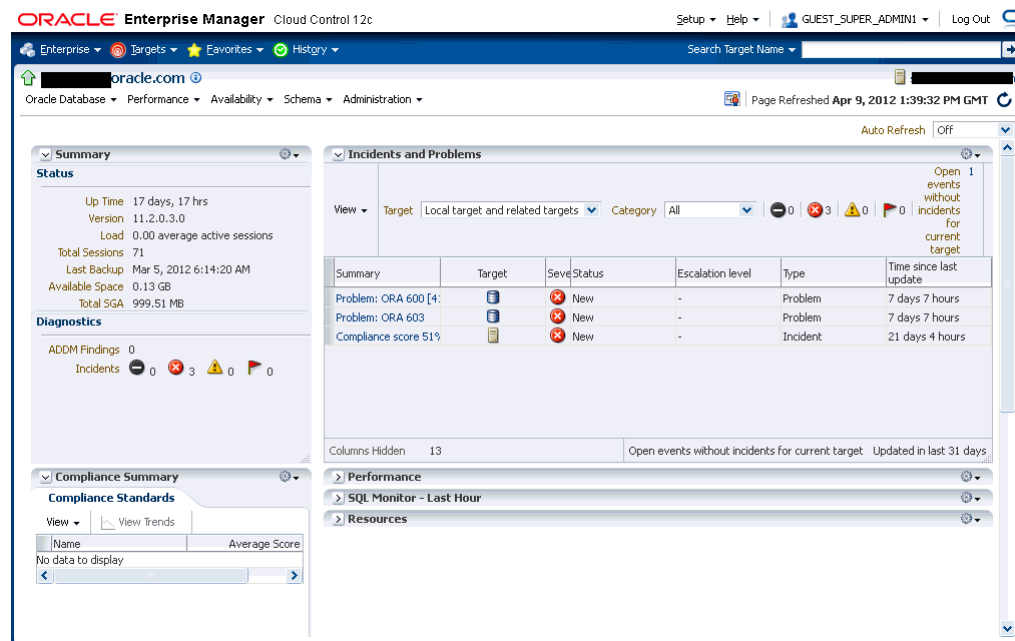
Custom Views (See ["Setting Up Custom Views"](#))



Group or System Homepages (See [Chapter 5, "Managing Groups"](#))



Target Homepages



Incident Manager (in context of a system or target)

Enterprise Manager Console

3.4.2 Searching for Incidents

1. Navigate to the Incident Manager page.

Using Incident Management 3-39

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**.

Validate that the list of incidents match what you are looking for. If not, change the search criteria as needed.
 - d. To view all the columns associated with this table, in the **View** menu, select **Columns**, then select **Show All**.

Searching for Incidents by Target Lifecycle Status

In addition to searching for incidents using high-level incident attributes, you can also perform more granular searches based on individual target lifecycle status. Briefly, lifecycle status is a target property that specifies a target's operational status. Status options for which you can search are:

- All
- Mission Critical
- Production
- Staging
- Test
- Development

For more discussion on lifecycle status, see [Section 3.5.5, "Event Prioritization."](#)

To search for incidents by target lifecycle status:

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**.
3. In the **Search** region, click **Add Fields**. A pop-up menu appears showing the available lifecycle statuses.
4. Choose on one or more of the lifecycle status options.
5. Enter any additional search criteria.
6. Click **Get Results**.

3.4.3 Setting Up Custom Views

Incident Manager also allows you to define custom 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 critical database incidents that you own. By specifying and saving view preferences to display only those incident attributes that you are interested in Enterprise Manager will show 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 incidents for daily triage. Accordingly, you can save the search criteria as a filter named "All priority 1 incidents

for my targets". The view becomes available in the UI for immediate use and will be available anytime you log in to access the specific incidents. The last view you used will be the default view used on your next login.

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**.
Validate that the list of incidents match what you are looking for. If not, change the search criteria as needed.
 - d. To view all the columns associated with this table, in the **View** menu, select **Columns**, then select **Show All**.
To select a subset of columns to display and also the order in which to display them, from the **View** menu, select **Columns**, then **Manage Columns**. A dialog displays showing a list of columns available to be added in the table.
 - e. Click the **Create View...** button.
 - f. Enter the view name.
 - g. Click **OK** to save the view.

3.4.4 Responding and Working on a Simple Incident

The following steps take you through one possible incident management scenario.

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, you should use **My Open Incidents and Problems** view to see incidents and problems assigned to you. You can then sort the list by priority.
3. To work on an incident, select the incident. In the **General** tab, click **Acknowledge** to indicate that you are working on this incident, and to stop receiving repeat notifications for the incident.

In addition to the acknowledging the incident, you can perform other incident management operations such as:

- Adding a comment.

- Managing the incident. See [Section 3.4.5, "Responding to and Managing Multiple Incidents, Events and Problems in Bulk"](#) for more information on incident management options.
- Editing the summary.
- Manually creating a ticket.
- Suppressing/unsuppressing the incident.
- Clearing the incident.

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.

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.

Alternatively, you can work with incidents for a specific target from that target's home page. From the *target* menu, select **Monitoring** and then select **Incident Manager** to access incidents for that target (or group).

3.4.5 Responding to and Managing Multiple Incidents, Events and Problems in Bulk

There may be situations where you want to respond to multiple incidents in the same way. For example, you find that a cluster of incidents that are assigned to you are due to insufficient tablespace issues on several production databases. Your manager suggests that these tablespaces be transferred to a storage system being procured by another administrator. In this situation, you want to set all of the tablespace incidents to a customized resolution state "Waiting for Hardware." You also want to assign the incidents to the other administrator and add a comment to explain the scenario. In this situation, you want to update all of these incidents in bulk rather than individually.

To respond to incidents in bulk:

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 to the subset of incidents you want to work on. For example, you can use **My Open Incidents and Problems** view to see incidents and problems assigned to you. You can then sort the list by priority.
3. Select the incidents to which you want to respond. You can select multiple incidents by holding down the Control key and selecting individual incidents or you can hold down the Shift key and select the first and last incidents to select a contiguous block of incidents.

4. From the **Action** menu, choose the desired response action.
 - **Acknowledge:** Indicate that you have viewed the incidents. This option also stops any repeat notifications sent out for the incidents. This sets the *Acknowledged* flag to *Yes* and also makes you the owner of the incident
 - **Manage:** Allows you to perform a multi-action response to the incidents.
 - *Acknowledge:* If 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. Acknowledgement also stops any repeat notifications for that incident
 - *Assign to:* Assign the incident(s) to the administrator who will take ownership of the incident.
 - *Prioritization:* The priority level of an incident can be set by selecting one of the out-of-the-box priority values: None, Urgent, Very High, High, Medium, Low
 - *Incident Status:* The resolution state for the incident can be set by selecting either Work in Progress or Resolved or to any custom status defined.
 - *Escalation Level:* Administrators can update incidents to set an escalation level: Level 1 through 5, in addition to the default value of None. An escalated issue can be de-escalated by setting the escalation to None. The appropriate *Escalation Level* depends on the IT procedures you have in place.
 - *Comment:* You can enter comments such as those you want to pass to the owner of the incident.
 - **Suppress:** Suppressing an incident stops corresponding notifications, and removes it from out-of-the-box views and default totals (such as those presented in the summary region). Suppression is typically performed when you want to defer action on the incident until a future time and in the meantime want to visually hide them from appearing in the console. Administrators can see suppressed incidents by explicitly searching for them such as performing a search on incidents where the search criteria includes the *Suppressed* search field
 Incidents can be suppressed until any of the following conditions are met:
 - Until the suppression is manually removed
 - Until specified date in the future
 - Until the severity state changes (incidents only)
 - Until it is closed
 - **Clear:** Administrators can clear incidents or problems manually. For incidents, this applies only to incidents containing incidents that can be manually cleared.
 - **Add Comment:** Users can add comments on incidents and events. Comments may be used for sharing information with other users or to provide tracking information on any actions being taken. Comments can be added even on closed issues.

Note: The single action **Acknowledge** and **Clear** buttons are enabled for open incidents and can be used for multiple incident selection.

If any of the above actions applies only to a subset of selected incidents (for example, if an administrator tries to acknowledge multiple incidents, of which some are already acknowledged), the action will be performed only where applicable. The administrator will be informed of the success or failure of the action.

When an administrator selects any of these actions, a corresponding annotation is added to the incident for future reference.

5. Click **OK**. Enterprise Manager displays a process summary and confirmation dialogs.
6. Continue working with the incidents as required.

3.4.6 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**.

If your My Oracle Support (MOS) login credentials have been saved as MOS Preferred Credentials, you do not need to log in manually. If not, you will need to sign in to My Oracle Support. To save your MOS login information as Preferred Credentials.

Setting MOS Preferred Credentials: From the **Setup** menu, select **Security** and then **Preferred Credentials**. From the My Oracle Support Preferred Credentials region, click **Set MOS Credentials**.

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.

3.4.7 Open Service Request (Problems-only)

There are times when you may need assistance from Oracle Support to resolve a *problem*. This procedure is not relevant for incidents or events.

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.

3.4.8 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 need to defer work on the incident until a future date (for example, until maintenance window). In order to avoid having it appear in the UI, you want to temporarily hide or suppress the incident until a future date. In order to find a suppressed incident, you must explicitly search for the incident using either the *Show all* or the *Only show suppressed* search option. In order to unhide a suppressed incident or problem, it must be manually unsuppressed.

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**.

3.4.9 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.

3.4.10 Reviewing Events on a Periodic Basis

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** standard view.

Perform the following steps:

1. From the **Enterprise** menu, select **Monitoring**, then select **Incident Manager**.
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 events for 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.

3.4.10.1 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".

3.5 Advanced Topics

The following sections discuss incident/event management features relating advanced applications or operational areas.

3.5.1 Defining Custom Incident Statuses

As discussed in ["Working with Incidents"](#) on page 3-36, one of the primary incident workflow attributes is *status*. For most conditions, these predefined status attributes will suffice. However, the uniqueness of your monitoring and management environment may require an incident workflow requiring specialized incident states. To address this need, you can define custom states using the *modify_resolution_state* EM CLI verb.

```
emcli modify_resolution_state
    -label="old label of the state to be changed"
    -new_label="New label for display"
    -position="New display position"
    [-applies_to=BOTH]
[ ] indicates that the parameter is optional
```

This verb modifies an existing resolution state that describes the state of incidents or problems. Only Super Administrators can execute this command. You need to specify the updated label as well as the updated position. The position can be between 2 and 98, and cannot be in use by another resolution state.

You can optionally indicate that the state should apply to both incidents and problems. A success message is reported if the command is successful. An error message is reported if the change fails.

3.5.2 Clearing Stateless Alerts for Metric Alert Event Types

For *metric alert* event types, an event (metric alert) is raised based on the metric threshold values. These metric alert events are called *stateful* alerts. For those metric alert events that are not tied to the state of a monitored system (for example, *snapshot too old*, or *resumable session suspended*), these alerts are called stateless alerts. Because stateless alerts are not cleared automatically, they need to be cleared manually. You can perform a bulk purge of stateless alerts using the *clear_stateless_alerts* EM CLI verb.

Note: For large numbers of incidents, you can manually clear incidents in bulk. See ["Responding to and Managing Multiple Incidents, Events and Problems in Bulk"](#).

clear_stateless_alerts clears the stateless alerts associated with the specified target. The clearing must be manually performed as the Management Agent does not automatically clear stateless alerts. To find the metric internal name associated with a stateless alert, use the EM CLI *get_metrics_for_stateless_alerts* verb.

Format

```
emcli clear_stateless_alerts -older_than=number_in_days -target_type=target_type
    -target_name=target_name [-include_members] [-metric_internal_name=target_type_
metric:metric_name:metric_column] [-unacknowledged_only] [-ignore_notifications]
```

`[-preview]`

`[]` indicates that the parameter is optional

Options

- **older_than**
Specify the age of the alert in days. (Specify 0 for currently open stateless alerts.)
- **target_type**
Internal target type identifier, such as `host`, `oracle_database`, and `emrep`.
- **target_name**
Name of the target.
- **include_members**
Applicable for composite targets to examine alerts belonging to members as well.
- **metric_internal_name**
Metric to be cleaned up. Use the `get_metrics_for_stateless_alerts` verb to see a complete list of supported metrics for a given target type.
- **unacknowledged_only**
Only clear alerts if they are not acknowledged.
- **ignore_notifications**
Use this option if you do not want to send notifications for the cleared alerts. This may reduce the notification sub-system load.
- **ignore_notifications**
Use this option if you do not want to send notifications for the cleared alerts. This may reduce the notification sub-system load.
- **preview**
Shows the number of alerts to be cleared on the target(s).

Example

The following example clears alerts generated from the database alert log over a week old. In this example, no notifications are sent when the alerts are cleared.

```
emcli clear_stateless_alerts -older_than=7 -target_type=oracle_database -tar get_
name=database -metric_internal_name=oracle_database:alertLog:genericErrStack
-ignore_notifications
```

3.5.3 User-reported Events

Users may create (publish) events manually using the EM CLI verb *publish_event*. A User-reported event is published as an event of the "User-reported event" class. Only users with Manage Target privilege can publish these events for a target. An error message is reported if the publish fails.

After an event is published with a severity other than CLEAR (see below), end-users with appropriate privileges can manually clear the event from the UI, or they can publish a new event using a severity level of CLEAR and the same details to report clearing of the underlying situation.

3.5.3.1 Format

```
emcli publish_event
    -target_name="Target name"
    -target_type="Target type internal name"
    -message="Message for the event"
    -severity="Severity level"
    -name="event name"
    [-key="sub component name"
    -context="name1=value1;name2=value2;.."
    -separator=context="alt. pair separator"
    -subseparator=context="alt. name-value separator"]
```

[] indicates that the parameter is optional

3.5.3.2 Options

- **target_name**

Target name.

- **target_type**

Target type name.

- **message**

Message to associate for the event. The message cannot exceed 4000 characters.

- **severity**

Numeric severity level to associate for the event. The supported values for severity level are as follows:

```
"CLEAR"
"MINOR_WARNING"
"WARNING"
"CRITICAL"
"FATAL"
```

- **name**

Name of the event to publish. The event name cannot exceed 128 characters.

This is indicative of the nature of the event. Examples include "Disk Used Percentage," "Process Down," "Number of Queues," and so on. The name must be repeated and identical when reporting different severities for the same sequence of events. This should not have any identifying information about a specific event; for example, "Process xyz is down." To identify any specific components within a target that the event is about, see the key option below.

- **key**

Name of the sub-component within a target this event is related to. Examples include a disk name on a host, name of a tablespace, and so forth. The key cannot exceed 256 characters.

- **context**

Additional context that can be published for a given event. This is a series of strings of format name:value separated by a semi-colon. For example, it might be useful to report the percentage size of a disk when reporting space issues on the disk. You can override the default separator ":" by using the sub-separator option, and the pair separator ";" by using the separator option.

The context names cannot exceed 256 characters, and the values cannot exceed 4000 characters.

- **separator**

Set to override the default ";" separator. You typically use this option when the name or the value contains ";". Using "=" is not supported for this option.

- **subseparator**

Set to override the default ":" separator between the name-value pairs. You typically use this option when the name or value contains ":". Using "=" is not supported for this option.

3.5.3.3 Examples

Example 1

The following example publishes a warning event for "my acme target" indicating that a HDD restore failed, and the failure related to a component called the "Finance DB machine" on this target.

```
emcli publish_event -target_name="my acme target" -target_type="oracle_acme"
-name="HDD restore failed" -key="Finance DB machine" -message="HDD restoration
failed due to corrupt disk" -severity=WARNING
```

Example 2

The following example publishes a minor warning event for "my acme target" indicating that a HDD restore failed, and the failure related to a component called the "Finance DB machine" on this target. It specifies additional context indicating the related disk size and name using the default separators. Note the escaping of the \ in the disk name using an additional "\\".

```
emcli publish_event -target_name="my acme target" -target_type="oracle_acme"
-name="HDD restore failed" -key="Finance DB machine" -message="HDD restoration
failed due to corrupt disk" -severity=MINOR_WARNING -context="disk
size":800GB\;"disk name":\\uddo0111245
```

Example 3

The following example publishes a critical event for "my acme target" indicating that a HDD restore failed, and the failure related to a component called the "Finance DB machine" on this target. It specifies additional context indicating the related disk size and name. It uses alternate separators, because the name of the disk includes the ":" default separator.

```
emcli publish_event -target_name="my acme target" -target_type="oracle_acme"
-name="HDD restore failed" -key="Finance DB machine" -message="HDD restoration
failed due to corrupt disk" -severity=CRITICAL -context="disk size"^800GB\;"disk
name"^\\sdd1245:2 -subseparator=context=^
```

3.5.4 Additional Rule Applications

Rules can be set up to perform more complicated tasks beyond straightforward notifications. The following tasks illustrate additional rule capabilities.

- [Setting Up a Rule to Send Different Notifications for Different Severity States of an Event](#)
- [Creating a Rule to Notify Different Administrators Based on the Event Type](#)

- [Creating a Rule to Create a Ticket for Incidents](#)
- [Creating a Rule to Send SNMP Traps to Third Party Systems](#)

3.5.4.1 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. Click **Create...**
4. From the Select Type of Rule to Create dialog, choose **Incoming events and updates to events**. Click **Continue**.
5. 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.
6. From the Rules section on the Edit Rules page, click **Create**.
7. Select **Newly created incidents or updates to incidents** as the rule type and click **Continue**.
8. Check **Specific Incidents**.
9. Check **Severity** and from the drop-down option selector, check **Critical** and **Warning**. Click **Next**.
10. On the Add Actions page, click **Add**. The Conditional Actions page displays.
11. In the **Conditions for actions** section, choose **Only execute the actions if specified conditions match**.

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

3.5.4.2 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. Search the list of enterprise rules matching the events from the production database.
3. 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.
4. From the Edit Rule Set page (Rules tab), select the rule which creates the incidents for the metric alert events for the database. Click **Edit**
5. From the Select Events page, click **Next**.
6. From the Add Actions page, click **+Add**. The Add Conditional Actions page displays.

7. In the Notifications area, enter the e-mail address of the DBA you want to be notified for this specific event type and click **Continue** to add the action. Enterprise Manager returns you to the Add Actions page.
8. Click **Next**.
9. On the Specify Name and Description page, enter an intuitive rule name and a brief description.
10. Click **Next**.
11. On the Review page, review the **Applies to**, **Actions** and **General** information for correctness .
12. Click **Continue** to create the rule.
13. Create/Edit additional rules to handle alternate additional administrator notifications according to event type.
14. Review the rules summary and make corrections as needed. Click **Save** to save your rule set changes.

3.5.4.3 Creating a Rule to Create a Ticket for Incidents

If your IT process requires a helpdesk ticket be created to resolve incidents, then you can use the helpdesk connector to associate the incident with a helpdesk ticket and have Enterprise Manager automatically open a ticket when the incident is created. Communication between Incident Manager and your helpdesk system is bidirectional, thus allowing you to check the changing 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.

For example, 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 set up to create a Remedy 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...**
4. Click **Next** to proceed to the **Add Actions** page.
5. Click **+Add** to access the **Add Conditional Actions** page.
 - a. Specify that a ticket should be generated for incidents covered by the rule.

- b. Specify the ticket template to be used.
6. Click **Continue** to return to the Add actions page.
7. On the Add Actions page, click **Next**.
8. On the Review page, click **Continue**.
9. On the Specify Name and Description page, click **Next**.
10. On the Review page, click **Continue**. A message displays indicating that the rule has been successfully modified. Click **OK** to close the message.
11. Repeat steps 3 through 10 until all appropriate rules have been edited.
12. Click **Save** to save your changes to the rule set.

3.5.4.4 Creating a Rule to Send SNMP Traps to Third Party Systems

As mentioned in [Chapter 4, "Using Notifications,"](#) Enterprise Manager supports integration with third-party management tools through the SNMP. Sending SNMP traps to third party systems is a two-step process:

Step 1: Create an advanced notification method based on an SNMP trap.

Step 2: Create an incident rule that invokes the SNMP trap notification method.

The following procedure assumes you have already created the SNMP trap notification method. For instruction on creating a notification method based on an SNMP trap, see ["Sending SNMP Traps to Third Party Systems"](#) on page 4-38.

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. On the Incident Rules - All Enterprise Rules page, click **Create Rule Set...**
3. Enter the rule set **Name**, a brief **Description**, and select the type of source object the rule **Applies to** (Targets).
4. Click on the **Rules** tab and then click **Create...**
5. On the Select Type of Rule to Create dialog, select **Incoming events and updates to events** and then click **Continue**.
6. On the Create New Rule : Select Events page, specify the criteria for the events for which you want to send SNMP traps and then click **Next**.

Note: You must create one rule per event type. For example, if you want to send SNMP traps for Target Availability events and Metric Alert events, you must specify two rules.

7. On the Create New Rule : Add Actions page, click **Add**. The Add Conditional Actions page displays.
8. In the Notifications section, under Advanced Notifications, select an existing SNMP trap notification method as shown in the following graphic.

ORACLE Enterprise Manager Cloud Control 12c Help

Add Actions

Add Conditional Actions Continue Cancel

Define actions to be taken when an event matches this rule.

Conditions for actions
You can define the actions to apply whenever the rule matches or apply them conditionally.
☒ Always execute the actions
☐ Only execute the actions if specified conditions match

Create Incident or Update Incident
If there is no incident associated with the event, you could create one and optionally, set the incident owner and priority. If an incident exists, you could update the incident.
☐ Create Incident (If not associated with one) ☐ Update Incident

Notifications
Assign recipients for notifications. Recipients for the "To" list can only be added or removed in this section. Users who subscribe to this rule will be added to the "Cc" list; users who unsubscribe to this rule will be removed from the "Cc" list. You could specify multiple users separated by commas. Recipients could be Enterprise Manager users, direct E-mail address or [predefined variables](#).

Basic Notifications
 E-mail To:
 E-mail Cc:
 Page:

Advanced Notifications

Name	Description	Supports Repeat
<input checked="" type="checkbox"/> HP OpenView Console (SNMP Trap)	Notification method to send trap to HP OpenView console	

Repeat Notifications(Not specified)

Clear events
For most events, Enterprise Manager detects when the underlying issue is cleared and will generate a clear event. These types of events cannot be cleared using this option. However, for some events, such as metric alerts that are generated by mining a log file, it is not feasible for Enterprise Manager to detect when the underlying issue is cleared. This type of events must be manually cleared by administrators. This action can be used to automate this behavior.
☐ Clear permanently

Forward to Event Connectors

For information on creating SNMP trap notification methods, see ["Sending SNMP Traps to Third Party Systems"](#) on page 4-38.

9. Click **Continue** to return to the Create New Rule : Add Actions page.
10. Click **Next** to go to the Create New Rule : Specify Name and Description page.
11. Specify a rule name and a concise description and then click **Next**.
12. Review the rule definition and then click **Continue** add the rule to the rule set. A message displays indicating the rule has been added to the rule set but has not yet been saved. Click **OK** to close the message.
13. Click **Save** to save the rule set. A confirmation is displayed. Click **OK** to close the message.

3.5.5 Event Prioritization

When working in a large enterprise, it is conceivable that when systems are under heavy load, a large number of incidents and events may be generated. All of these need to be processed in a timely and efficient manner in accordance with your business priorities. An effective prioritization scheme is needed to determine which events/incidents should be resolved first.

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_property_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 events (highest priority)
- Non-informational events.
- Informational events

3.5.6 Root Cause Analysis (RCA) and Target Down Events

Root Cause Analysis (RCA) tries to identify the root causes of issues that cause operational events. Beginning with Enterprise Manager Cloud Control 12.1.0.3, Incident Manager automatically performs RCA over *target down events*, thus actively identifying whether the target down event is the cause or symptom of other target down events.

The term target down event specifically pertains to Target Availability events that are raised when the targets are detected to be down.

3.5.6.1 How RCA Works

RCA is an ongoing process that identifies whether a target down event is root cause or symptom. It uses the Causal Analysis Update attribute of the event to store the results of its analysis, i.e. identifying whether or not the target down event is root cause or symptom. Whenever a new target availability event comes in, RCA is automatically performed on the incoming event and existing target down events that are related to it. Based on the analysis, it updates the Causal Analysis Update attribute value if the incoming event is a target down event. It also updates the Causal Analysis Update attribute for the related target down events if there is a change.

Two types of target relationships are used for identifying the related targets: *dependency* and *containment*.

When one target depends on another target for its availability, dependency relationship exists between them. For example, J2EE application target depends on the Weblogic Server target over which it is deployed.

When a target has a containment relationship over its member targets from operational point of view, i.e., the member target does not exist outside of the container, a containment relationship exists between the container target and each member. For example, Real Application Cluster (RAC) database target has containment relationship with each of its RAC database instance targets.

The causal analysis update attribute is used only for target down events (i.e. Target Availability event for target down) and can have be assigned any one of the following values by the RCA process:

- *Symptom* -- this means the target down event has been caused by another target down event. In the case of containment relationships, target down events of targets that are part of a container target are set as symptoms to de-emphasize the events in favor of the target down event at the target container level.
- *Root Cause* - this means the target down event has caused another target down event and it is not the symptom of any other target down event.
- *Composite Down* - this applies to a container target that is down and it is not a symptom of any other container target down event. It is marked as composite down to emphasize its importance over the target down events of targets that are part of the container target.
- *N/A* - root cause analysis is not applicable to this event .
- *Not a cause and not a symptom* - the target down event is not root cause and not a symptom of other target down events This is shown in Incident Manager as a "-".

The following rules describe the RCA process:

- **Rule 1:** Down event on a non-container target (a target that does not have members) is marked as the cause if a dependent target is down and it is not symptom of other target down events.

Examples:

- You have a J2EE applications deployed on a standalone Weblogic Server. If both J2EE application and WebLogic Server targets are down, the Weblogic Server down event is the cause for the J2EE applications deployed on it.
- You have a J2EE application deployed on couple of Weblogic Servers, which are part of a Weblogic Cluster. If one WebLogic Server is down along with its J2EE application, then the Weblogic Server down event is the cause of the J2EE application target down. This assumes the Weblogic Cluster is not down.
- **Rule 2:** Down event on a non-container target (a target that does not have members) is marked as a symptom if a target it depends on is down or if the target containing it is down.

Examples:

- You have a J2EE application deployed on a standalone Weblogic Server. If both J2EE application and WebLogic Server targets are down, J2EE application down event is the symptom of Weblogic Server being down.
- You have a couple of Weblogic Servers which are part of a WebLogic Cluster. Each WebLogic Server has a J2EE application deployed on it. If the Weblogic Cluster is down, this means both WebLogic Servers are down. This in turn means the J2EE applications that are deployed on them are also down. The Weblogic Server down events would be marked as the symptoms of Weblogic Cluster being down. Logically, Weblogic Servers down events are both a cause (for the J2EE application down events) and a symptom (of the Weblogic Cluster down event). However, the WebLogic Server down is marked as symptom instead of root cause because root cause is meant to represent the main issue that administrators should focus on. In this scenario, the main focus is on resolving the WebLogic Cluster down event (i.e. restoring the WebLogic Cluster). The WebLogic Cluster down event is marked as composite down. See Rule 3 for details.
- You have a couple of RAC database instance targets that are part of cluster database target. If the cluster database is down, then means all RAC instances are also down. The RAC instance down events would be marked as the

symptoms of cluster database being down. The RAC instance down is marked as a symptom because the focus should be on resolving the cluster database down event. The cluster database down event is marked as composite down. See Rule 3 for details.

- **Rule 3:** Down event on a container target is marked as composite down if all member targets are down and any target containing it is not down.

Examples:

- You have a couple of Weblogic Servers, which are part of a standalone Weblogic Cluster. A Weblogic Cluster down event would be marked as composite down, if both the Weblogic Servers are down. The focus should be on the WebLogic Cluster down event (restoring the WebLogic Cluster) so that is why it is marked as composite down.
 - You have a couple of RAC database instance targets that are part of a cluster database target. The cluster database target down event would be marked as composite down, if both database instances are down. The focus should be on the cluster database down event (i.e. restoring the cluster database) so that is why it is marked as composite down.
- **Rule 4:** Down event on a container target is marked as symptom if the target containing it is down.

Example:

You have a couple of Weblogic Clusters that are part of a Weblogic Domain target. If the Weblogic domain is down, this means the WebLogic Clusters are also down. The Weblogic Cluster target down events would be the symptoms of Weblogic Domain being down. The WebLogic Domain down event would be marked as composite down because that should be the focus,

3.5.6.2 Leveraging RCA Results in Incident Rule Sets

As described above, RCA is an ongoing process which results in marking target down events as *cause*, *symptom*, *composite down* or *neither* as new target down events come in and are processed. So a target down event may be marked as a cause or symptom or composite down as it comes in or after some time when RCA has analyzed additional event information.

Most datacenters automatically create incidents for target down events since these are important events that need to be resolved right away. This is recommended best practice and also implemented by the out-of-the-box rule sets. However, in terms of notifying response teams or creating trouble tickets, it is not desirable to do so for symptom incidents. Some datacenters may also choose to not create incidents for symptom events.

So the RCA results can be leveraged to do the following:

1. Notify or create tickets only for non-symptom events:

This can be achieved in 2 ways:

- Create two separate event rules, one event rule to create incidents for all relevant events, but take no further action (no notification or ticket creation) and another one to create incidents for non-symptom events only and also send notifications and create tickets. See "[Creating Incidents On Non-symptom Events](#)" on page 3-63 for instructions.
- Create an event rule that creates incidents for all target down events. Create another rule to update the incident priority, send notifications and create

tickets only for incidents stemming from non-symptom events. Once the incident priority is set to say "Urgent", customer can also create additional incident rules to take additional actions on the Urgent priority incidents. See ["Creating a Rule to Update Incident Priority for Non-symptom Events"](#) on page 3-62.

2. Only create incidents after a suitable wait for events that are not initially marked as neither a cause nor a symptom:

As mentioned previously, RCA is an iterative process whereby incoming target down events are continually being evaluated, resulting in updates to causal analysis state of existing events. Over a period of time (minutes), a target down event that was initially marked as a root cause may or may not remain a root cause depending on other incoming target down events. The original target down event may later be classified as a symptom.

To avoid prematurely creating an incident and opening a ticket for an event which may later turn out to be a symptom event, you can set up your rules as follows:

- In addition to the rules already defined in the previous step, create an additional event rule to act upon RCA updates to events and when the RCA update indicates that the event is marked as a symptom, lower the priority of the incident to "Low". This will also send an update to the ticket automatically. This is recommended. See ["Introducing a Time Delay"](#) on page 3-64 for instructions.

OR

- To allow time for target down events to be reported, analyzed, and then acted upon (such as creating an incident or updating an incident), you can add a delay in the rule actions. This is useful when customer have some tolerance to take action after some minimum delay (typically 5 minutes).

3. Only create incidents for non-symptom events.

Some datacenters may choose not to create any incidents for symptom events. This can be achieved by changing the rules to only create incidents for events marked as cause or neither a cause nor symptom. See ["Creating Incidents On Non-symptom Events"](#) on page 3-63 for instructions.

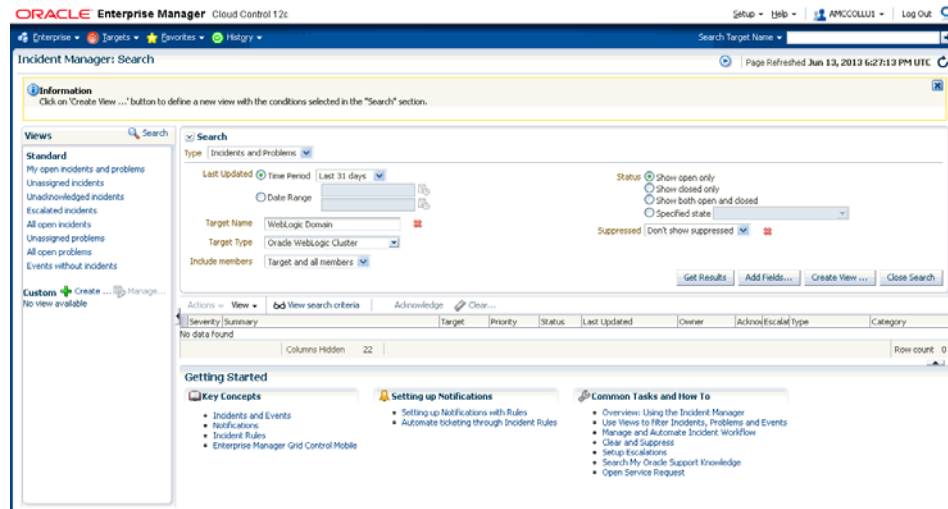
Please note that, even in this approach, it is possible that an event that was originally marked as cause or neither a cause nor symptom, may be marked as a symptom when more information is received. Customers can use an approach similar to that of the second option in step 2 to build some delay in creating the incidents.

Even with this, it is still feasible but a bit unlikely, that newer information shows up after the pre-set delay and ends up marking the event as symptom. So it is recommended to use the approach of setting incident priority and using that as a way to manage workflow.


3.5.6.3 Leveraging RCA Results in Incident Manager

You can use the RCA results to focus on the non-symptom incidents in Incident Manager. This involves using the Causal Analysis Update incident attribute when creating custom views.


1. From the **Enterprise** menu, select **Monitoring**, and then **Incidents**. The Incident Manager page displays.
2. From the Views region, click **Create**. The Search page displays.



3. Click **Add Fields...** and then choose **Causal analysis update**. The Causal analysis update displays as additional search criteria.



4. Choose 'Do Not Show Symptoms' from the list of available criteria. This will automatically exclude incidents that have been marked as 'symptom'. Incidents that are not marked as symptom or root cause will be included as long as it matches any other criteria you may have specified.

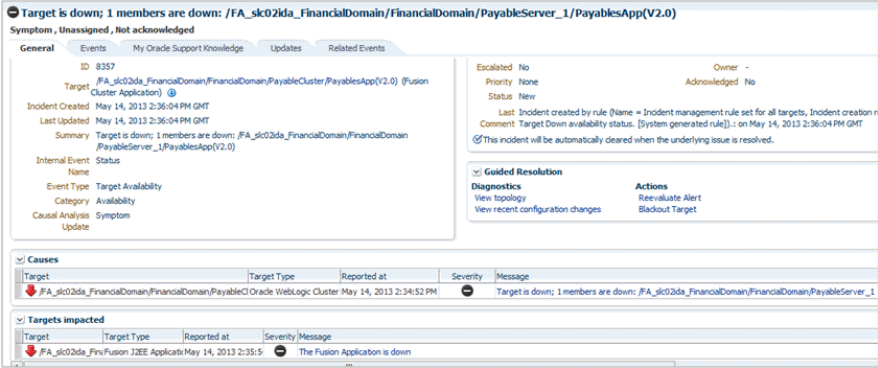


5. Click **Create View**, enter a **View Name** when prompted, and then click **OK**.

Showing RCA Results in an Incident Detail

An incident that is a root cause or symptom will be identified prominently as part of the details of the incident in Incident Manager. In addition, in case the incident is a symptom, a **Causes** section will be added to identify the root cause(s) of the incident. In case the incident has, in turn, caused other target down incidents, an **Impacted Targets** section will also be added to show the targets that have been affected, that is, other targets that are down as a result of the original target down. The following figure shows the incident detail.

Figure 3–8 Incident that is a Root Cause

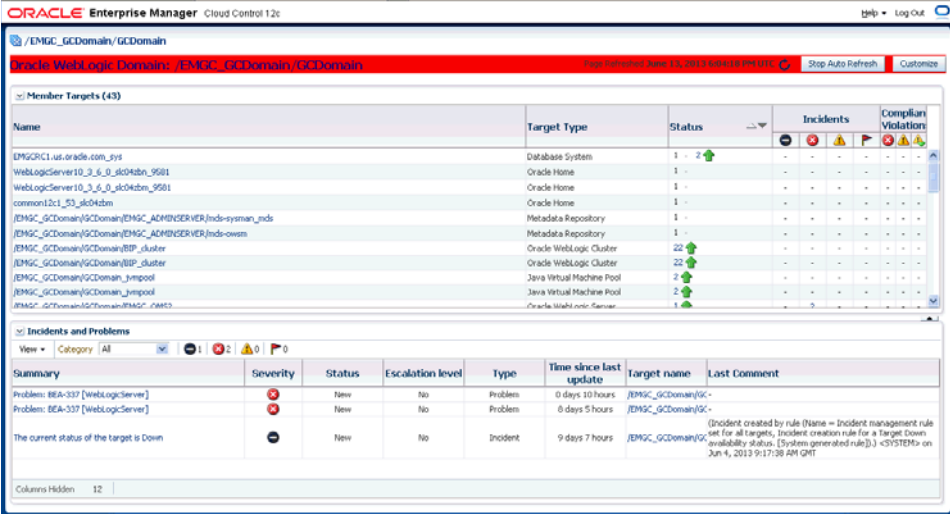


3.5.6.4 Leveraging RCA Results in the System Dashboard

In the System Dashboard, you can use the RCA results to exclude symptom incidents from the Incidents table so administrators can focus their attention on incidents that are root cause or have not been caused by other target down events.

To exclude Symptom Incidents:


1. In the System Dashboard, click on the **View** option that is accessible from the upper left hand corner of the **Incidents and Problems** table.



2. Choose the option to 'Exclude symptoms'. Alternatively, you can also choose the option 'Cause only' show only shows target down incidents that have been identified as cause of other target down incidents. Regardless of the option chosen, incidents that have not been marked as symptom or root cause will continue to be displayed.



2. When adding an action, select the priority to be set for incidents associated with the non-symptom events selected above.

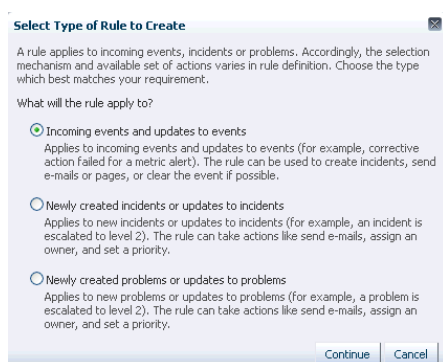


3.5.6.6 Creating Incidents On Non-symptom Events

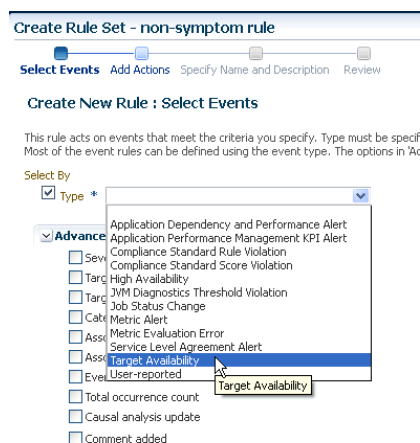
You can leverage Incident Manager's RCA capability creating rule sets that generate incidents. For monitoring situations where a high number of *symptom* target down events are generated, but only a few non-symptom target down events, you can create rule sets that generate incidents and send notifications only for non-symptom events.

To create a rule set that creates incidents for non-symptom target down events:

1. From the **Setup** menu, select **Incidents**, then select **Incident Rules**.
2. Click **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 "Incoming events and updates to events."



3. Click **Continue**. The *Create New Rule : Select Events* displays.



4. In the Advanced Selection Options region, choose *Causal analysis update*. Three causal event options display.

event is marked as cause or composite down: A target down is considered a cause if other targets depending on it are down or if its a composite.

event is marked as a symptom: A target down is considered a symptom if a target it depends on is also down, or if the target is part of a composite which is down.

event is not a cause and not a symptom: A target down is neither a cause or symptom.

By selecting one or more options, you can filter out extraneous target down events and focus on those target availability events that pertain to targets with interdependencies. To create an incident for only non-symptom events, choose *event is marked as cause or composite down* and *event is not a cause and not a symptom*.

Advanced Selection Options

- ☐ Severity
- ☐ Target type
- ☐ Target Lifecycle Status
- ☐ Category
- ☐ Associated with incident
- ☐ Associated incident acknowledged
- ☐ Event name
- ☐ Total occurrence count
- ☒ Causal analysis update
 - ☒ event is marked as cause or composite down
 - ☐ event is marked as a symptom
 - ☒ event is not a cause and not a symptom
- ☐ Comment added

Click **Next**.

5. On the *Create New Rule : Add Actions* page, click **Add**. The *Add Conditional Actions* page displays.
6. In the *Create Incident or Update Incident* region, choose **Create Incident**.
7. Specify the remaining assignment and notification details and click **Continue**.
8. Complete the remaining *Create Rule Set* wizard pages. See ["Creating a Rule Set"](#) on page 3-30 for more information on creating rule sets.

3.5.6.7 Introducing a Time Delay

As mentioned previously, Incident Manager RCA is an iterative process whereby incoming target down events are continually being evaluated, resulting in updates to causal analysis states. Over a period of time (minutes), a root cause may or may not remain a root cause depending on incoming target down events. The original target down event may later be classified as a symptom. To allow time for target down events to be reported, analyzed, and then acted upon (such as creating an incident), you can define an event evaluation time delay when creating a rule set.

In the previous example, where incidents are created for non-symptom events, without a time delay in the rule, there could potentially be an incident created for a non-symptom event that eventually becomes a symptom.

To add a time delay to the rule:

1. From the *Create Rule Set* wizard *Add Actions* page, click **Add** or **Edit** (modify an existing rule). The *Add Conditional Actions* page displays.
2. In the *Conditions for Actions* region, choose **Only execute the actions if specified conditions match**. A list of conditions displays.
3. Choose **Event has been open for specified duration**.
4. Specify the desired time delay.

5. Click **Continue** and complete the remaining steps in the wizard.

3.6 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. For example, critical metric alerts will be available as critical incidents.

Rules

When you migrate to Enterprise Manager 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.

Using 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](#)
- [Sending Notifications Using OS Commands and Scripts](#)
- [Sending Notifications Using PL/SQL Procedures](#)
- [Sending SNMP Traps to Third Party Systems](#)
- [Management Information Base \(MIB\)](#)
- [Passing Corrective Action Status Change Information](#)
- [Passing Job Execution Status Information](#)
- [Passing User-Defined Target Properties to Notification Methods](#)
- [Troubleshooting Notifications](#)
- [EMOMS Properties](#)
- [Passing Event, Incident, Problem Information to an OS Command or Script](#)

- [Passing Information to a PL/SQL Procedure](#)

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, you can also define notifications for 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

Setup

Enterprise Manager requires the following information to send e-mail notifications by means of Incident Rules. When specifying multiple SMTP servers, separate each server by a comma or space. Revert Apply Test Mail Servers

Outgoing Mail (SMTP) Server

Use the format SERVER:PORT (Example: SMTP1:587). Port 25 is used if no port is specified for the server. (Example: SMTP1, MyServer:587).

User Name

Specify user name if your SMTP server requires authentication.

Password

Specify the authentication password. The name and password will be used for all SMTP servers.

Confirm Password

Identify Sender As

Sender's E-mail Address

Use Secure Connection ☒ No ☐ TLS, if available ☐ SSL

Scripts and SNMP Traps

Before Enterprise Manager can send notifications by means of OS commands, PL/SQL procedures, or SNMP traps, they must first be defined as Notification Methods. Administrators can then use these methods in Incident Rules.

Add OS Command Go

Name	Type	Support Repeat Notifications
No notification methods found.		

TIP Remember to create Incident Rules in order to send notifications by means of these methods.

Repeat Notifications

Repeat notifications allow you to be notified repeatedly about the same events, incidents or problems. Once enabled, you will still need to choose the repeat notification option in each Incident Rule that will use it. If you disable repeat notifications on this page, all repeat notifications will stop.

☐ Send Repeat Notifications

Repeat Frequency (minutes)

Maximum Repeat Notifications

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**.

As standard practice, each user should have their own e-mail address.

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.

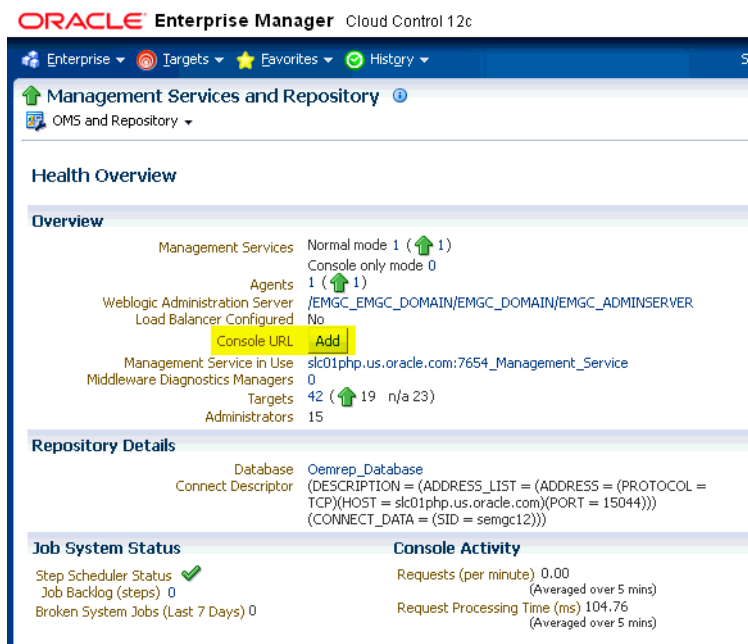
Defining multiple mail servers will improve the reliability of e-mail notification delivery. E-mail notifications will be delivered if at least one e-mail server is up. The notification load is balanced across multiple e-mail servers by the OMS, which switches through them (servers are allocated according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

Setting the Cloud Control Console URL when Using an SLB

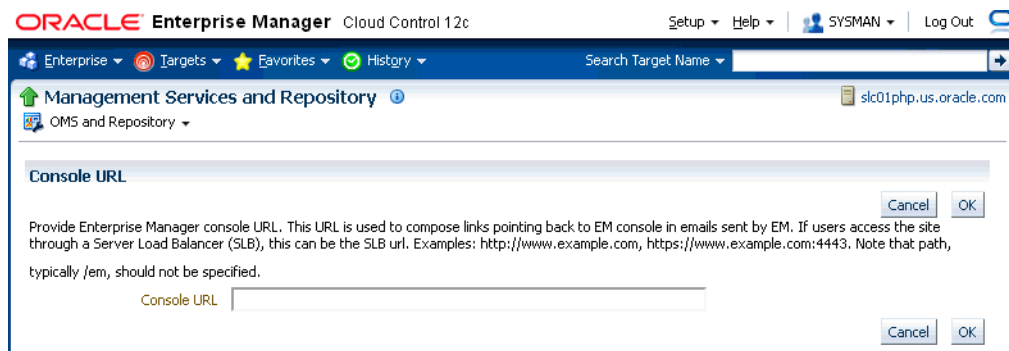
If you have a multi-OMS environment with a Server Load Balancer (SLB) configured for the OMS instances, you should update the console URL to ensure that any emails from Enterprise Manager direct you to the Enterprise Manager console through the SLB URL and not the specific OMS URL from which the email may have originated.

To change the console URL:

1. From the **Setup** menu, select **Manage Cloud Control**, and then **Health Overview**. The Management Services and Repository page displays.
2. On the Management Services and Repository page, in the Overview section, click **Add/Edit** against the *Console URL* label.



The *Console URL* page displays.



3. Modify the Console URL to the SLB URL.

Examples:

<http://www.example.com>

<https://www.example.com:4443>.

Note that path, typically */em*, should not be specified.

4. Click **OK**.

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 an 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 *username* drop-down menu, 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.

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. Only e-mail addresses that have been specified with your user preferences (**Enterprise Manager Password and Email** page) can be used in the 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.

Figure 4–2 Notification Schedule Page

ORACLE Enterprise Manager Cloud Control 12c

Setup Help SYSMAN Log Out

Enterprise Targets Favorites History Search Target Name

Setup

Notification Schedule

You need to set up your notification schedule in order to receive e-mail notifications for alerts. For each day in the schedule, specify the e-mail addresses that should be used to send you e-mail notifications. [Edit Schedule Definition](#)

Administrator: **SYSMAN** [Change](#)

Suspend Notification: **None** [Edit](#)

Rotation Frequency: **1 Week**
 Current Week: **1**
 Time Zone: **(UTC-07:00)**

View Week Containing: **Oct 9, 2012** [Go](#)
 Example: Oct 31, 2002

☐ Indicates time slots with E-mail Addresses ☐ Indicates Suspended Notification

	Oct 8, 2012 Monday	Oct 9, 2012 Tuesday	Oct 10, 2012 Wednesday	Oct 11, 2012 Thursday	Oct 12, 2012 Friday	Oct 13, 2012 Saturday	Oct 14, 2012 Sunday
12:00 am - 1:00 am	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...
1:00 am - 2:00 am	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...
2:00 am - 3:00 am	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...
3:00 am - 4:00 am	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...
4:00 am - 5:00 am	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...
5:00 am - 6:00 am	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...
6:00 am - 7:00 am	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...
7:00 am - 8:00 am	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...
8:00 am - 9:00 am	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...
9:00 am - 10:00 am	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...
10:00 am - 11:00 am	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...	smptestemail@o...

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 - 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.

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. 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**.

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 "[Setting Up Rule Sets](#)" 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 whose 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 desired 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 4–1 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 `[METRIC_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 4–2 Control Structures

Control Structure	Description
Pipe " "	<p>Two or more attributes can be separated by ' ' character. For example,</p> <pre>[METRIC_NAME SEVERITY]</pre> <p>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 attribute is applicable, only the left-most attribute is used.</p>

Table 4–2 (Cont.) Control Structures

Control Structure	Description
IF	<p>Allows you to make a block of text conditional. Only one level of IF and ELSIF is supported. Nested IF constructs are not supported.</p> <p>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.</p> <p>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,</p> <pre>[IF REPEAT_COUNT EQ "1" AND SEVERITY EQ "CRITICAL" THEN]</pre> <p>The statement above will be true when the attributes of the alert match the following condition:</p> <ul style="list-style-type: none"> ■ Attribute Name: REPEAT_COUNT ■ Attribute Value: 1 ■ Attribute Name: SEVERITY ■ Attribute Value: CRITICAL <p>Example IF Block:</p> <pre>[IF EXECUTION_STATUS NEQ NULL] [JOB_NAME_LABEL] = [EXECUTION_STATUS] [JOB_OWNER_LABEL] = [JOB_OWNER] [ENDIF]</pre> <pre>[IF SEVERITY_CODE EQ CRITICAL] [METRIC_NAME_LABEL] = [METRIC_GROUP] [METRIC_VALUE_LABEL] = [METRIC_VALUE] [TARGET_NAME_LABEL] = [TARGET_NAME] [KEY_VALUES] [ENDIF]</pre> <p>Example IF and ELSEIF Block:</p> <pre>[IF SEVERITY_CODE EQ CRITICAL] statement1 [ELSIF SEVERITY_CODE EQ WARNING] statement2 [ELSIF SEVERITY_CODE EQ CLEAR] statement3 [ELSE] statement4 [ENDIF]</pre>

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: |

4.1.5 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).

Configuring Repeat Notifications Globally

To enable repeat notifications for a notification method (globally), select the **Send Repeat Notifications** option on the Notification Methods page . 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.

Configuring Repeat Notifications Via Incident Rules

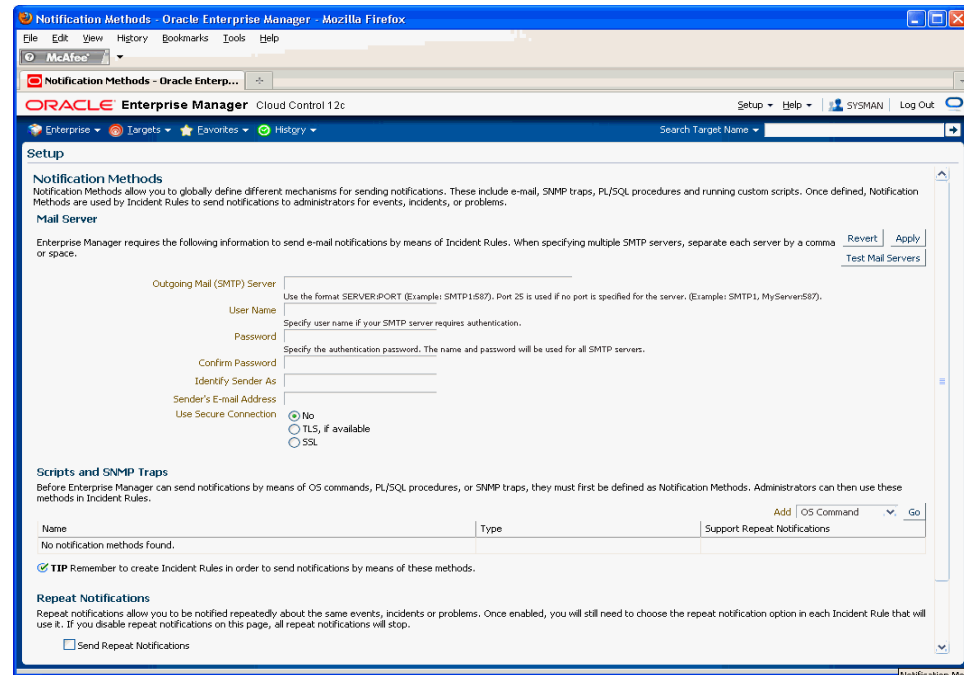
Setting repeat notifications globally at the notification method level may not provide sufficient flexibility. For example, you may want to have different repeat notification settings based on event type. Enterprise Manager accomplishes this by allowing you to set repeat notifications for individual 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 will only be sent if the **Send Repeat Notifications** option is enabled in the Notification Methods page.

Non-E-mail Repeat Notifications

For non-e-mail repeat notifications (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 4–4 Enabling Repeat Notification for an OS Command Notification Method



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.

4.3 Sending Notifications Using OS Commands and Scripts

Notification system can invoke a custom script when an incident rule matches the OS Command advanced notification action. A 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 `oracle.sysman.core.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.

Registering a Custom Script

In order to use a custom script, you must first register the script with the notification system. This is performed in four steps:

1. Define your OS command or script.
2. Deploy the script on each Management Service host.
3. Register your OS Command or Script as a new Notification Method.
4. Assign the notification method to an incident rule.

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 4-52
- ["Passing Corrective Action Execution Status to an OS Command or Script"](#) on page 4-44
- ["Passing Job Execution Status to an OS Command or Script"](#) on page 4-48

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.

Important: Both scripts and OS Commands should be specified using absolute paths. 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, eventually progressing to 30 minutes. From here, the procedure is retried every 30 minutes until the notification is a 24 hours old. The notification will be then 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 prevents OS Commands from running for an inordinate length of time and blocks the delivery of other notifications. By default the command is allowed to run for 30 seconds before it is killed. The `oracle.sysman.core.notification.os_cmd_timeout` emoms property can be configured to change the default timeout value.

Example 4-4 Changing the `oracle.sysman.core.notification.os_cmd_timeout` emoms Property

```
emctl set property -name oracle.sysman.core.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 "[Sending Notifications Using OS Commands and Scripts](#)".

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.

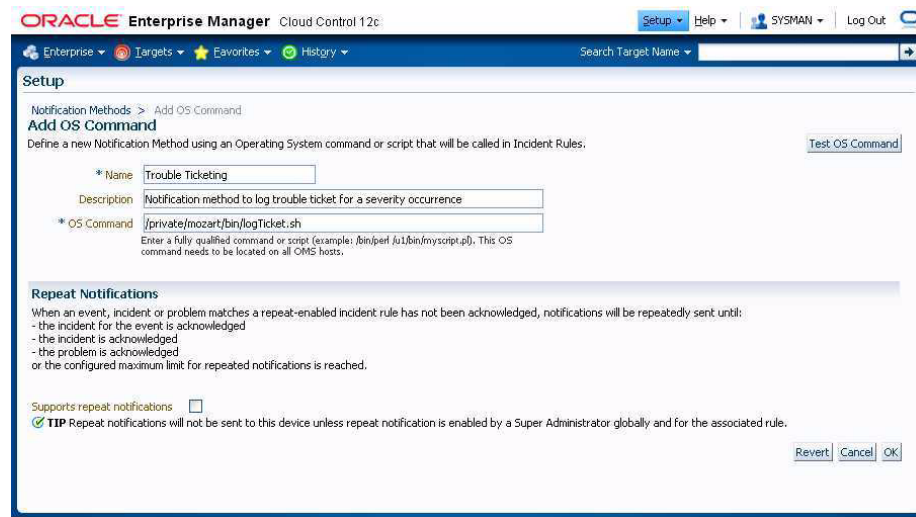
Example 4-5 OS Command Notification Method

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.

Figure 4–5 Notification Method: Add OS Command



Step 4: Assign the notification method to an incident 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.

For detailed reference information on passing event, incident, and problem information to an OS Command or script, see ["Passing Event, Incident, Problem Information to an OS Command or Script"](#) on page 4-52.

4.3.1 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_REPORTED_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"
```

4.3.2 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.

Note: Policy Violations are no longer supported beginning with the Enterprise Manager 12c release.

4.3.2.1 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 4–3 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
METRIC_VALUE	VALUE
VIOLATION_CONTEXT	EVENT_CONTEXT_ATTRS
VIOLATION_GUID	SEVERITY_GUID
POLICY_RULE	No mapping, Obsolete as of Enterprise Manager 12c release.

To obtain KEY_VALUE_NAME and KEY_VALUE, 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_COLUMN_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.

4.3.2.2 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 4–4 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	""

4.3.2.3 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 4–5 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
TARGET_TYPE1	TARGET_TYPE

Table 4–5 (Cont.) pre-12c/12c job_status_change Environment Variable Mappings

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

4.3.2.4 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 4–6 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

4.3.2.5 Notification Type Mapping

Table 4–7 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

4.4 Sending Notifications Using PL/SQL Procedures

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

Note: When upgrading from pre-12c to 12c versions of Enterprise Manager, existing pre-12c PL/SQL advanced notification methods will continue to work without modification. You should, however, update the procedures to use new signatures.

New PL/SQL advanced notification methods created with Enterprise Manager 12c must use the new signatures documented in the following sections.

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

4.4.1 Defining a PL/SQL-based Notification Method

Creating a PL/SQL-based notification method consists of four steps:

1. Define the PL/SQL procedure.
2. Create the PL/SQL procedure on the Management Repository.
3. Register your PL/SQL procedure as a new notification method.
4. Assign the notification method to an incident rule.

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 4-61

["Passing Corrective Action Status Change Information"](#) on page 4-44

["Passing Job Execution Status Information"](#) on page 4-46

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 4.4, "Sending Notifications Using PL/SQL Procedures"](#).

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

[Figure 4–6](#) illustrates how to add a PL/SQL-based notification method from the Enterprise Manager UI.

Figure 4–6 Adding a PL/SQL Procedure

ORACLE Enterprise Manager Cloud Control 12c

Setup ▾ Help ▾ SYSMAN ▾ Log Out

Enterprise ▾ Targets ▾ Favorites ▾ History ▾ Search Target Name

Setup

Notification Methods > Add PL/SQL Procedure

Add PL/SQL Procedure

Define a new Notification Method using a PL/SQL procedure that will be called in Incident Rules: Test PL/SQL Procedure

* Name:

Description:

* PL/SQL Procedure:

Enter a fully qualified procedure name (example: SCOTT.PRGRAME.PROCNAME). This PL/SQL procedure needs to be in the repository database.

Repeat Notifications

When an event, incident or problem matches a repeat-enabled incident rule has not been acknowledged, notifications will be repeatedly sent until:

Supports repeat notifications ☐

TIP Repeat notifications will not be sent to this device unless repeat notification is enabled by a Super Administrator globally and for the associated rule.

Revert Cancel OK

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.

See ["Passing Information to a PL/SQL Procedure"](#) on page 4-61 for more information about how incident, event, and problem information is passed to the PLSQL procedure.

Example 4–10 PL/SQL Script

```
-- Assume log_table is created by following DDL
-- CREATE TABLE log_table (message VARCHAR2(4000)) ;
-- Define PL/SQL notification method for Events
CREATE OR REPLACE PROCEDURE log_table_notif_proc(s IN GC$NOTIF_EVENT_MSG)
IS
    l_categories gc$category_string_array;
    l_category_codes gc$category_string_array;
    l_attrs gc$notif_event_attr_array;
    l_ca_obj gc$notif_corrective_action_job;
BEGIN
    INSERT INTO log_table VALUES ('notification_type: ' || s.msg_info.notification_
type);
    INSERT INTO log_table VALUES ('repeat_count: ' || s.msg_info.repeat_count);
    INSERT INTO log_table VALUES ('ruleset_name: ' || s.msg_info.ruleset_name);
    INSERT INTO log_table VALUES ('rule_name: ' || s.msg_info.rule_name);
    INSERT INTO log_table VALUES ('rule_owner: ' || s.msg_info.rule_owner);
    INSERT INTO log_table VALUES ('message: ' || s.msg_info.message);
    INSERT INTO log_table VALUES ('message_url: ' || s.msg_info.message_url);
    INSERT INTO log_table VALUES ('event_instance_guid: ' || s.event_payload.event_
instance_guid);
    INSERT INTO log_table VALUES ('event_type: ' || s.event_payload.event_type);
    INSERT INTO log_table VALUES ('event_name: ' || s.event_payload.event_name);
    INSERT INTO log_table VALUES ('event_msg: ' || s.event_payload.event_msg);
    INSERT INTO log_table VALUES ('source_obj_type: ' || s.event_
payload.source.source_type);
```



```

INSERT INTO log_table VALUES ('source_obj_name: ' || s.event_
payload.source.source_name);
INSERT INTO log_table VALUES ('source_obj_url: ' || s.event_
payload.source.source_url);
INSERT INTO log_table VALUES ('target_name: ' || s.event_payload.target.target_
name);
INSERT INTO log_table VALUES ('target_url: ' || s.event_payload.target.target_
url);
INSERT INTO log_table VALUES ('severity: ' || s.event_payload.severity);
INSERT INTO log_table VALUES ('severity_code: ' || s.event_payload.severity_
code);
INSERT INTO log_table VALUES ('event_reported_date: ' || to_char(s.event_
payload.reported_date, 'D MON DD HH24:MI:SS'));

l_categories := s.event_payload.categories;
IF l_categories IS NOT NULL
THEN
FOR c IN 1..l_categories.COUNT
LOOP
INSERT INTO log_table VALUES ('category ' || c || ' - ' || l_categories(c));
END LOOP;
END IF;

l_category_codes := s.event_payload.category_codes;
IF l_category_codes IS NOT NULL
THEN
FOR c IN 1..l_category_codes.COUNT
LOOP
INSERT INTO log_table VALUES ('category_code ' || c || ' - ' || l_category_
codes(c));
END LOOP;
END IF;

l_attrs := s.event_payload.event_attrs;
IF l_attrs IS NOT NULL
THEN
FOR c IN 1..l_attrs.COUNT
LOOP
INSERT INTO log_table VALUES ('EV.ATTR name=' || l_attrs(c).name || '
value=' || l_attrs(c).value || ' nls_value=' || l_attrs(c).nls_value);
END LOOP;
END IF;

COMMIT ;
END ;
/

```

Example 4-11 PL/SQL Script to Log Events to a Table

```

CREATE TABLE event_log (
notification_type    VARCHAR2(32),
repeat_count        NUMBER,
ruleset_name        VARCHAR2(256),
rule_owner          VARCHAR2(256),
rule_name           VARCHAR2(256),
message             VARCHAR2(4000),
message_url         VARCHAR2(4000),
event_instance_guid RAW(16),
event_type          VARCHAR2(20),

```

```
event_name          VARCHAR2(512),
event_msg           VARCHAR2(4000),
categories          VARCHAR2(4000),
source_obj_type     VARCHAR2(120),
source_obj_name     VARCHAR2(256),
source_obj_url      VARCHAR2(4000),
severity            VARCHAR2(128),
severity_code       VARCHAR2(32),
target_name         VARCHAR2(256),
target_type        VARCHAR2(128),
target_url          VARCHAR2(4000),
host_name           VARCHAR2(256),
timezone            VARCHAR2(64),
occured            DATE,
ca_guid             RAW(16),
ca_name             VARCHAR2(128),
ca_owner            VARCHAR2(256),
ca_type             VARCHAR2(256),
ca_status           VARCHAR2(64),
ca_status_code      NUMBER,
ca_job_step_output  VARCHAR2(4000),
ca_execution_guid   RAW(16),
ca_stage_change_guid RAW(16)
)
;

CREATE OR REPLACE PROCEDURE log_event(s IN GC$NOTIF_EVENT_MSG)
IS
    l_categories gc$category_string_array;
    l_ca_obj gc$notif_corrective_action_job;
    l_categories_new VARCHAR2(1000);
BEGIN
    -- save event categories
    l_categories := s.event_payload.categories;
    IF l_categories IS NOT NULL
    THEN
        FOR c IN 1..l_categories.COUNT
        LOOP
            l_categories_new := (l_categories_new|| c || ' - ' || l_
categories(c)||',');
        END LOOP;
    END IF;

    -- save event message
    IF s.msg_info.notification_type = 'NOTIF_CA' AND s.event_payload.corrective_
action IS NOT NULL
    THEN
        l_ca_obj := s.event_payload.corrective_action;
        INSERT INTO event_log (notification_type, repeat_count, ruleset_name, rule_
name, rule_owner, message, message_url, event_instance_guid, event_type, event_
name, event_msg, categories, source_obj_type, source_obj_name, source_obj_url,
severity, severity_code, target_name, target_type, target_url, host_name,
timezone, occured, ca_guid, ca_name, ca_owner, ca_type, ca_status, ca_status_code,
ca_job_step_output, ca_execution_guid, ca_stage_change_guid)
        VALUES (s.msg_info.notification_type, s.msg_info.repeat_count, s.msg_
info.ruleset_name, s.msg_info.rule_name,s.msg_info.rule_owner, s.msg_info.message,
s.msg_info.message_url, s.event_payload.event_instance_guid, s.event_
payload.event_type, s.event_payload.event_name, s.event_payload.event_msg, l_
categories_new, s.event_payload.source.source_type, s.event_payload.source.source_
name, s.event_payload.source.source_url, s.event_payload.severity, s.event_
```

```

payload.severity_code, s.event_payload.target.target_name, s.event_
payload.target.target_type, s.event_payload.target.target_url, s.event_
payload.target.host_name, s.event_payload.target.target_timezone, s.event_
payload.occurrence_date, l_ca_obj.JOB_GUID, l_ca_obj.JOB_NAME, l_ca_obj.JOB_OWNER,
l_ca_obj.JOB_TYPE, l_ca_obj.JOB_STATUS, l_ca_obj.JOB_STATUS_CODE, l_ca_obj.JOB_
STEP_OUTPUT, l_ca_obj.JOB_EXECUTION_GUID, l_ca_obj.JOB_STATE_CHANGE_GUID);
ELSE
    INSERT INTO event_log (notification_type, repeat_count, ruleset_name, rule_
name, rule_owner, message, message_url, event_instance_guid, event_type, event_
name, event_msg, categories, source_obj_type, source_obj_name, source_obj_url,
severity, severity_code, target_name, target_type, target_url, host_name,
timezone, occurred, ca_guid, ca_name, ca_owner, ca_type, ca_status, ca_status_code,
ca_job_step_output, ca_execution_guid, ca_stage_change_guid)
    VALUES (s.msg_info.notification_type, s.msg_info.repeat_count, s.msg_
info.ruleset_name, s.msg_info.rule_name, s.msg_info.rule_owner, s.msg_
info.message, s.msg_info.message_url, s.event_payload.event_instance_guid,
s.event_payload.event_type, s.event_payload.event_name, s.event_payload.event_msg,
l_categories_new, s.event_payload.source.source_type, s.event_
payload.source.source_name, s.event_payload.source.source_url, s.event_
payload.severity, s.event_payload.severity_code, s.event_payload.target.target_
name, s.event_payload.target.target_type, s.event_payload.target.target_url,
s.event_payload.target.host_name, s.event_payload.target.target_timezone, s.event_
payload.occurrence_date, null,null,null,null,null,null,null,null,null);
END IF;
COMMIT;
END log_event;
/

```

Example 4-12 PL/SQL Script to Log Incidents to a Table

```

CREATE TABLE incident_log (
    notification_type    VARCHAR2(32),
    repeat_count         NUMBER,
    ruleset_name         VARCHAR2(256),
    rule_owner           VARCHAR2(256),
    rule_name            VARCHAR2(256),
    message              VARCHAR2(4000),
    message_url          VARCHAR2(4000),
    incident_id          VARCHAR2(128),
    ticket_url           VARCHAR2(4000),
    assoc_event_cnt      NUMBER,
    severity             VARCHAR2(128),
    severity_code        VARCHAR2(32),
    priority             VARCHAR2(128),
    priority_code        VARCHAR2(32),
    status               VARCHAR2(32),
    categories           VARCHAR2(1000),
    target_name          VARCHAR2(256),
    target_type          VARCHAR2(128),
    host_name            VARCHAR2(256),
    timezone             VARCHAR2(64),
    occurred             DATE
)
;

CREATE OR REPLACE PROCEDURE log_incident(s IN GC$NOTIF_INCIDENT_MSG)
IS
    l_src_info_array GC$NOTIF_SOURCE_INFO_ARRAY;
    l_src_info GC$NOTIF_SOURCE_INFO;
    l_categories gc$category_string_array;

```

```

l_target_obj GC$NOTIF_TARGET;
l_target_name VARCHAR2(256);
l_target_type VARCHAR2(256);
l_target_timezone VARCHAR2(256);
l_hostname VARCHAR2(256);
l_categories_new VARCHAR2(1000);
BEGIN
    -- Save Incident categories
    IF l_categories IS NOT NULL
    THEN
        FOR c IN 1..l_categories.COUNT
        LOOP
            l_categories_new := (l_categories_new || c || ' - ' || l_
categories(c) || ',');
        END LOOP;
    END IF;

    -- GET target info
    l_src_info_array := s.incident_payload.incident_attrs.source_info_arr;
    IF l_src_info_array IS NOT NULL
    THEN
        FOR I IN 1..l_src_info_array.COUNT
        LOOP
            IF l_src_info_array(I).TARGET IS NOT NULL
            THEN
                l_target_name := l_src_info_array(I).TARGET.TARGET_NAME;
                l_target_type := l_src_info_array(I).TARGET.TARGET_TYPE;
                l_target_timezone := l_src_info_array(I).TARGET.TARGET_TIMEZONE;
                l_hostname := l_src_info_array(I).TARGET.HOST_NAME;
            END IF;
        END LOOP;
    END IF;

    -- save Incident notification message
    INSERT INTO incident_log(notification_type, repeat_count, ruleset_name, rule_
owner, rule_name, message, message_url, incident_id, ticket_url, assoc_event_cnt,
severity, severity_code, priority, priority_code, status, categories, target_name,
target_type, host_name, timezone, occurred)
    VALUES (s.msg_info.notification_type, s.msg_info.repeat_count, s.msg_
info.ruleset_name, s.msg_info.rule_owner, s.msg_info.rule_name, s.msg_
info.message, s.msg_info.message_url, s.incident_payload.incident_attrs.id,
s.incident_payload.ticket_url, s.incident_payload.assoc_event_count, s.incident_
payload.incident_attrs.severity, s.incident_payload.incident_attrs.severity_code,
s.incident_payload.incident_attrs.priority, s.incident_payload.incident_
attrs.priority_code, s.incident_payload.incident_attrs.STATUS, l_categories_new,
l_target_name, l_target_type, l_hostname, l_target_timezone, s.incident_
payload.incident_attrs.creation_date);
    COMMIT;
END log_incident;
/

```

Example 4-13 PL/SQL Script to Log Problems to a Table

```

CREATE TABLE problem_log (
    notification_type    VARCHAR2(32),
    repeat_count         NUMBER,
    ruleset_name         VARCHAR2(256),
    rule_owner           VARCHAR2(256),

```

```

rule_name          VARCHAR2(256),
message            VARCHAR2(4000),
message_url        VARCHAR2(4000),
problem_key        VARCHAR2(850),
assoc_incident_cnt NUMBER,
problem_id         NUMBER,
owner              VARCHAR2(256),
severity           VARCHAR2(128),
severity_code      VARCHAR2(32),
priority           VARCHAR2(128),
priority_code      VARCHAR2(32),
status             VARCHAR2(32),
categories         VARCHAR2(1000),
target_name        VARCHAR2(256),
target_type        VARCHAR2(128),
host_name          VARCHAR2(256),
timezone           VARCHAR2(64),
occured            DATE
)
;

CREATE OR REPLACE PROCEDURE log_problem(s IN GC$NOTIF_PROBLEM_MSG)
IS
  l_src_info_array GC$NOTIF_SOURCE_INFO_ARRAY;
  l_src_info       GC$NOTIF_SOURCE_INFO;
  l_categories      gc$category_string_array;
  l_target_obj      GC$NOTIF_TARGET;
  l_target_name     VARCHAR2(256);
  l_target_type     VARCHAR2(256);
  l_target_timezone VARCHAR2(256);
  l_hostname        VARCHAR2(256);
  l_categories_new  VARCHAR2(1000);
BEGIN
  -- Save Problem categories
  l_categories := s.problem_payload.problem_attrs.categories;
  IF l_categories IS NOT NULL
  THEN
    FOR c IN 1..l_categories.COUNT
    LOOP
      l_categories_new := (l_categories_new || c || ' - ' || l_
categories(c) || ',');
    END LOOP;
  END IF;

  -- GET target info
  l_src_info_array := s.problem_payload.problem_attrs.source_info_arr;
  IF l_src_info_array IS NOT NULL
  THEN
    FOR I IN 1..l_src_info_array.COUNT
    LOOP
      IF l_src_info_array(I).TARGET IS NOT NULL
      THEN
        l_target_name := l_src_info_array(I).TARGET.TARGET_NAME;
        l_target_type := l_src_info_array(I).TARGET.TARGET_TYPE;
        l_target_timezone := l_src_info_array(I).TARGET.TARGET_TIMEZONE;
        l_hostname := l_src_info_array(I).TARGET.HOST_NAME;
      END IF;
    END LOOP;
  END IF;

  -- save Problem notification message

```

```

INSERT INTO problem_log(notification_type, repeat_count, ruleset_name, rule_
owner, rule_name, message, message_url, problem_key, assoc_incident_cnt, problem_
id, owner, severity, severity_code, priority, priority_code, status, categories,
target_name, target_type, host_name, timezone, occurred)
VALUES (s.msg_info.notification_type, s.msg_info.repeat_count, s.msg_
info.ruleset_name, s.msg_info.rule_owner, s.msg_info.rule_name, s.msg_
info.message, s.msg_info.message_url, s.problem_payload.problem_key,
s.problem_payload.ASSOC_INCIDENT_COUNT, s.problem_payload.problem_attrs.id,
s.problem_payload.problem_attrs.owner, s.problem_payload.problem_attrs.severity,
s.problem_payload.problem_attrs.severity_code, s.problem_payload.problem_
attrs.PRIORITY, s.problem_payload.problem_attrs.PRIORITY_CODE, s.problem_
payload.problem_attrs.status, l_categories_new, l_target_name, l_target_type, l_
hostname, l_target_timezone, s.problem_payload.problem_attrs.CREATION_DATE);
COMMIT;
END log_problem;
/

```

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

Pre-12c notifications map to event notifications in Enterprise Manager 12c. The event types `metric_alert`, `target_availability` and `job_status_alert` correspond to the pre-12c notification functionality.

Note: Policy Violations are no longer available beginning with Enterprise Manager 12c.

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.

4.4.2.1 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 4–8 Metric Alert Mapping

MGMT_NOTIFY_SEVERITY	12c Notification Payload
TARGET_NAME	<code>gc\$notif_target.target_name</code>
TARGET_TYPE	<code>gc\$notif_target.target_type</code>
TIMEZONE	<code>gc\$notif_target.target_timezone</code>
HOST_NAME	<code>gc\$notif_target.host_name</code>
MERTIC_NAME	<code>gc\$notif_event_attr.value</code> where its name='metric_group' in <code>gc\$notif_event_attr_array</code> .
METRIC_DESCRIPTION	<code>gc\$notif_event_attr.value</code> where its name='metric_description' in <code>gc\$notif_event_attr_array</code> .

Table 4–8 (Cont.) Metric Alert Mapping

MGMT_NOTIFY_SEVERITY	12c Notification Payload
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 Table 4–9, "Severity Code Mapping" .
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' 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 illustrates how to obtain similar pre-12c KEY_VALUE and KEY_VALUE_NAME from an Enterprise Manager 12c notification payload.

Example 4–14 Extracting KEY_VALUE and KEY_VALUE_NAME

```
-- Get the pre-12c KEY_VALUE and KEY_VALUE_NAME from an Enterprise Manager 12c
-- notification payload
-- parameters
--   IN Parameters:
--       event_msg : The event notification payload
--   OUT Parameters
--       key_value_name_out : the KEY_VALUE_NAME backward compitable to pre-12c
--                           notification payload
--       key_value_out      : the KEY_VALUE backward compitable to pre-12c
--                           notification payload
--
CREATE OR REPLACE PROCEDURE get_pre_12c_key_value(
    event_msg IN GC$NOTIF_EVENT_MSG,
```

```
        key_value_name_out OUT VARCHAR2,
        key_value_out OUT VARCHAR2)
IS
    l_key_columns MGMT_SHORT_STRING_ARRAY := MGMT_SHORT_STRING_ARRAY();
    l_key_column_values MGMT_MEDIUM_STRING_ARRAY := MGMT_MEDIUM_STRING_ARRAY();
    l_key_value VARCHAR2(1790) := NULL;
    l_num_keys NUMBER := 0;
    l_attrs gc$notif_event_attr_array;
    l_key_value_name VARCHAR2(512);
BEGIN
    l_attrs := event_msg.event_payload.event_attrs;
    key_value_name_out := NULL;
    key_value_out := NULL;

    IF l_attrs IS NOT NULL AND
        l_attrs.COUNT > 0
    THEN
        l_key_columns.extend(7);
        l_key_column_values.extend(7);
        FOR c IN 1..l_attrs.COUNT
        LOOP
            CASE l_attrs(c).name
                WHEN 'num_keys' THEN
                    BEGIN
                        l_num_keys := to_number(l_attrs(c).value);
                    EXCEPTION
                        WHEN OTHERS THEN
                            -- should never happen, but guard against it l_num_keys := 0;
                    END;
                WHEN 'key_value' THEN
                    l_key_value := substr(l_attrs(c).nls_value,1,1290);
                WHEN 'key_column_1' THEN
                    l_key_columns(1) := substr(l_attrs(c).nls_value,1,64);
                WHEN 'key_column_2' THEN
                    l_key_columns(2) := substr(l_attrs(c).nls_value,1,64);
                WHEN 'key_column_3' THEN
                    l_key_columns(3) := substr(l_attrs(c).nls_value,1,64);
                WHEN 'key_column_4' THEN
                    l_key_columns(4) := substr(l_attrs(c).nls_value,1,64);
                WHEN 'key_column_5' THEN
                    l_key_columns(5) := substr(l_attrs(c).nls_value,1,64);
                WHEN 'key_column_6' THEN
                    l_key_columns(6) := substr(l_attrs(c).nls_value,1,64);
                WHEN 'key_column_7' THEN
                    l_key_columns(7) := substr(l_attrs(c).nls_value,1,64);
                WHEN 'key_column_1_value' THEN
                    l_key_column_values(1) := substr(l_attrs(c).nls_value,1,256);
                WHEN 'key_column_2_value' THEN
                    l_key_column_values(2) := substr(l_attrs(c).nls_value,1,256);
                WHEN 'key_column_3_value' THEN
                    l_key_column_values(3) := substr(l_attrs(c).nls_value,1,256);
                WHEN 'key_column_4_value' THEN
                    l_key_column_values(4) := substr(l_attrs(c).nls_value,1,256);
                WHEN 'key_column_5_value' THEN
                    l_key_column_values(5) := substr(l_attrs(c).nls_value,1,256);
                WHEN 'key_column_6_value' THEN
                    l_key_column_values(6) := substr(l_attrs(c).nls_value,1,256);
                WHEN 'key_column_7_value' THEN
                    l_key_column_values(7) := substr(l_attrs(c).nls_value,1,256);
```



```

        ELSE
            NULL;
        END CASE;
    END LOOP;

    -- get key_value and key_value_name when l_num_keys > 0

    IF l_num_keys > 0
    THEN
        -- get key value name
        IF l_key_columns IS NULL OR l_key_columns.COUNT = 0
        THEN
            key_value_name_out := NULL;
        ELSE
            l_key_value_name := NULL;
            FOR i IN l_key_columns.FIRST..l_num_keys
            LOOP
                IF i > 1
                THEN
                    l_key_value_name := l_key_value_name || ' ';
                END IF;
                l_key_value_name := l_key_value_name || l_key_columns(i);
            END LOOP;
            key_value_name_out := l_key_value_name;
        END IF;

        -- get key_value
        IF l_num_keys = 1
        THEN
            key_value_out := l_key_value;
        ELSE
            l_key_value := NULL;
            IF l_key_column_values IS NULL OR l_key_column_values.COUNT = 0
            THEN
                key_value_out := NULL;
            ELSE
                FOR i IN l_key_column_values.FIRST..l_num_keys
                LOOP
                    IF i > 1
                    THEN
                        l_key_value := l_key_value || ' ';
                    END IF;
                    l_key_value := l_key_value || l_key_column_values(i);
                END LOOP;
                -- max length for key value in pre-12c = 1290
                key_value_out := substr(l_key_value,1,1290);
            END IF;
        END IF;
    END IF; -- l_num_keys > 0
END IF; -- l_attrs IS NOT NULL
END get_pre_12c_key_value;
/

```

When the event type is metric_alert:

Use the following severity code mapping from 12c to pre-12c when the event type is *metric_alert*.

Table 4–9 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 4–10 Target Availability Mapping

MGMT_NOTIFY_SEVERITY	12c Notification Payload
TARGET_NAME	<code>gc\$notif_target.target_name</code>
TARGET_TYPE	<code>gc\$notif_target.target_type</code>
TIMEZONE	<code>gc\$notif_target.target_timezone</code>
HOST_NAME	<code>gc\$notif_target.host_name</code>
MERTIC_NAME	Use fixed value "Response".
METRIC_DESCRIPTION	NULL
METRIC_COLUMN	Use fixed value "Status".
METRIC_VALUE	<code>gc\$notif_event_attr.value</code> where its name='target_status' in <code>gc\$notif_event_attr_array</code> .
KEY_VALUE	NULL
KEY_VALUE_NAME	NULL
KEY_VALUE_GUID	NULL
CTXT_LIST	<code>gc\$notif_event_context_array</code>
COLLECTION_TIMESTAMP	<code>gc\$notif_event_payload.reported_date</code>
SEVERITY_CODE	<code>gc\$notif_event_attr.value</code> where its name='avail_severity' in <code>gc\$notif_event_attr_array</code> .
MESSAGE	<code>gc\$notif_msg_info.message</code>
SEVERITY_GUID	<code>gc\$notif_event_attr.value</code> where its name='severity_guid' in <code>gc\$notif_event_attr_array</code> .
METRIC_GUID	<code>gc\$notif_event_attr.value</code> where its name='metric_guid_id' in <code>gc\$notif_event_attr_array</code> .
TARGET_GUID	<code>gc\$notif_target.target_guid</code>
RULE_OWNER	<code>gc\$notif_msg_info.rule_owner</code>
RULE_NAME	<code>gc\$notif_msg_info.ruleset_name</code>

4.4.2.2 Mapping for MGMT_NOTIFY_JOB

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

Table 4–11 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
RULE_NAME	gc\$notif_msg_info.ruleset_name
OCCURRED_DATE	gc\$notif_event_payload.reported_date

4.4.2.3 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 4–8, "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 4–12 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

Table 4–12 (Cont.) Corrective Action Mapping

MGMT_NOTIFY_CORRECTIVE_ACTION	12c Notification Payload
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.

4.5 Sending SNMP Traps to Third Party Systems

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.

Important: In order for a third-party system to interpret traps sent by the OMS, the omstrap.v1 file must first be loaded into the third-party SNMP console. For more information about this file and its location, see ["MIB Definition"](#) on page 4-43.

The Enterprise Manager 12c version of the MIB file incorporates the 10g and 11g MIB content, thus ensuring backward compatibility with earlier Enterprise Manager releases.

The trap is an SNMP Version 1 trap and is described by the MIB definition [Appendix B, "Enterprise Manager MIB Definition."](#) See ["Management Information Base \(MIB\)"](#) on page 4-42 for an explanation of how the MIB works. If you are using Enterprise Manager 12c, see [Appendix A, "Interpreting Variables of the Enterprise Manager MIB"](#) and [Appendix B, "Enterprise Manager MIB Definition."](#) If you are upgrading from a pre-12c version of Enterprise Manager, see [Appendix C, "SNMP Trap Mappings"](#) for specific version mappings.

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

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. As shown in, the SNMP host will receive your SNMP traps.

Figure 4–7 *SNMP Trap Required Information*

ORACLE Enterprise Manager Cloud Control 12c

Setup ▾ Help ▾ SYSMAN ▾ Log Out

Enterprise ▾ Targets ▾ Favorites ▾ History ▾ Search Target Name ▾

Setup

Notification Methods > Add SNMP Trap

Add SNMP Trap

Define a new Notification Method so that SNMP traps can be sent by means of Incident Rules.

*** Name:** HP OpenView Console

Description: Notification method to send trap to HP OpenView console

*** SNMP Trap Host Name:** machine1.oracle.com

SNMP Trap Host Port: 162
If left blank, default port of 162 will be used.

SNMP Community: public
If left blank, the default community, public, will be used.

! Your notification method needs to be event compatible in order to be executed successfully.

Repeat Notifications

When an event matches a repeat enabled incident rule and its associated incident is not acknowledged, notification will be sent repeatedly until the configured maximum limit is reached.

Supports repeat notifications ☐

TIP Repeat notifications will not be sent to this device unless repeat notification is enabled by a Super Administrator globally and for the associated rule.

Revert Cancel OK

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

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

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. For instructions on setting up incident rules using SNMP traps, see ["Creating a Rule to Send SNMP Traps to Third Party Systems"](#) on page 3-54.

Example SNMP Trap Implementation

In this scenario, you want to identify the unique issues from the SNMP traps that are sent. Keep in mind that all events that are related to the same issue are part of the same event sequence. Each event sequence has a unique identification number.

An event sequence is a sequence of related events that represent the life of a specific issue from the time it is detected and an event is raised to the time it is fixed and a corresponding *clear* event is generated. For example, a warning metric alert event is raised when the CPU utilization of a host crosses 80%. This starts the event sequence representing the issue *CPU Utilization of the host is beyond normal level*. Another critical event is raised for the same issue when the CPU utilization goes above 90% and the event is added to the same event sequence. After a period of time, the CPU utilization returns to a normal level and a *clear* event is raised. At this point, the issue is resolved and the event sequence is closed.

The SNMP trap sent for this scenario is shown in [Example 4–15](#). Each piece of information is sent as a variable embedded in the SNMP Trap.

Example 4–15 SNMP Trap

```
*****V1 TRAP***[1]*****
Community : public
Enterprise :1.3.6.1.4.1.111.15.2
Generic :6
Specific :3
TimeStamp :67809
Agent address :10.240.36.109
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: CPU Utilization is 92.658%, crossed warning (80) or
critical (90) threshold.
1.3.6.1.4.1.111.15.3.1.1.4.1:
https://sampleserver.oracle.com:5416/em/redirect?pageType=sdk-core-event-console-d
etailEvent&issueID=C77AE9E578F00773E040F00A6D242F90
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
1.3.6.1.4.1.111.15.3.1.1.8.1:
1.3.6.1.4.1.111.15.3.1.1.9.1:
1.3.6.1.4.1.111.15.3.1.1.10.1: Aug 17, 2012 3:26:36 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:cpuUtil
1.3.6.1.4.1.111.15.3.1.1.15.1: 281
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: sampleserver.oracle.com
1.3.6.1.4.1.111.15.3.1.1.22.1:
https://sampleserver.oracle.com:5416/em/redirect?pageType=TARGET_
HOMEPAGE&targetName=sampleserver.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: sampleserver.oracle.com
1.3.6.1.4.1.111.15.3.1.1.25.1: SYSMAN
1.3.6.1.4.1.111.15.3.1.1.26.1:
1.3.6.1.4.1.111.15.3.1.1.27.1: 5.8.0.0.0
1.3.6.1.4.1.111.15.3.1.1.28.1: Operating System=Linux, Platform=x86_64,
1.3.6.1.4.1.111.15.3.1.1.29.1:
1.3.6.1.4.1.111.15.3.1.1.30.1:
1.3.6.1.4.1.111.15.3.1.1.31.1:
1.3.6.1.4.1.111.15.3.1.1.32.1:
1.3.6.1.4.1.111.15.3.1.1.33.1:
1.3.6.1.4.1.111.15.3.1.1.34.1:
1.3.6.1.4.1.111.15.3.1.1.35.1:
1.3.6.1.4.1.111.15.3.1.1.36.1:
1.3.6.1.4.1.111.15.3.1.1.37.1:
1.3.6.1.4.1.111.15.3.1.1.38.1:
1.3.6.1.4.1.111.15.3.1.1.39.1: SnmpNotifRuleset
1.3.6.1.4.1.111.15.3.1.1.40.1: SnmpNotifRuleset, SnmpNotifEvent
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: C77AE9E578F00773E040F00A6D242F90
1.3.6.1.4.1.111.15.3.1.1.43.1:
1.3.6.1.4.1.111.15.3.1.1.44.1:
```

```

1.3.6.1.4.1.111.15.3.1.1.45.1:
1.3.6.1.4.1.111.15.3.1.1.46.1: CPU Utilization is 92.658%, crossed warning (80) or
critical (90) threshold., Incident created by rule (Name = Incident management
Ruleset for all targets, Incident creation Rule for metric alerts.; Owner =
<SYSTEM>).
1.3.6.1.4.1.111.15.3.1.1.61.1: Metric GUID=0C71A1AFAC2D7199013837DA35522C08
1.3.6.1.4.1.111.15.3.1.1.62.1: Severity GUID=C77AE9E578EC0773E040F00A6D242F90
1.3.6.1.4.1.111.15.3.1.1.63.1: Cycle GUID=C77AE9E578EC0773E040F00A6D242F90
1.3.6.1.4.1.111.15.3.1.1.64.1: Collection 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=CPU 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: Metric value=92.658
1.3.6.1.4.1.111.15.3.1.1.69.1: Key Value=
1.3.6.1.4.1.111.15.3.1.1.70.1:
1.3.6.1.4.1.111.15.3.1.1.71.1:
1.3.6.1.4.1.111.15.3.1.1.72.1:
1.3.6.1.4.1.111.15.3.1.1.73.1:
1.3.6.1.4.1.111.15.3.1.1.74.1:
1.3.6.1.4.1.111.15.3.1.1.75.1:
1.3.6.1.4.1.111.15.3.1.1.76.1:
1.3.6.1.4.1.111.15.3.1.1.77.1:
1.3.6.1.4.1.111.15.3.1.1.78.1:
1.3.6.1.4.1.111.15.3.1.1.79.1:
1.3.6.1.4.1.111.15.3.1.1.80.1:
1.3.6.1.4.1.111.15.3.1.1.81.1:
1.3.6.1.4.1.111.15.3.1.1.82.1:
1.3.6.1.4.1.111.15.3.1.1.83.1:
1.3.6.1.4.1.111.15.3.1.1.84.1: Number of keys=0
1.3.6.1.4.1.111.15.3.1.1.85.1:
*****END V1 TRAP*****

```

This following example illustrates how OIDs are used during the lifecycle of an event. Here, for one event (while the event is open), the event sequence OID remains the same even though the event severity changes.

The OID for the event sequence is:

```
1.3.6.1.4.1.111.15.3.1.1.42.1: C77AE9E578F00773E040F00A6D242F90
```

The OID for the event severity code is:

```
1.3.6.1.4.1.111.15.3.1.1.6.1: CRITICAL
```

When the event clears, these OIDs show the same event sequence with a different severity code:

The OID for the event sequence is:

```
1.3.6.1.4.1.111.15.3.1.1.42.1: C77AE9E578F00773E040F00A6D242F90
```

The OID for the event severity code is:

```
1.3.6.1.4.1.111.15.3.1.1.6.1: CLEAR
```

The length of the SNMP OID value is limited to 2560 bytes by default. Configure the `emoms` property `oracle.sysman.core.notification.snmp.max_oid_length` to change the default limit.

4.5.1 SNMP Traps: Moving from Previous Enterprise Manager Releases to 12c

Note: When you upgrade from a pre-Enterprise Manager 12c release to 12c, SNMP advanced notification methods defined using previous versions of Enterprise Manager (pre-12c) will continue to function without modification.

For Enterprise Manager 11g and earlier, there were two types of SNMP traps:

- Alerts
- Job Status

For Enterprise Manager 12c there is now a single, comprehensive SNMP trap type that covers all available event types such as metric alerts, target availability, compliance standard violations, or job status changes. For more information about pre-12c to 12c SNMP trap mappings, see [Appendix C, "SNMP Trap Mappings."](#) Traps will conform to the older Enterprise Manager MIB definition. Hence, pre-Enterprise Manager 12c traps will continue to be sent. See [Appendix C, "SNMP Trap Mappings"](#) for more information.

Also, for Enterprise Manager 12c, size of SNMP trap has increased in order to accommodate all event types and provide more comprehensive information. By default, the maximum SNMP packet size is 5120 bytes. If the third party system has a limit in the size of SNMP trap it can receive, you can change the default size of SNMP trap that Enterprise Manager sends. To change the default packet size, set this *emoms* `oracle.sysman.core.notification.snmp_packet_length` parameter, and then bounce the OMS.

Note: When limiting the SNMP trap packet size, Oracle recommends not setting the `oracle.sysman.core.notification.snmp_packet_length` parameter any lower than 3072 bytes (3K).

The Enterprise Manager 12c MIB includes all pre-Enterprise Manager 12c MIB definitions. Hence, if you have an Enterprise Manager 12c MIB in your third party system, you can receive SNMP traps from both pre-Enterprise Manager 12c as well as Enterprise Manager 12c sites. For detailed information on version mapping, see [Appendix C, "SNMP Trap Mappings."](#)

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 MIB Definition

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

`OMS_HOME/network/doc/omstrap.v1`

The file *omstrap.v1* is the OMS MIB.

For more information, see [Appendix A, "Interpreting Variables of the Enterprise Manager MIB."](#)

A hardcopy version of *omstrap.v1* can be found in [Appendix B, "Enterprise Manager MIB Definition."](#)

The length of the SNMP OID value is limited to 2560 bytes by default. Configure *emoms* property `oracle.sysman.core.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.

4.6.3 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 1, is not included in these MIB variable descriptions. Since Oracle has implemented all MIB variables as CURRENT, this value does not vary.

4.6.3.1 Variable Name

Syntax

Maps to the SYNTAX element of SNMP MIB definition, Version 1.

Max-Access

Maps to the MAX-ACCESS element of SNMP MIB definition, Version 1.

Status

Maps to the STATUS element of SNMP MIB definition, Version 1.

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 1.

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.

4.7 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 may 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.

4.7.1 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 4–13 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

4.7.2 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 4–44, "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 4–14 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
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

Table 4–14 (Cont.) Corrective Action Status Codes

Name	Datatype	Value
SKIPPED_STATUS	NUMBER(2)	18
REASSIGNED_STATUS	NUMBER(2)	20

4.8 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 may 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 4.3, "Sending Notifications Using OS Commands and Scripts"](#) and [Section 4.4, "Sending Notifications Using PL/SQL Procedures"](#).

4.8.1 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 4–15](#) lists all corrective action status change attributes that can be passed:

Table 4–15 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.
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.

Table 4–15 (Cont.) Job Status Attributes

Attribute	Datatype	Additional Information
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 4–16 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
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

Table 4–16 (Cont.) Job Status Codes

Name	Datatype	Value
SKIPPED_STATUS	NUMBER(2)	18
REASSIGNED_STATUS	NUMBER(2)	20

Example 4–16 PL/SQL Procedure Using a Status Code (Job)

```

CREATE TABLE job_log (jobid RAW(16), status_code NUMBER(2), occurred 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;
    occurred_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;

        occurred_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, occurred)
            VALUES (job_guid, exec_status_code, occurred_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;
/

```

4.8.2 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 4–17 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.

4.9 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 4–8 Host Target Properties

The screenshot shows the 'Monitoring Configuration' dialog for a host target in Oracle Enterprise Manager. The dialog has a title bar with 'ORACLE Enterprise Manager Cloud Control 12c' and a 'Help' button. The main content area is divided into two sections: 'Properties' and 'Monitoring'. The 'Properties' section contains a table with columns 'Name' and 'Value'. The 'Monitoring' section contains a text box with information about automatic monitoring and a 'Cancel' button.

Name	Value
SNMP Community String (Default: public)	<input type="text"/>
SNMP Hostname	<input type="text"/>
SNMP Timeout (Default: 10 seconds)	<input type="text"/>
Host Username for WBEM Access	<input type="text"/>
Host Password for WBEM Access	<input type="text"/>
Port number for WBEM Access Default: 5988	<input type="text"/>
Disk Activity Metrics Collection Max Rows Upload(>0) Default:16	<input type="text"/>
Monitor Loopback Filesystems (true/false) Default:false	<input type="text"/>
Use pseudo-memory for Swap utilization (true/false) Default:true	<input type="text"/>

Monitoring
Oracle has automatically enabled monitoring for this target's availability and performance, so no further monitoring configuration is necessary. You can edit the metric thresholds from the target's homepage.

4.10 Notification Reference

This section contains the following reference material:

- [EMOMS Properties](#)
- [Passing Event, Incident, Problem Information to an OS Command or Script](#)
- [Passing Information to a PL/SQL Procedure](#)
- [Troubleshooting Notifications](#)

4.10.1 EMOMS Properties

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 4–18 *emoms Properties for Notifications*

Property Name	Default	
	Value	Description
oracle.sysman.core.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.
oracle.sysman.core.notification.commands_per_minute	100	OS Command delivery limits per minute. The Notification system uses this value to throttle number of OS Command delivery per minutes.
oracle.sysman.core.notification.os_command_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.
oracle.sysman.core.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.
oracle.sysman.core.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.
oracle.sysman.core.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.
oracle.sysman.core.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.

Table 4–18 (Cont.) emoms Properties for Notifications

Property Name	Default	
	Value	Description
oracle.sysman.core.notification.locale.oscmd	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.
oracle.sysman.core.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.
oracle.sysman.core.notification.oscmd.max_env_var_length	512	The maximum length of OS Common environment variable value.
oracle.sysman.core.notification.snmp.max_oid_length	2560	The maximum length of SNMP OID value.
oracle.sysman.core.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.
oracle.sysman.core.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.
oracle.sysman.core.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.
oracle.sysman.core.notification.snmp_packet_length	<=1 (5120)	The maximum size of SNMP Protocol Data unit.
oracle.sysman.core.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), and 7-bit(BASE64).
oracle.sysman.core.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.
oracle.sysman.core.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 `emctl get/set/delete/list property` command.

Note: The following commands require an OMS restart in order for the changes to take place.

Get Property Command

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

Set Property Command

```
emctl set property -name oracle.sysman.core.notification.short_format_length
-value 155
```

Emoms Properties Entries for a Short E-mail Format

```
emctl set property -name oracle.sysman.core.notification.short_format_length
-value 155
emctl set property -name oracle.sysman.core.notification.short_format -value both
```

4.10.2 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.

4.10.2.1 Environment Variables Common to Event, Incident and Problem

Table 4–19 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 4–20 Category-Related Environment Variables

Environment Variable	Description
CATEGORIES_COUNT	Number of categories in this notification. This value is equal to 1 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

Table 4–21 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 4–21 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

4.10.2.2 Event Notification-Specific Environment Variables

Table 4–22 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.

Table 4–22 (Cont.) Event Notification-Specific Environment Variables

Environment Variable	Description
LAST_UPDATED_TIME	Last updated time
SEQUENCE_ID	The unique event sequence identifier. An event sequence may consist of one or more events. All events in this sequence have the same event sequence ID.
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
RCA_DETAILS	If RCA is associated with this events.

Table 4–23 lists the environment variables for the incident associated with an event. They are populated when the event is associated with an incident.

Table 4–23 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.

Table 4–24 lists the common environment variables related to the Source Object. They are populated when \$SOURCE_OBJ_TYPE is not TARGET.

Table 4–24 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.

Table 4–24 (Cont.) Source Object-Related Environment Variables

Environment Variable	Description
SOURCE_OBJ_OWNER	Owner of the Source object.

Table 4–25 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 4–25 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.
TARGET_LIFECYCLE_STATUS	Life Cycle Status of the target. Possible values: Production, Mission Critical, Stage, Test, and Development. It is null if not defined.
TARGET_VERSION	Target Version of the target

4.10.2.3 Environment Variables Specific to Event Types

Events are classified into multiple types. For example, the `mertc_alert` event type is used for modeling metric alerts. You can use SQL queries to list the event types in your deployment as well as their event-specific payload. The following SQL example can be used to list all internal event type names that are registered in Enterprise Manager.

```
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;
```

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

Table 4–26 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.

Table 4–26 (Cont.) Environment Variables Specific to Metric Alert Event Type

Environment Variable	Description
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.
SEVERITY_GUID	The GUID of the severity record associated with this metric alert.
CYCLE_GUID	A unique identifier for a metric alert cycle, which starts from the time the metric alert is initially generated until the time it is clear.
VALUE	Value of the metric when the event triggered.

Table 4–27 Environment variables specific to Target Availability Event Type

Environment Variable	Description
AVAIL_SEVERITY	<p>The transition severity (0-6) that resulted in the status of the target to change to the current availability status.</p> <p>Possible Values for AVAIL_SEVERITY</p> <ul style="list-style-type: none"> ■ 0 (Target Down) ■ 1 (Target Up) ■ 2 (Target Status Error) ■ 3 (Agent Down) ■ 4 (Target Unreachable) ■ 5 (Target Blackout) ■ 6 (Target Status Unknown)
AVAIL_SUB_STATE	The substatus of a target for the current status.
CYCLE_GUID	A unique identifier for a metric alert cycle, which starts from the time the metric alert is initially generated until the time it is clear.
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 4–28 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. For possible values, see Table 4–16, "Job Status Codes" .
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 Enterprise 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.

4.10.2.4 Environment Variables Specific to Incident Notifications

Table 4–29 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, High, Medium, Low, None
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

[Table 4–30](#) lists the associated problem's environment variables, when the incident is associated with a problem.

Table 4–30 Associated Problem Environment Variables for Incidents

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.

4.10.2.5 Environment Variables Specific to Problem Notifications

Table 4–31 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.
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..

Table 4–31 (Cont.) Problem-Specific Environment Variables

Environment Variable	Description
PRIORITY_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.

4.10.2.6 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.

4.10.2.6.1 Environment Variables Related to Event Sources The number of event sources is set by the 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 4–32 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.

Table 4–33 lists the environment variables for a target of xth Event Source.

Table 4–33 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.

4.10.3 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 three 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 three 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 4–34 Event Notification Payload

Attribute	Datatype	Additional Information
EVENT_PAYLOAD	<code>gc\$notif_event_payload</code>	Event notification payload. See <code>gc\$notif_event_payload</code> type definition for detail.
MSG_INFO	<code>gc\$notif_msg_info</code>	Notification message. See <code>gc\$notif_msg_info</code> 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 4–35 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 4–36 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 4–37 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	VARCHA2(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 4–38 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.
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*)

Contains the incident specific attributes, associated problem and ticket information.

Table 4–39 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*)

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

Table 4–40 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.
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*)

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

Table 4–41 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

Table 4–41 (Cont.) Payload Common to Incidents and Problems

Attribute	Datatype	Additional Information
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.
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 \$gcnofif_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 incident and problem payloads)

Represents the associated incident summary in the event payload, or associated problem summary in the incident payload, respectively.

Table 4–42 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.

Table 4–42 (Cont.) Payload

Attribute	Datatype	Additional Information
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
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. The 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. Each element of this array is of the type *gc\$notif_event_context*.

gc\$notif_event_context: This object represents the detail of event context data which is additional contextual information captured by the source system at the time of event generation that may have diagnostic value. The context for an event should consist of a set of keys and values along with data type (Number or String only).

Table 4–43 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

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 4–44 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 Enterprise Manager. For more information on status codes, see Table 4–14, "Corrective Action Status Codes" .
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

Provides access to the multiple sources to which an incident or a problem could be related. 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 4–45 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

Used for referencing source objects other than a job target.

Table 4–46 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 4–47 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

Array of *gc\$notif_udtp* type with a maximum size of 20.

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

gc\$notif_udtp

Used for referencing User-defined target properties. UDTP should consist of a set of property key names and property values.

Table 4–48 Payload

Attribute	Datatype	Additional Information
NAME	VARCHAR2(64),	The name of property.
VALUE	VARCHAR2(1024)	Property value.
LABEL	VARCHAR(256)	Property label.
NLS_ID	VARCHAR(64)	Property nls id

4.10.3.1 Notification Payload Elements Specific to Event Types**gc\$notif_event_attr_array**

Array of *gc\$notif_event_attr* is used for referencing event-specific attributes. The array has a maximum size of 25. Each element of the array is of type *gc\$notif_event_attr* (used for referencing event type-specific attributes).

Table 4–49 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.

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

```
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
order by event_type, env_var_name;
```

You should convert the attribute name to upper case before using the name for comparison.

There is an attribute variable payload specific to each event type that can be accessed from a *gc\$notif_event_attr_array* database type. The following tables list notification attributes for the most critical event types. You should convert the attribute name to uppercase before using the name for comparison.

Table 4–50 Environment variables specific to Metric Alert Event Type

Environment Variable	Description
COLL_NAME	The name of the collection collecting the metric.

Table 4–50 (Cont.) Environment variables specific to Metric Alert Event Type

Environment Variable	Description
KEY_COLUMN_X	Internal 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_DESCRIPTION	Brief description of the metric.
METRIC_GROUP	The name of the metric.
NUM_KEYS	The number of key metric columns in the metric.
SEVERITY_GUID	The GUID of the severity record associated with this metric alert.
VALUE	Value of the metric when the event triggered.

Table 4–51 Environment variables specific to Target Availability Event Type

Environment Variable	Description
AVAIL_SEVERITY	<p>The transition severity (0-6) that resulted in the status of the target to change to the current availability status.</p> <p>Possible Values for AVAIL_SEVERITY</p> <ul style="list-style-type: none"> ■ 0 (Target Down) ■ 1 (Target Up) ■ 2 (Target Status Error) ■ 3 (Agent Down) ■ 4 (Target Unreachable) ■ 5 (Target Blackout) ■ 6 (Target Status Unknown)
AVAIL_SUB_STATE	The substatus of a target for the current status.
CYCLE_GUID	A unique identifier for a metric alert cycle, which starts from the time the metric alert is initially generated until the time it is clear.
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.

Table 4–52 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.
EXEC_STATUS_CODE	Execution status code of job execution. For possible values, see Table 4–16, "Job Status Codes" .
STATE_CHANGE_GUID	Unique ID of last status change

Example 4-17 PL/SQL Script: Event Type Payload Elements

```

-- log_table table is created by following DDL to demonstrate how to access
-- event notification payload GC$NOTIF_EVENT_MSG.

CREATE TABLE log_table (message VARCHAR2(4000)) ;

-- Define PL/SQL notification method for Events
CREATE OR REPLACE PROCEDURE log_table_notif_proc(s IN GC$NOTIF_EVENT_MSG)
IS
    l_categories gc$category_string_array;
    l_category_codes gc$category_string_array;
    l_attrs gc$notif_event_attr_array;
    l_ca_obj gc$notif_corrective_action_job;
BEGIN
    INSERT INTO log_table VALUES ('notification_type: ' || s.msg_info.notification_
type);
    INSERT INTO log_table VALUES ('repeat_count: ' || s.msg_info.repeat_count);
    INSERT INTO log_table VALUES ('ruleset_name: ' || s.msg_info.ruleset_name);
    INSERT INTO log_table VALUES ('rule_name: ' || s.msg_info.rule_name);
    INSERT INTO log_table VALUES ('rule_owner: ' || s.msg_info.rule_owner);
    INSERT INTO log_table VALUES ('message: ' || s.msg_info.message);
    INSERT INTO log_table VALUES ('message_url: ' || s.msg_info.message_url);
    INSERT INTO log_table VALUES ('event_instance_guid: ' || s.event_payload.event_
instance_guid);
    INSERT INTO log_table VALUES ('event_type: ' || s.event_payload.event_type);
    INSERT INTO log_table VALUES ('event_name: ' || s.event_payload.event_name);
    INSERT INTO log_table VALUES ('event_msg: ' || s.event_payload.event_msg);
    INSERT INTO log_table VALUES ('source_obj_type: ' || s.event_
payload.source.source_type);
    INSERT INTO log_table VALUES ('source_obj_name: ' || s.event_
payload.source.source_name);
    INSERT INTO log_table VALUES ('source_obj_url: ' || s.event_
payload.source.source_url);
    INSERT INTO log_table VALUES ('target_name: ' || s.event_payload.target.target_
name);
    INSERT INTO log_table VALUES ('target_url: ' || s.event_payload.target.target_
url);
    INSERT INTO log_table VALUES ('severity: ' || s.event_payload.severity);
    INSERT INTO log_table VALUES ('severity_code: ' || s.event_payload.severity_
code);
    INSERT INTO log_table VALUES ('event_reported_date: ' || to_char(s.event_
payload.reported_date, 'D MON DD HH24:MI:SS'));

    IF s.event_payload.target.TARGET_LIFECYCLE_STATUS IS NOT NULL
    THEN
        INSERT INTO log_table VALUES ('target lifecycle_status: ' || s.event_
payload.target.TARGET_LIFECYCLE_STATUS);
    END IF;

    -- Following block illustrates the list of category codes to which the event
    -- belongs.

    l_category_codes := s.event_payload.category_codes;
    IF l_categories IS NOT NULL
    THEN
        FOR c IN 1..l_category_codes.COUNT
        LOOP
            INSERT INTO log_table VALUES ('category_code ' || c || ' - ' || l_category_
codes(c));
        END LOOP;
    END IF;

```

```

END IF;

--
-- Each event type has a specific set of attributes modeled. Examples of
-- event types include metric_alert, target_availability, job_status_change.
-- Following block illustrates how to access the attributes for job_status
change
-- event type
--
IF s.event_payload.event_type = 'job_staus_chage'
THEN
  l_attrs := s.event_payload.event_attrs;
  IF l_attrs IS NOT NULL
  THEN
    FOR c IN 1..l_attrs.COUNT
    LOOP
      INSERT INTO log_table VALUES ('EV.ATTR name=' || l_attrs(c).name || '
value=' || l_attrs(c).value || ' nls_value=' || l_attrs(c).nls_value);
    END LOOP;
  END IF;
END IF;

-- Following block illustrates how to access corrective action job's attributes
IF s.msg_info.notification_type = GC$NOTIFICATION.NOTIF_CA AND s.event_
payload.corrective_action IS NOT NULL
THEN
  l_ca_obj := s.event_payload.corrective_action;
  INSERT INTO log_table VALUES ('CA JOB_GUID: ' || l_ca_obj.JOB_GUID);
  INSERT INTO log_table VALUES ('CA JOB_NAME: ' || l_ca_obj.JOB_NAME);
  INSERT INTO log_table VALUES ('CA JOB_OWNER: ' || l_ca_obj.JOB_OWNER);
  INSERT INTO log_table VALUES ('CA JOB_TYPE: ' || l_ca_obj.JOB_TYPE);
  INSERT INTO log_table VALUES ('CA JOB_STATUS: ' || l_ca_obj.JOB_STATUS);
  INSERT INTO log_table VALUES ('CA JOB_STATUS_CODE: ' || l_ca_obj.JOB_STATUS_
CODE);
  INSERT INTO log_table VALUES ('CA JOB_STEP_OUTPUT: ' || l_ca_obj.JOB_STEP_
OUTPUT);
  INSERT INTO log_table VALUES ('CA JOB_EXECUTION_GUID: ' || l_ca_obj.JOB_
EXECUTION_GUID);
  INSERT INTO log_table VALUES ('CA JOB_STATE_CHANGE_GUID: ' || l_ca_obj.JOB_
STATE_CHANGE_GUID);
  INSERT INTO log_table VALUES ('CA OCCURRED_DATE: ' || l_ca_obj.OCCURRED_DATE);
END IF;

COMMIT ;
END ;
/

```

4.10.4 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.

4.10.4.1 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, PL/SQL 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.

4.10.4.2 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.

4.10.4.3 Notification System Trace Messages

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

Tracing is configured by setting the `log4j.category.oracle.sysman.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 -name log4j.category.oracle.sysman.em.notification -value
DEBUG -module logging
```

Note: The system will prompt you for the SYSMAN password.

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

```

4.10.4.4 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: rgmemeasmtplib.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

4.10.4.5 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.

4.10.4.6 SNMP Trap Errors

Use the Test button to make sure SNMP Trap configuration is valid.

The OMS will not report an error if the SNMP trap cannot reach the third party SNMP console as this is sent via UDP. If the SNMP trap encounters problems when trying to reach the third party SNMP console, possible SNMP trap problems include: invalid host name, port, community for a machine running an SNMP Console or a network issue such as a firewall problem.

Other possible SNMP trap problems include: invalid host name, port, or community for a machine running an SNMP Console.

4.10.4.7 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 plsqli_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 plsqli_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 plsqli_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.

Managing Groups

This chapter introduces the concept of group management and contains the following sections:

- [Introduction to Groups](#)
- [Managing Groups](#)
- [Using Out-of-Box Reports](#)

5.1 Introduction to Groups

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

Group Management Tasks

You can use Enterprise Manager to perform the following group management functions:

- Edit the configuration of a selected group, remove groups, and, in the case of an Administration Group, associate or disassociate a Template Collection.
- View the status and health of the group from the System Dashboard
- Drill down from a specific group to collectively monitor and manage its member targets.
- View a roll-up of member statuses and open incidents for members of the group.
- Apply blackouts to all targets in a group.
- Run jobs against a group
- Run a report
- Apply monitoring templates
- Associate compliance standards

In addition to creating groups, you can also create specific types of groups, such as redundancy groups, privilege propagating groups, and dynamic groups. The following sections explain the different types of groups.

5.1.1 Overview of Groups

Groups enable you to collectively monitor and administer many targets as a single logical unit. For example, you can define a group to contain all the databases serving an enterprise application, and define another group to contain all the hosts in a host farm. You can then use these groups to perform administrative operations. To create a group, you can manually select and add the members of the group. If you add an aggregate target, such as a Cluster Database, all of its member targets are automatically added to the group.

A group can include targets of the same type, such as all your production databases, or it could include all the targets on a host which would be comprised of different target types. You can nest static groups inside each other. In the target selector when you are selecting group members, choose Group as the target type, or choose a parent group as part of the process of creating a group.

If a system target is added to a group, it automatically pulls in its member targets. This could be the case of a regular group where a system such as WebLogic Server is added and also pulls in its members, or in a dynamic group where you specify a Target type to be an Oracle WebLogic Server and it also pulls in members of the WebLogic Server even though it does not match the dynamic group criteria. In this scenario, the group operations (for example, running jobs, blackouts, and so on) apply to all members of the group.

After you configure a group, you can perform various administrative operations, such as:

- View a summary status of the targets within the group.
- View a roll-up of member statuses and open incidents for members of the group.
- View a summary of critical patch advisories.
- View configuration changes during the past 7 days.
- Create jobs and view the status of job executions.
- Create blackouts and view the status of current blackouts.

5.1.2 Overview of Privilege Propagating Groups

Privilege propagating groups enable administrators to propagate privileges to members of a group. You can grant a privilege on a group once to an administrator or a role and have that same privilege automatically propagate to any new member of the group. For example, granting *operator* privilege on a privilege propagating 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. Any aggregate that you add to a privilege propagating group must also be privilege propagating as well. For example, any group that you add to a privilege propagating group must also be privilege propagating.

Privileges on the group can be granted to an Enterprise Manager administrator or a role. Use a role if the privileges you want to grant are to be granted to a group of Enterprise Manager administrators.

For example, suppose you create a 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.

5.1.3 Overview of Dynamic Groups

The membership management for groups is typically manual or static in nature. Manually managing memberships works well for small deployments but not necessarily in large, dynamic environments where new targets come into the system frequently. Groups whose members are added frequently would be easier to maintain if they were to be defined by membership criteria instead of adding targets directly into the group. When the membership criteria is defined once, Enterprise Manager will automatically add targets.

A dynamic group is a group whose membership is determined by membership criteria. The owner of a dynamic group specifies the membership criteria during dynamic group creation (or modification) and membership in the group is determined solely by the criteria specified. Membership in a dynamic group cannot be modified directly because targets cannot be directly added to a dynamic group. Enterprise Manager automatically adds targets that match membership criteria when a dynamic group is created. It also updates group membership as new targets are added or target properties are changed and the target matches the group's membership criteria.

It is important to note that static groups can contain dynamic groups as members but not the other way around. You cannot include a static group as a member of a dynamic group.

Use the Define Membership Criteria function of Dynamic Groups to define the criteria for group membership. Once you have defined criteria, the targets selected by the criteria will be displayed in a read-only table in the Members region of the Groups page. Since dynamic groups are defined by criteria, you can intentionally or unintentionally define criteria that could result in very large groups.

The following requirements apply to dynamic groups:

- Dynamic groups cannot contain static groups, other dynamic groups, or administration groups.
- Administration groups cannot contain dynamic groups, however, a static group can contain dynamic groups as a member.
- OR-based criteria is not supported. All criteria selected on the criteria page are AND-based.
- Supported properties are limited to global properties plus other attributes specifically supported for administration groups such as Version, Platform, Target Name and Type, and so on. Specifically user-defined properties and other instance properties, plus config data elements are not supported as criteria.
- The View Any Target and Add Any Target privileges are required to create a dynamic group.
- The Full Any Target, Add Any Target, and Create Privilege Propagating Group privileges are required to create a privilege-propagating dynamic group.

5.1.4 Overview of 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.

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.

5.2 Managing Groups

By combining targets in a group, Enterprise Manager provides management features that enable you to efficiently manage these targets as one group. Using the Group functionality, you can:

- View a summary status of the targets within the group.
- Monitor 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.

When you choose Groups from the Targets menu in the Enterprise Manager, the Groups page appears. 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 member status summary and rollup of incidents for members in a group.
- Create Groups, Dynamic Groups or the Administration Group hierarchy, edit the configuration of a selected group, remove groups, and, in the case of an Administration Group, associate or disassociate a Template Collection.
- 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.
- Customize the homepage of a specific group

Redundancy systems and special high availability groups are not accessed from this Groups page. You can access them from the All Targets page or you can access Redundancy Systems and other systems from the Systems page.

5.2.1 Creating and Editing Groups

Enterprise Manager Groups enable administrators to logically organize distributed targets for efficient and effective management and monitoring.

To create a group, follow these steps:

1. From the Enterprise Manager Console, choose **Targets** then choose **Groups**. Alternately, you can choose **Add Target** from the **Setup** menu and choose the menu option to add the specific type of group.
2. Click **Create**. The Enterprise Manager Console displays a set of Create Group pages that function similarly to a wizard.
3. On the General tab of the Create Group page, enter the **Name** of the Group you want to create. If you want to make this a privilege propagating group, then enable the Privilege Propagation option by clicking **Enabled**. If you enable Privilege Propagation for the group, the target privileges granted on the group to an administrator or a role are propagated to the member targets. As with regular groups with privilege propagation, the Create Privilege Propagating Group privilege is required for creation of privilege propagating dynamic groups. In addition, the Full any Target privilege is required to enable privilege propagation because only targets on which the owner has Full Target privileges can be members, and any target can potentially match the criteria and a system wide Full privilege is required. To create a regular dynamic group, the View any Target system wide privilege is required as the group owner must be able to view any target that can potentially match the membership criteria.
4. Configure each page, then click **OK**. You should configure all the pages before clicking **OK**. For more information about those steps, see the online help.

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

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

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

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

5.2.2 Creating Dynamic Groups

The owner of a dynamic group specifies the membership criteria during dynamic group creation (or modification) and membership in the group is determined solely by the criteria specified. Membership in a dynamic group cannot be modified directly. Enterprise Manager automatically adds targets that match membership criteria when a dynamic group is created. It also updates group membership as new targets are added or target properties are changed and the targets match the group's membership criteria.

To create a dynamic group, follow these steps:

1. From the Groups page, click **Create** and then select **Dynamic Group** from the drop-down list. Alternately, you can choose **Add Target** from the **Setup** menu and then select **Group**.
2. On the General tab of the Create Dynamic Group page, enter the **Name** of the Dynamic Group you want to create. If you want to make this a privilege propagating dynamic group, then enable the Privilege Propagation option by clicking **Enabled**. If you enable Privilege Propagation for the group, the target

privileges granted on the group to an administrator or a role are propagated to the member targets. As with regular groups with privilege propagation, the Create Privilege Propagating Group privilege is required for creation of privilege propagating dynamic groups. In addition, the Full any Target privilege is required to enable privilege propagation because only targets on which the owner has Full Target privileges can be members, and any target can potentially match the criteria and a system wide Full privilege is required. To create a regular dynamic group, the View any Target system wide privilege is required as the group owner must be able to view any target that can potentially match the membership criteria.

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

Grant this privilege to an administrator or role that enables him to become group administrator for the group. This means he can perform operations on the group, share privileges on the group with other administrators, etc.

The Group Administration Privilege is available for both Privilege Propagating Groups and conventional groups. If you are granted this privilege, you can grant full privilege access to the group to other Enterprise Manager users without having to be the SuperAdministrator to grant the privilege.

3. In the Define Membership Criteria section, define the criteria for the dynamic group membership by clicking **Define Membership Criteria**.

The Define Membership Criteria page appears where you can Add or Remove properties of targets to be included in the group. Group members must match one value in each of the populated target properties. Use the Member Preview section to review a list of targets that match the criteria. Click **OK** to return to the General page.

At least one of the criteria on the Define Membership Criteria page must be specified. You cannot create a Dynamic group without at least one of the target types, on hosts or target properties specified. Use the following criteria for dynamic groups:

- Target type(s)
- Department
- On Host
- Target Version
- Lifecycle Status
- Operating System
- Line of Business
- Platform
- Location
- CSI
- Cost Center

- Contact
- Comment

You can add or remove properties using the **Add** or **Remove** Target Properties button on the Define Membership Criteria page.

4. Enter the **Time Zone**. The time zone you select is used for scheduling operations such as jobs and blackouts on this group. The groups statistics charts will also use this time zone.
5. Click the **Charts** tab. Specify the charts that will be shown in the Dynamic Group Charts page. By default, the commonly used charts for the target types contained in the Dynamic Group are added.
6. Click the **Columns** tab to add columns and abbreviations that will be seen in the Members page and also in the Dashboard.
7. Click the **Dashboard** tab to specify the parameters for the System Dashboard. The System Dashboard displays the current status and incidents and compliance violations associated with the members of the Dynamic Group in graphical format.
8. Click the **Access** tab. Use the Access page to administer access privileges for the group. On the Access page you can grant target access to Enterprise Manager roles and grant target access to Enterprise Manager administrators.
9. Click **OK** to create the Dynamic Group.

5.2.3 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-aggregate target as the member of a privilege propagating group. For aggregate targets in Cloud Control version 12c, cluster and RAC databases and other propagating groups can be added as members. Cloud Control version 12c supports more aggregate target types, such as redundancy systems, systems and services.

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.

5.2.4 Converting Conventional Groups to Privilege Propagating Groups

In Enterprise Manager release 12c 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_aggregate_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*.

5.2.5 Viewing and Managing Groups

Enterprise Manager 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 Group Home page provides the following sections:

- A General section that displays the general information about the group, such as the Owner, Group Type, and whether the group is privilege propagating. You can drill down to the Edit Group page to enable or disable privilege propagating by clicking on the Privilege Propagating field.
- 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.
- An Overview of Incidents and Problems section that displays the summary of incidents on members of the group that have been updated in the recent period of time. It also shows a count of open problems as well as problems updated in recent period of time.

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 the compliance of members of the group against the compliance standards defined for the group. This section also shows a rollup of violations by severity (critical, warning, minor warning) as well as the average compliance score(%).

- A Job Activity section that displays a summary of jobs for the targets in the group whose start date is within the last 7 days. You can click Show to see the latest run or all runs. Click View to select and reorder the columns that appear in the table or to adjust scrolling and expanding the table.
- A Blackouts section that displays information about current or pending blackouts. You can also create a blackout from this section.
- A Patch Recommendations section that displays the Oracle patch recommendations that are applicable to your enterprise. You can view patch recommendations by classification or target type.

You can navigate to My Oracle Support to view all recommendations by clicking the All Recommendations link.

- 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 also view inventory summary information in the context of different dimensions. From here you can click See Details to display the Inventory and Usage page.
- A Configuration Changes section that displays the number of configuration changes to the group in the previous 7 days. You can click the number to display a page that displays detailed information about the changes. Enterprise Manager automatically collects configuration information for group targets and changes to configurations are recorded and may be viewed from that page.

Viewing a Group

To view a group, follow these steps:

1. From the Enterprise Manager Console, choose **Targets** then choose **Groups**. A summary table lists all defined groups.
2. Click the desired group to go to the Home page of that group.

You can use View By filters (located in the upper right corner of the home page) to change the view of the homepage to members of targets of a specific type. When you do this, the Group homepage refreshes to only show information for targets of that type. Additional regions of interest might display. For example, DBAs might switch to the Database filter to view information specifically on Database targets in the group.

You can also personalize the home page by clicking the Actions icon in the upper right corner of each region on the home page to move that region up or down on the page. You can also expand or contract a region by clicking the arrow icon in the upper left corner of each region.

You can also navigate to other management operations on the group using the Group menu. For example, you can view all the members in a group by choosing **Member** from the Group menu. Likewise you can view the **Membership History** of the group by choosing Membership History from the Group menu.

5.2.6 Overview of Group Charts

Group Charts enable you to monitor the collective performance of a 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.

5.2.7 Overview of Group Members

Enterprise Manager allows you to summarize information about the member targets in a group. It provides information on their current availability status, roll-up of open incidents and compliance violations, 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 rank members by a certain criterion (for example, database targets in order of decreasing wait time percentage). You can display default key performance metrics based on the targets you select, but you can customize these to include additional metrics that are important for managing your group.

You can view the members of a group by choosing **Members** from the Group menu. Enterprise Manager displays the Members page where you can view the table of members filtered by All Members, Direct Members, or Indirect Members. Direct members are targets directly added to the group. Indirect members are targets that are members of a direct member target, and are automatically included into the group because their parent target was added to the group. The page provides the option to **Export** or **Edit** the group.

You can also access information about membership history by choosing **Membership History** from the Group menu. The Membership History page displays changes in the group membership over time.

5.2.8 Viewing Group Status History

You can view Status History for a group to see the historical availability of a member during a specified time period or view the current status of all group members. You can access the Status History page by choosing **Monitoring** from the Group menu and then selecting **Status History**.

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 of a member, you can click a Status icon on the View Group Status History page 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.

5.2.9 About the System Dashboard

The System Dashboard enables you to proactively monitor the status, incidents and compliance violations in the group as they occur. The color-coded interface is designed to highlight problem areas — 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 determine the problem areas for any target and drill down for details as needed. An incident table is also included to provide a summary for all open incidents in the group. The incidents in the table are presented in reverse chronological order to show the most recent incidents first, but you can also click any column in the table to change the sort order. The colors in top bar of the Member Targets table change based on the incident's critical level. The priority progresses from

warning to critical to fatal. If the group has at least one fatal incident (irrespective of critical or warning incidents), the top bar becomes dark red. If the group has at least one critical incident (irrespective of warning incidents), the top bar becomes faint red. If the group has only warning incidents, the top bar turns yellow. If the group has no incidents, the top bar remains colorless.

The Dashboard auto-refreshes based on the Refresh Frequency you set on the Customize Dashboard page.

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
- Status icon for a metric with key values to access the metric page with a list of all key values
- Dashboard header to access the group home page
- Incidents and Problems table to view summary information about all incidents or specific categories of incidents.

Click **Customize** to access the Customize Dashboard page. This page allows you to change the refresh frequency and display options for the Member Targets table at the top of the dashboard. You can either show all individual targets or show by target type. There is also the option to expand or contract the Incidents and Problems table at the bottom. To change the columns shown in the Member Targets table, go to the Columns tab of the Edit Group page which you can access by choosing Target Setup from the Group menu.

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 incidents displayed are rolled up for the targets in that specific target type.

If you minimize the dashboard window, pertinent alert information associated with the group or system is still displayed in the Microsoft Windows toolbar.

You can use Information Publisher reports to make the System Dashboard available to non-Enterprise Manager users. First, create a report and include the System Monitoring Dashboard reporting element. In the report definition, choose the option, Allow viewing without logging in to Enterprise Manager. Once this is done, you can view it from the Enterprise Manager Information Publisher Reports website.

5.3 Using 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 14, "Using Information Publisher"](#)

Using 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:

Auto-Applying Monitoring Settings to Targets through Administration Groups and Template Collections

Step 1.

The Administrator sets the target property Lifecycle Status to "Production".




Step 2.

Enterprise Manager adds the target to the Administration Group.



Step 3.

Enterprise Manager applies the monitoring template to the target.



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 Manager 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 *global 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
- Platform
- Operating System
- Target Version
- Customer Support Identifier
- Target Type (Allowed but not a global target property.)

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.

Enterprise Manager Administrators and Target Properties

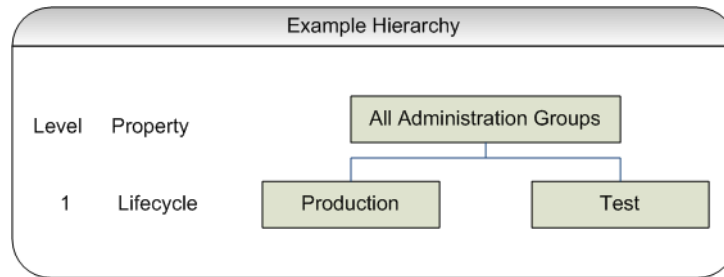
When creating an Enterprise Manager administrator, you can associate properties such as Contact, Location, and Description. However, there are additional resource allocation properties that can be associated with their profile. These properties are:

- Department
- Cost Center
- Line of Business

It is important to note that these properties are persistent--when associated with an administrator, the properties (which mirror, in part, the target properties listed above) are automatically passed to any targets that are discovered or created by the administrator.

Example

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.

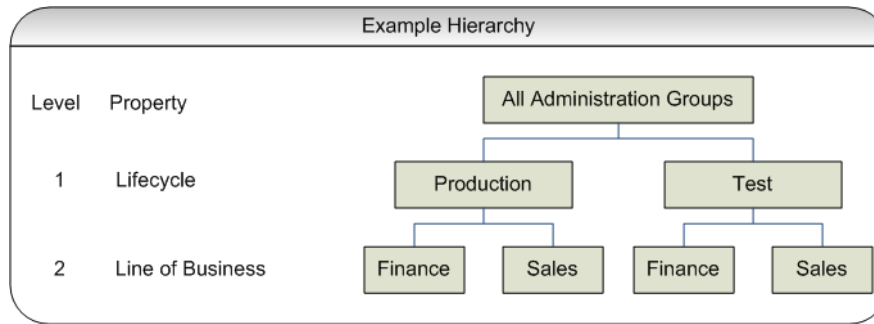
Tree Root Level	Level 1 Target Property Lifecycle Status	Level 2 Target Property Line of Business
All Targets	Production or Mission Critical	Finance
		Sales
	Staging, Test, or Development	Finance
		Sales

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

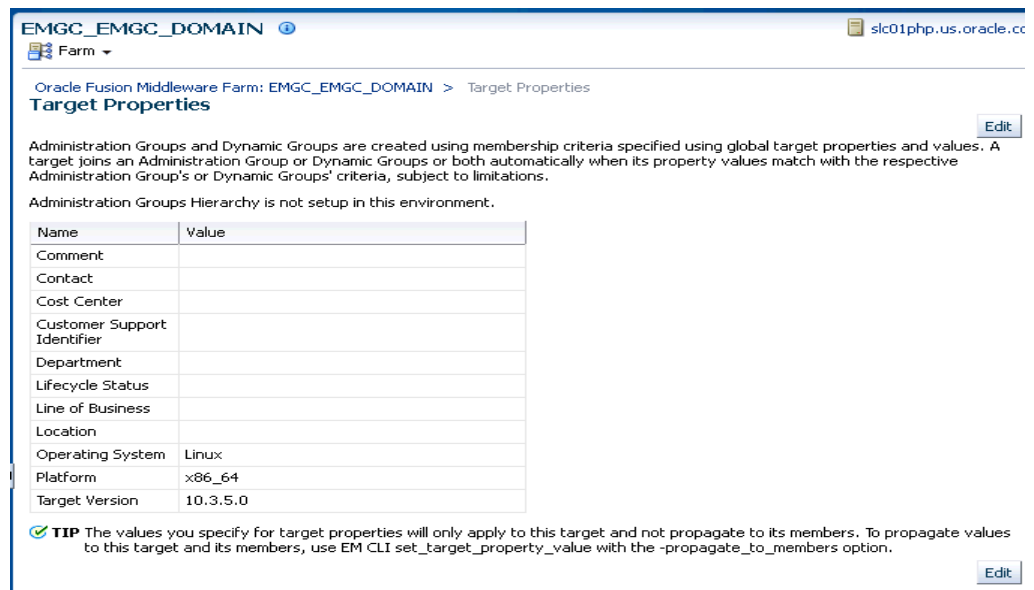
After establishing the desired administration group hierarchy, you must make sure properties are set correctly for each target to ensure they join the correct administration group. Using target properties, Enterprise Manager automatically places targets into the appropriate administration group without user intervention. For targets that have already been added to Enterprise Manager, you can also set the target properties via the console or using the EM CLI verb `set_target_property_value`. See the Enterprise Manager Command Line Interface guide for more information. Note that when running `set_target_property_value`, any prior values of the target property are overwritten. 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.



To help the you specify the appropriate target property values used as administration group criteria, pay attention to the instructional verbiage at the top of the page.

3. Once you have set the target properties, click **OK**.

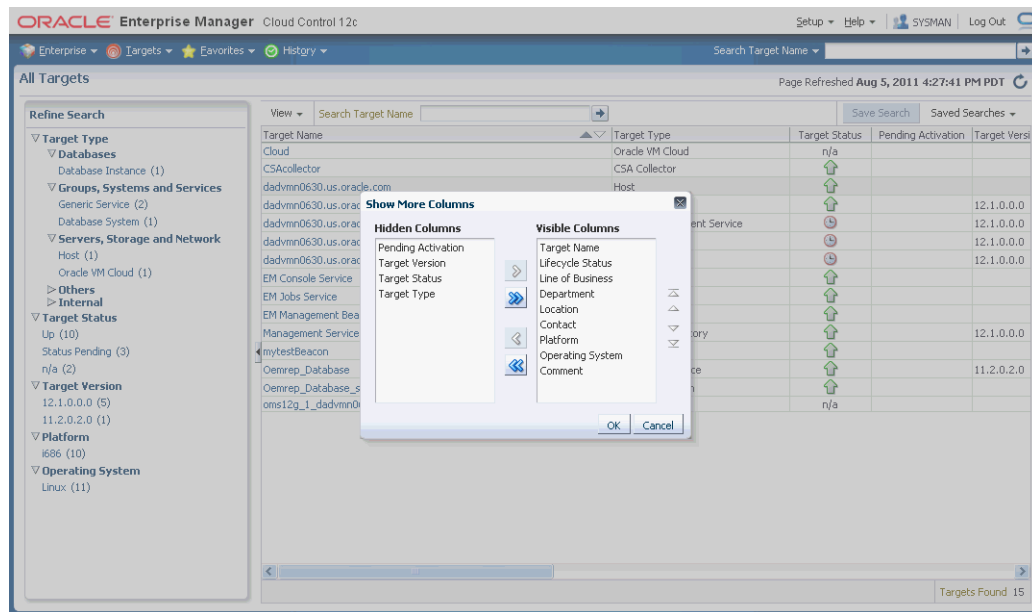
For large numbers of targets, it is best to use the Enterprise Manager Command Line Interface (EM CLI) `set_target_property_value` verb to perform a mass update. For more information about this EM CLI verb, see the Enterprise Manager Command Line Interface guide.

Administration groups, like other aggregate targets, are privilege-propagating: Once you grant an access privilege to an aggregate target, all member targets of that

aggregate target inherit the same access privilege. Any aggregate target that is added to an administration group must also be privilege propagating. A good example of aggregate target is the Privilege Propagating Group. See ["Managing Groups"](#) on page 5-1 for more information.

At any time, you can use the **All Targets** page to view properties across all targets. To view target properties:

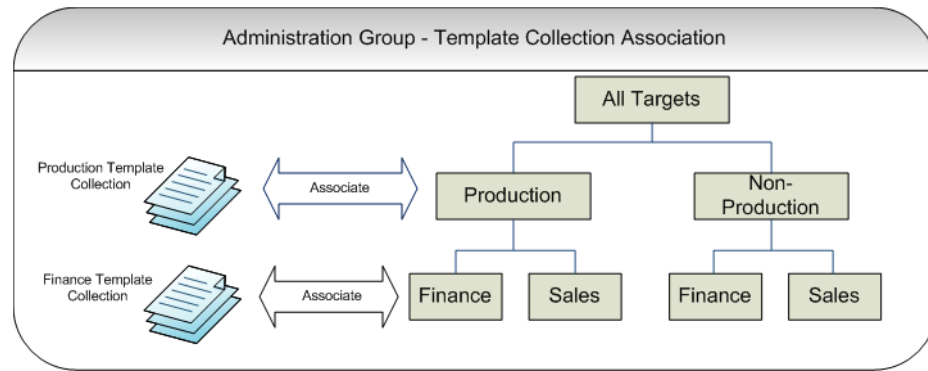
1. From the **Targets** menu, select **All Targets** to display the All Targets page.
2. From the **View** menu, select **Columns**, then select **Show All**.
3. Alternatively, if you are interested in specific target properties, choose **Columns** and then select **Show More Columns** to display column 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.

Auto-Applying Monitoring Settings to Targets through Administration Groups and Template Collections

Step 1.

The Administrator sets the target property Lifecycle Status to "Production".



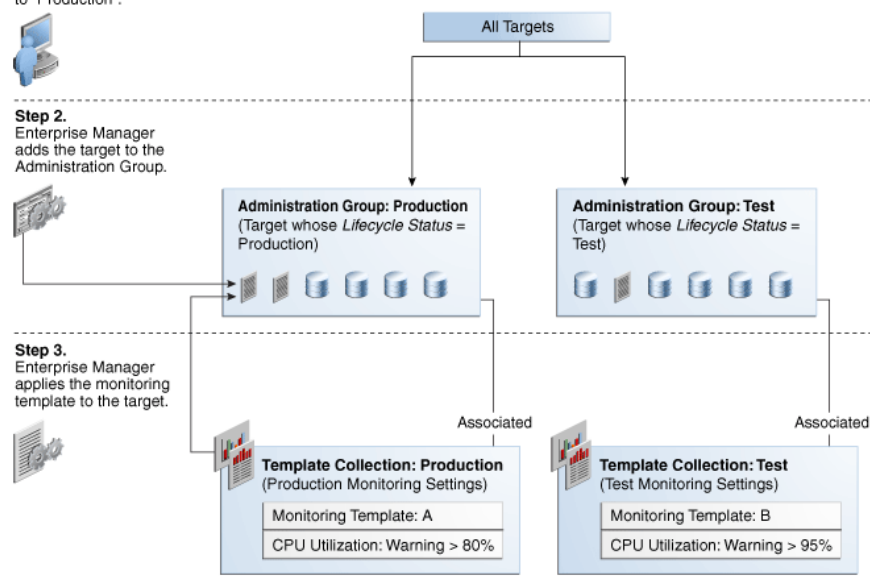
Step 2.

Enterprise Manager adds the target to the Administration Group.



Step 3.

Enterprise Manager applies the monitoring template to the target.



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"](#).

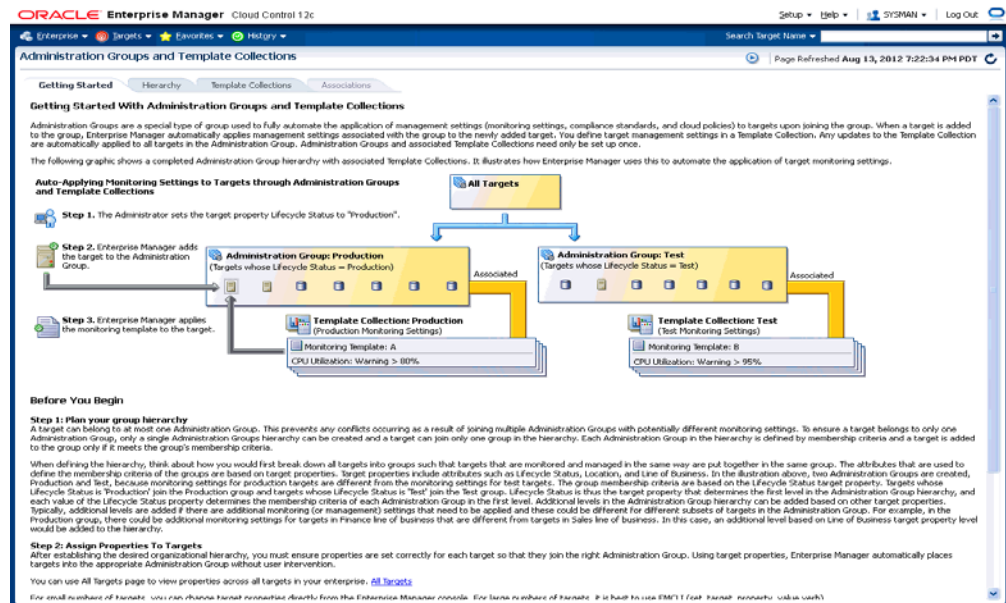
Task 4: Associate the Template Collections with the Administration Groups

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 home page displays.

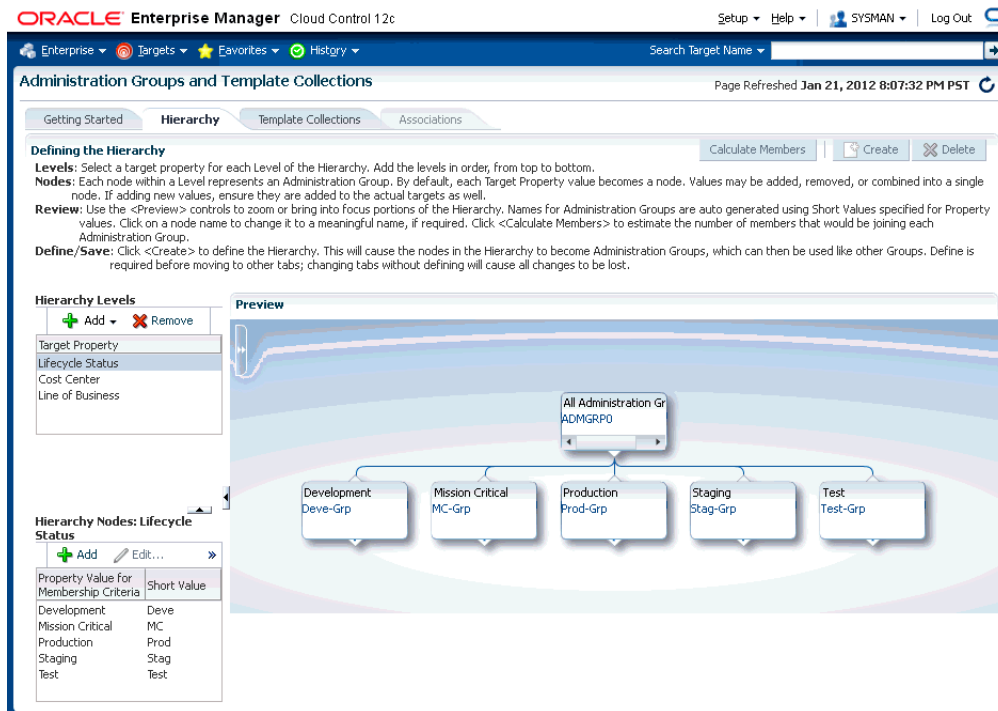


Read the relevant information on the **Getting Started** page. The information contained in this page summarizes the steps outlined in this chapter. For your convenience, links are provided that take you to appropriate administration group functions, as well as the Enterprise Manager **All Targets** 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. The Short Value column displays abbreviated value names that are used to auto-generate group names.

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.



Adding Target Property Values

If you need to display more values for a target property, you can do so by specifying an OMS property `admin_groups_width_limit`. Specify the maximum number of values that should be displayed for a target property. For example, to support showing up to 30 values for a target property that will be used in administration group criteria, set the `admin_groups_width_limit` as follows (using the OMS `emctl` utility):

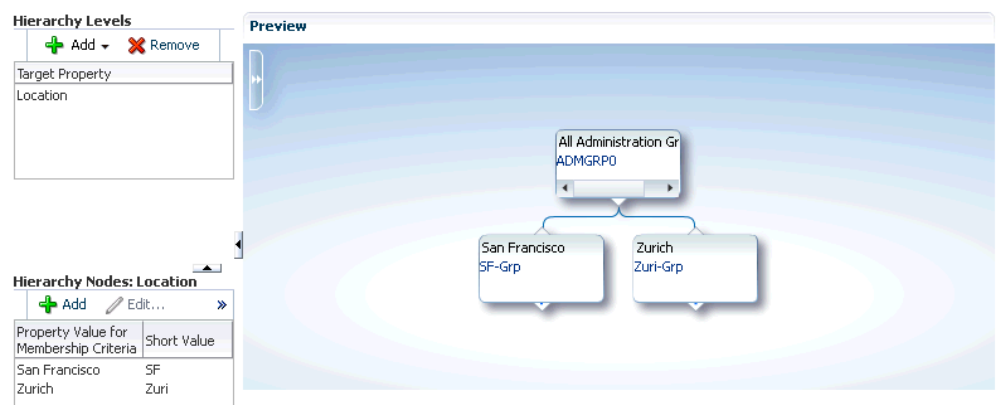
```
emctl set property -name admin_groups_width_limit -value 30 -module emoms
```

You can also add up to four levels after the root node of an administration group hierarchy. If there is a need to add additional level, you will first need to change the OMS `admin_groups_height_limit` property to the maximum height limit. For example, if you want to create to administration group hierarchy consisting of five levels after the root node, set the `admin_groups_height_limit` property as follows (using the OMS `emctl` utility):

```
emctl set property -name admin_groups_height_limit -value 5 -module emoms
```

This is a global property and only needs to be set once using the `emctl` utility of any OMS. This is also a dynamic property and does not require a stop/restart of the OMS in order to take effect.

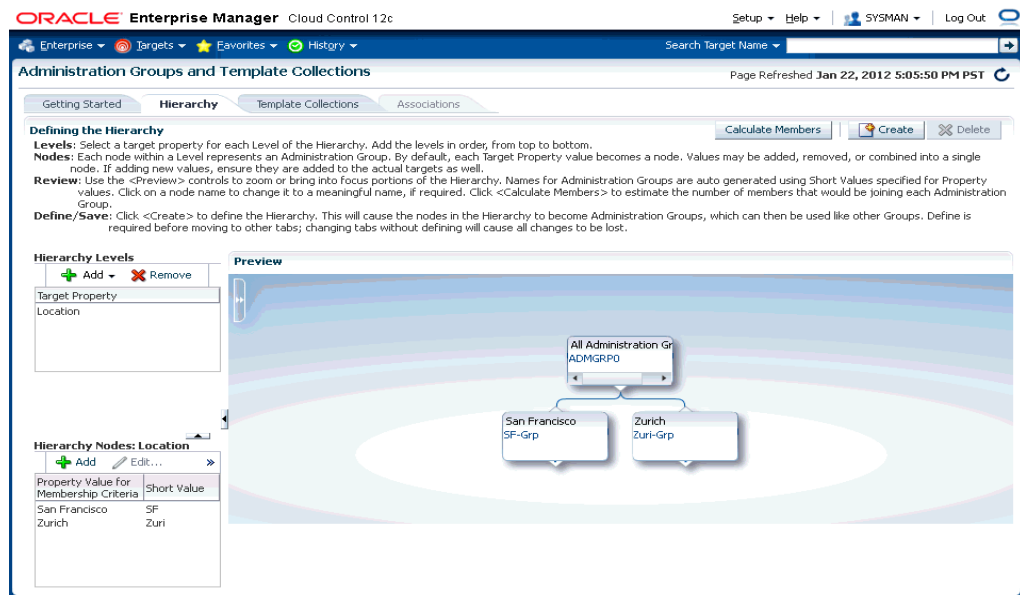
Click **OK**. The two locations "San Francisco" and "Zurich" appear as nodes in the **Preview** pane as shown in the following graphic.



Under certain circumstances, it may be useful to treat multiple property values as one: Targets may have different target property values, but should belong to the same administration group because they have same monitoring profile/settings. 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**.
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. Set the time zone for the group.
 1. Click on the group name. The Administration Group Details dialog displays allowing you to select the appropriate time zone.

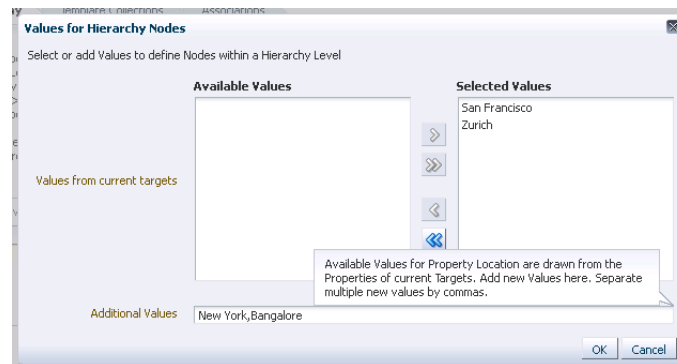


The administration group time zone is used for displaying group charts and also for scheduling operations on the group. Because this is also the default time zone for all subgroups, you should specify the time zone at the highest level group in the administration group hierarchy. Note that child group time zone settings override the default time zone. The auto-generated name can also be changed.

7. Click **Create** to define the hierarchy.

IMPORTANT: Review and define the complete hierarchy before clicking **Create**.

Even after your administration group hierarchy has been created, you can always make future updates if organizational needs change. For example, adding/removing 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. For more information about changing the administration group hierarchy, see ["Changing the Administration Group Hierarchy"](#) on page 6-27.



Click **Update** to save your changes.

6.3.1.3 Defining Template Collections

A template collection is an assemblage of 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.

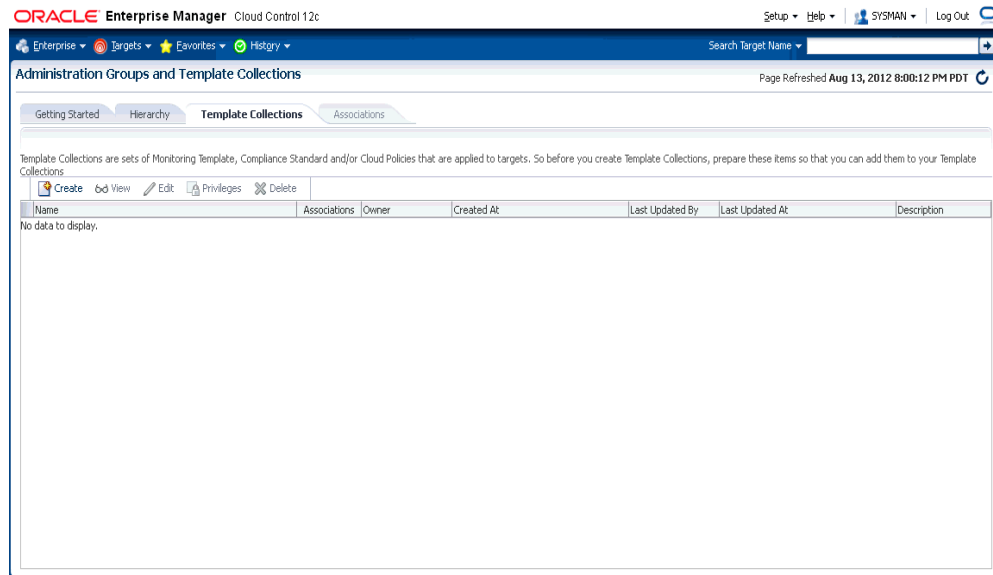
You create template collections when you define administration groups. Template collections may consist of three types of monitoring/management setting categories:

- 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 Policies). 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** 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 **Cancel**. This restores the template collection to its state when it was last saved.

6.3.1.3.1 Required Privileges To create a template collection, you must have the *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
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 Template Collection Operations (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

6.3.1.3.2 Corrective Action Credentials A corrective action is an automated task that is executed in response to a metric alert. When a corrective action is part of a monitoring template/template collection, the credentials required to execute the corrective action will vary depending on how the template is applied.

The two situations below illustrate the different credential requirements.

- *The corrective action is part of a monitoring template that is manually applied to a target.*

When the corrective action runs, it can use one of the following:

- The preferred credentials of the user who is applying the template
- or
- The user-specified named credentials.

The user selects the desired credential option during the template apply operation.

- *The corrective action is part of a monitoring template within a template collection that is associated with an administration group.*

When the corrective action runs, the preferred credentials of the user who is associating the template collection with the administration group is used.

6.3.1.4 Associating Template Collections with Administration Groups

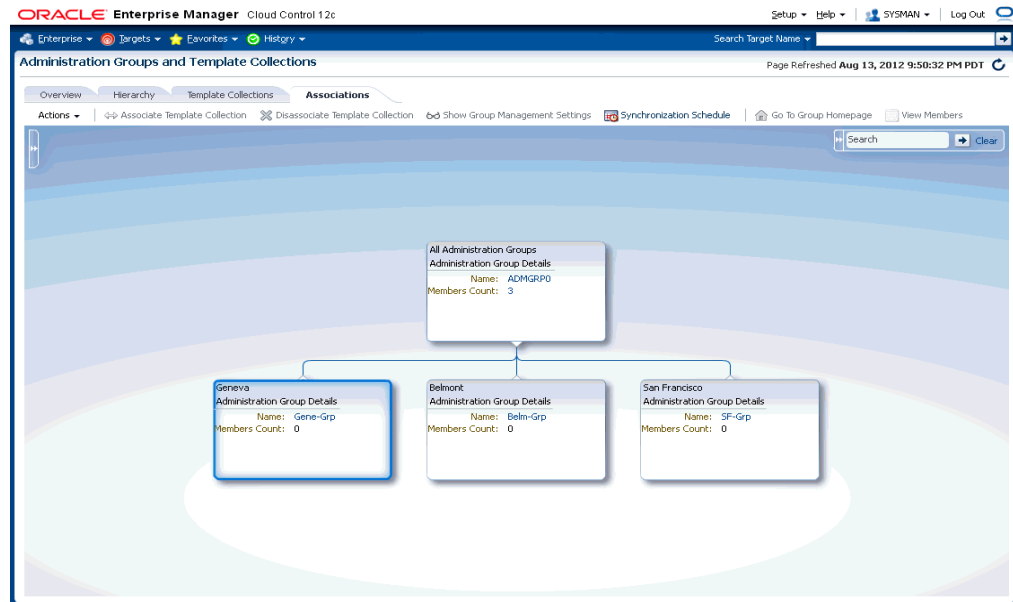
Once you have defined one or more template collections, you need to associate them to administration groups in the hierarchy. You can associate a template collection with one or more administration groups. As a rule, you should associate the template collection with the applicable administration group 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

Note: For users that do not have View privilege on all administration groups, you can also perform the association/disassociation operation from the Groups page (from the **Targets** menu, select **Groups**).

1. Click the **Associations** tab. The 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 list of targets affected by this operation is displayed. Confirm or discard the operation.

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.


Searching for Administration Groups

While the administration group UI is easy to navigate, there may be cases where the administration group hierarchy is inordinately large, thus making it difficult to find individual groups. At the upper right corner of the Associations page is a search function that greatly simplifies finding groups in a large hierarchy.


Figure 6–1 Administration Group Search Dialog

To search for a specific administration group:

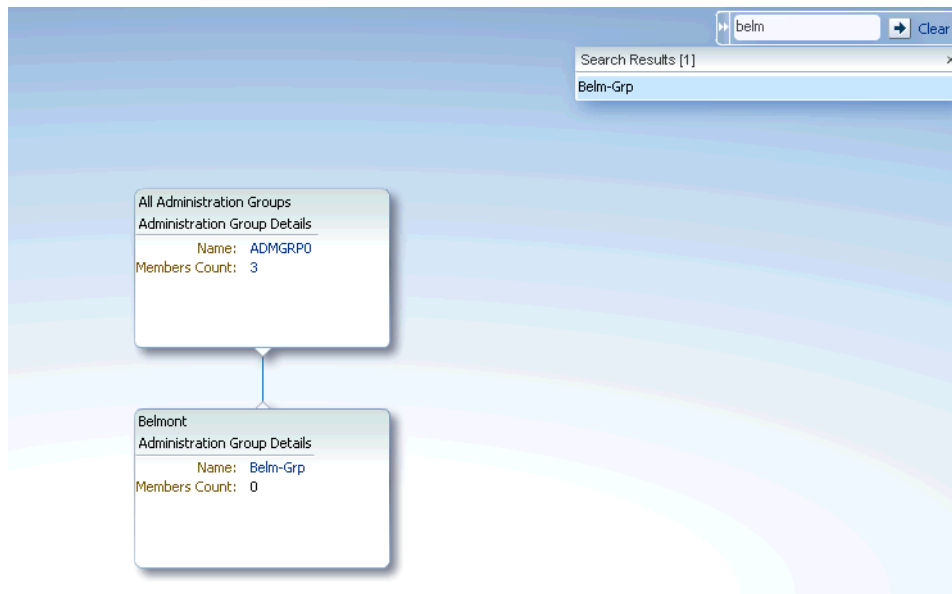
1. If not already displayed, expand the Search interface.



2. Enter either a full or partial group name and click **Search**.



As shown in the graphic, the search results display a list of administration groups that match the search criteria. You can then choose an administration group from the list by double-clicking on the entry. The administration group hierarchy will then display a vertical slice (subset) of the administration group hierarchy from the root node to the group you selected.

Figure 6–2 Administration Group Search: Graphical Display

To restore the full administration group hierarchy, click **Clear**.

Group Names and Searches

In order to perform effective searches for specific administration groups, it is helpful to know how Enterprise Manager constructs an administration group name: Enterprise Manager uses the administration group criteria to generate names. For example, you have an administration group with the following criteria:

- Lifecycle Status: Development or Mission Critical
- Department: DEV
- Line of Business: Finance or HR
- Location: Bangalore

Enterprise Manager assembles a group name based on truncated abbreviations. In this example, the generated administration group name is *DC-DEV-FH-Bang-Grp*

During the design phase of an administration group, you have the option of specifying a custom name. However, if there is large number of groups, it is easier to allow Enterprise Manager to generate unique names.

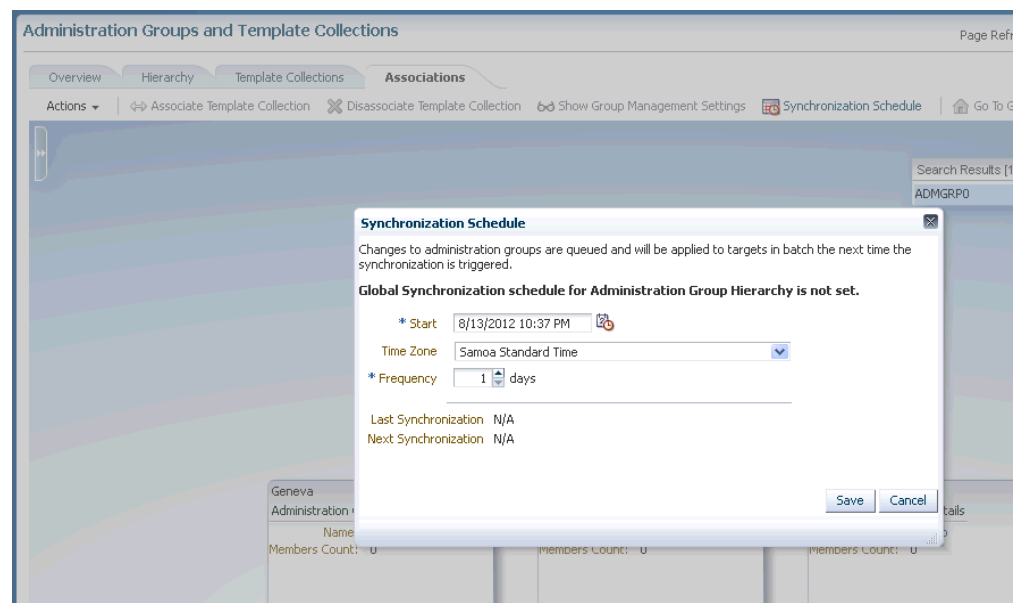
Setting the Global Synchronization Schedule

In order to apply the template collection/administration group association, you must set up a global synchronization 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.

Important: You **must** 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.



- Click **Edit** and then choose a date and time you want the Administration Group-Template Collection 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 this schedule.

- Click **Save**.

Synchronization Sequence of Events

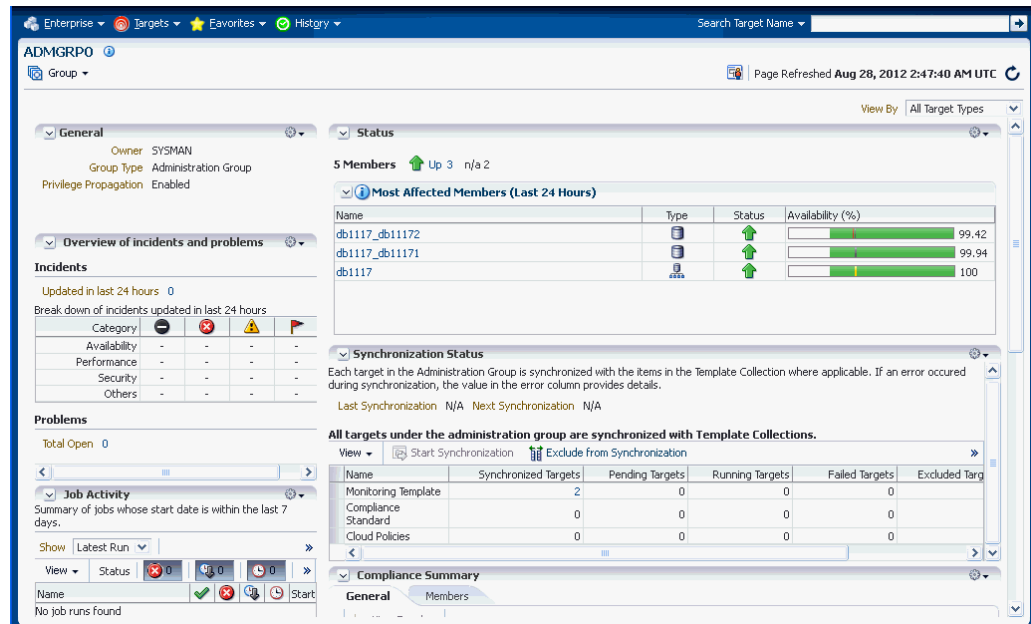
The following table summarizes when Template Collection Synchronization operations (such as apply operations) occur on targets in administration groups.

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.
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. Immediate synchronization operation occurs.
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. Targets under the administration group will be synchronized based on next schedule date in Global Synchronization Schedule.
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.

- Select an administration group in the hierarchy.
- Click **Goto Group Homepage**.
- 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**.

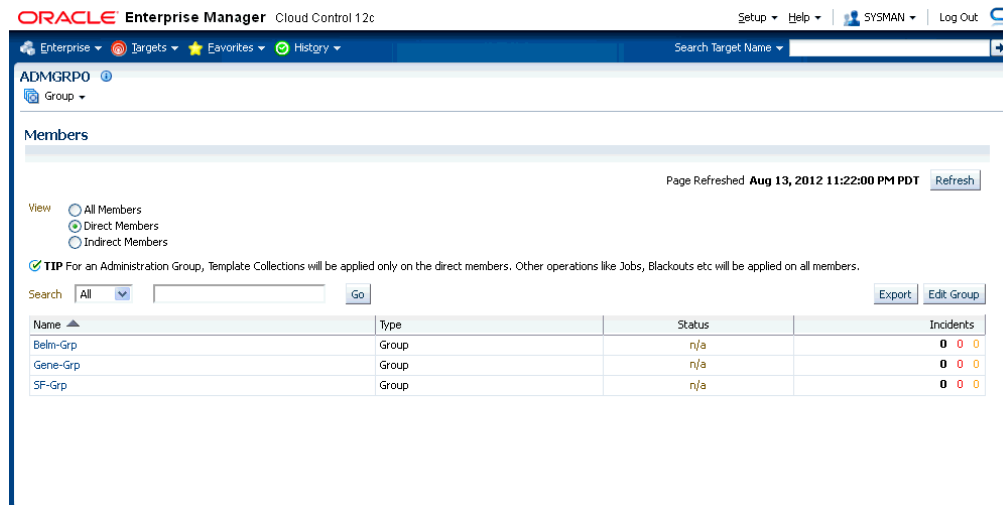
Group Member Type and Synchronization

There are two types of administration group members: Direct and Indirect

- **Direct Members:** Group members whose target properties are set to match the administration group criteria. Monitoring settings, compliance standards, cloud policies from the associated template collection are applied to direct members.
- **Indirect Members:** Aggregate targets (targets composed of member targets) that, although part of the group, will not have management settings from the associated template collection applied. Group functions such as jobs, blackouts or reports will, however, apply to indirect members.

Only direct members are represented in the targets count in the Synchronization Status region. To see which group members are direct or indirect members, from the Enterprise menu, choose menu--> Members page.

1. From the hierarchy diagram, click on a group name to access the group's home page.
2. From the **Group** menu, select **Members**. The Members page displays.



System Targets and Administration Groups

If a system target gets added to an administration group because it matches group criteria, then the system target and its constituent members are also added. However, for template apply purposes, it will only operate on the direct members that also match the administration group criteria. All other group operations, such as jobs and blackouts, will apply on all members, both direct and indirect.

Disassociating a Template Collection from a Group

To remove a template collection from an administration group.

1. From the **Setup** menu, select **Add Target** and then **Administration Groups**. The Administration Group home page displays.
2. Click on the **Associations** tab to view the administration group hierarchy diagram.
3. From the hierarchy diagram, select the administration group with the template collection you wish to remove. If necessary, use the **Search** option to locate the administration group.
4. Click **Disassociate Template Collection**. The number of targets affected by this operation is displayed. Click **Continue** or **Cancel**.

The template collection is immediately removed. See "[Synchronization Sequence of Events](#)" on page 6-22 for more information.

Viewing Aggregate (Group Management) Settings

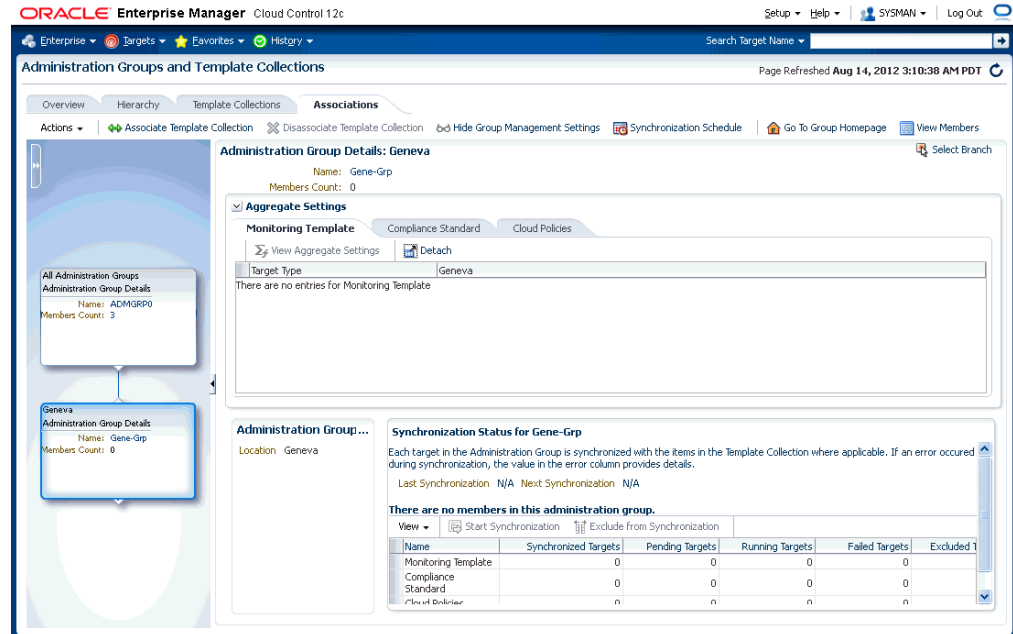
For any administration group, you can easily view what template collection components (monitoring templates, compliance standards, and/or cloud policies) are associated with individual group members.

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 **Setup** menu, select **Add Target** and then **Administration Groups**. The Administration Group home page displays.

2. Click on the **Associations** tab to view the administration group hierarchy diagram.
3. From the hierarchy diagram, select the desired administration group.
4. Click **Show Group Management Settings**.

The **Administration Group Details** page displays.



This page displays all aggregate settings for monitoring templates, compliance standards and cloud policies that will be applied to members of the selected administration group (listed by target type). The page also displays the synchronization status of group members.

To change the display to show a different branch of the administration group hierarchy, click **Select Branch**. This function lets you display hierarchy branches by choosing different target property values.

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, details of the Associated Template Collection for the group selected in hierarchy viewer, job activity, or critical patch advisories. To view administration group home pages:

1. From the hierarchy diagram, select an administration group.
2. Click **Goto Group Homepage**. The homepage for that particular administration group displays.

General
 Owner: SYSMAN
 Group Type: Administration Group
 Privilege Propagation: Enabled

Status
 5 Members Up 3 n/a 2

Most Affected Members (Last 24 Hours)

Name	Type	Status	Availability (%)
db1117_db11172	Database Instance	Up	99.43
db1117_db11171	Database Instance	Up	99.94
db1117	Database Instance	Up	100

Synchronization Status
 Each target in the Administration Group is synchronized with the items in the Template Collection where applicable. If an error occurred during synchronization, the value in the error column provides details.
 Last Synchronization: N/A Next Synchronization: N/A

All targets under the administration group are synchronized with Template Collections.

Name	Synchronized Targets	Pending Targets	Running Targets	Failed Targets	Excluded Targets
Monitoring Template	2	0	0	0	0
Compliance Standard	0	0	0	0	0
Cloud Policies	0	0	0	0	0

Compliance Summary
 General Members

Alternatively, from the Enterprise Manager **Targets** menu, choose **Groups**. From the table, you can expand the group hierarchy.

Groups
 Page Refreshed Jan 24, 2012 2:48:48 AM UTC

Groups allow users to monitor and manage many targets as one. Users creating Privilege Propagating Groups must have full privilege on all member targets. When privileges (e.g. View) on Privilege Propagating Group are granted by the owner to any administrator, the grantee gets the same privilege on all the member targets as well. It is possible to create regular groups that are not Privilege Propagating. Administration Groups are hierarchical in nature and their membership is only through criteria defined using global target properties. All Administration Groups are Privilege Propagating.

Search
 Name:
 Search Advanced Search Save Search Criteria

View Create Create Like Edit Remove View Members Customize Page Associate Template Collection

Name	Group Type	Template Collection	Members	Member Status Summary	Incidents
adc6140250_SBAI	Database Instance	n/a	Siebel Component(8)	1 7 - - 1 1 - -	- - - -
adc6140250_SBAI	Database Instance	n/a	Siebel Component(3)	- 3 - - - - - -	- - - -
ADMGRP0	Administration Group	n/a	Group(9), Database Instance(4), Host(2), Cluster Database(1)	1 6 - - 1 5 - -	- - - -
Stag-Grp	Administration Group	-	Group(2)	- - - - - - - -	- - - -
Stag-Test-G	Administration Group	-	-	- - - - - - - -	- - - -
Stag-prod-G	Administration Group	-	-	- - - - - - - -	- - - -
Prod-Grp	Administration Group	TC02	Database Instance(4), Group(2), Host(1), Cluster Database(1)	1 5 - - 1 4 - -	- - - -
Prod-prod-G	Administration Group	-	-	- - - - - - - -	- - - -
Prod-Test-G	Administration Group	-	Database Instance(4), Host(1), Cluster Database(1)	1 5 - - 1 4 - -	- - - -
Deve-Grp	Administration Group	-	Group(2), Host(1)	- 1 - - - - - -	- 1 - -
Group1	Administration Group	n/a	Host(25), Group(1)	- 17 - - - 23 10 - -	- - - -
Group2	Administration Group	n/a	Host(1), Database Instance(1), Listener(1), Oracle WebLogic Server(1)	- 4 - - - 1 - -	- - - -

Columns Hidden: 9

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.

Enterprise Targets Favorites History Search Target Name

Unassigned Targets Page Refreshed Jan 24, 2012 2:54:26 AM UTC

This table lists all the targets that are not part of any Administration Group. The values for those properties that are used to form Administrations Groups Hierarchy are shown. Non privilege propagating aggregate targets cannot become member of Administration Groups.

Search Target Name Target Type Advanced Supply Chain Planning Go

View Detach

Target Name	Target Type	Non Privilege Propagating Aggregate	Department	Lifecycle Status
slc00awo.us.oracle.com:4889_Management_Service	Oracle Management Service	<input checked="" type="checkbox"/>		
slc00awo.us.oracle.com:4889_Management_Service_PBS	OMS Platform			
EM Management Beacon	Beacon			
OH1457874305_slc00azb	Oracle Home			
OH556752328_slc00azb	Oracle Home			
agent12g1_slc00awo	Oracle Home			
EM Jobs Service	EM Service	<input checked="" type="checkbox"/>		
EMGC_ADMINSERVER	Oracle WebLogic Server			
mds-owsm	Metadata Repository			
mds-sysman_mds	Metadata Repository			
EMGC_OMS1	Oracle WebLogic Server			
emgc	Application Deployment	<input checked="" type="checkbox"/>		
empbs	Application Deployment	<input checked="" type="checkbox"/>		
OCMRpeater	Application Deployment	<input checked="" type="checkbox"/>		
oracle.security.apm(11.1.1.3.0)	Oracle Authorization Policy Manager			
ohs1	Oracle HTTP Server			
EMGC_GCDomain	Oracle Fusion Middleware Farm	<input checked="" type="checkbox"/>		
WebLogicServer10.3.5.0_slc00awo	Oracle Home			
slc00azb.us.oracle.com:4889_Management_Service	Oracle Management Service	<input checked="" type="checkbox"/>		

Columns Hidden 5

- 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.

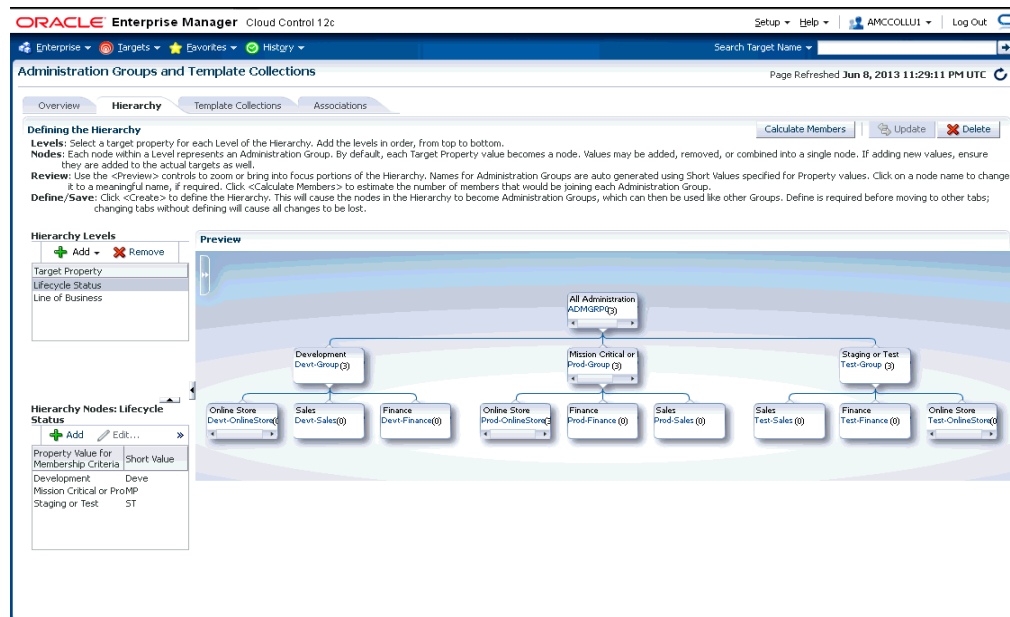
For information on changing target properties, see ["Planning"](#) on page 6-3.

- Click your browser *back* button to return to the **Administration Groups and Template Collections** homepage.

6.4 Changing the Administration Group Hierarchy

Organizations are rarely static--new lines of business may be added or perhaps groups are reorganized due to organizational expansion. To accommodate these changes, you may need to make changes to the existing administration group hierarchy.

Beginning with Enterprise Manager 12c Release 12.1.0.3, you can change the administration group hierarchy without having to rebuild the entire hierarchy. You can easily perform administration group alterations such as adding more groups to each hierarchy level, merging two or more groups, or adding/deleting entire hierarchy levels. All of these operations can be performed from the Hierarchy page.



Important: After making any change to the administration group hierarchy, click **Update** to save your changes.

6.4.1 Adding a New Hierarchy Level

Adding a new hierarchy level equates to adding a new target property to the administration group criteria. For this reason, you must set the value of this target property for all your targets in order for them to continue to be part of the administration group hierarchy.

See "[Adding a Hierarchy Level](#)" on page 6-12 for step-by-step instructions on adding a new level.

6.4.2 Removing a Hierarchy Level

Removing a hierarchy level equates to deleting a target property, which in turn causes groups at that level to be deleted.

To remove a hierarchy level:

1. On the **Administration Group** page, click the **Hierarchy** tab.
2. From the **Hierarchy Levels** table, select a hierarchy level and click **Remove**
3. Click **Update** to save your changes.

If any of those groups have an associated template collection, then the monitoring settings of the subgroups of the deleted group will be impacted since the subgroups obtained monitoring settings from the associated template collection. You may need to review the remaining template collections and re-associate the template collection with the appropriate administration group.


6.4.3 Merging Administration Groups

If you want to merge two or more administration groups, you merge their corresponding target property criteria in the administration group hierarchy definition. The group merge operation consists of retaining one of the groups to be merged and then moving over the targets from the other groups into the group that is retained. Once the targets have been moved, the other groups will be deleted.

You choose which group is retained by choosing its corresponding target property value. The group(s) containing the selected target property value as part of its criteria is retained. If the retained target property criteria corresponds to multiple groups, i.e. group containing subgroups, the movement of targets will actually occur at the lowest level administration groups since the targets only reside in the lowest level administration groups. The upper-level administration groups' criteria will be updated to include the criteria of the other groups that have been merged into it.

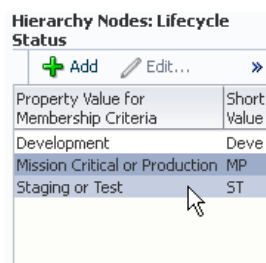
To merge groups:

1. Select a target property from the list of chosen properties in the **Hierarchy Levels** table. You choose the target property corresponding to the groups you want to merge.




For example, let us assume you want to merge <Dev-Group> with <Test or Stage Group>. In the hierarchy, this corresponds to target property Lifecycle Status. The associated property values are displayed.

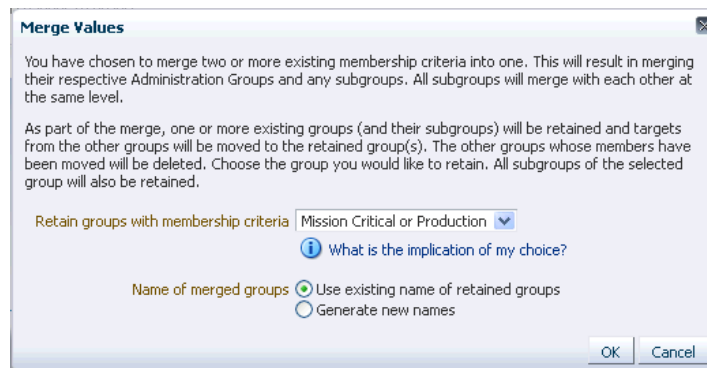
2. Select two or more property values corresponding to the groups you would like to merge by holding down the *Shift* or *CTRL* key and clicking on the desired values.



3. Click Merge.



The Merge Values dialog displays.



Again, by merging membership criteria (target properties), you are merging administration groups and their respective subgroups. You choose the administration group to be retained. The other groups will be merged into that group.

4. Choose the group you wish to retain and specify whether you want to use the existing name of the retained group or specify a new name.

Important: When deciding which group to retain, consider choosing the group that is used in most group operations such as incident rule sets, system dashboard, or roles. These groups will be retained and the members of the other merged groups will join the retained groups. After the merge, group operations on the retained groups will also now apply to the members from the other merged groups. Doing so minimizes the impact of the merge.

5. Click **OK** to merge the groups.
6. Click **Update** to save the new hierarchy.

Example

Your administration group consists of the following:

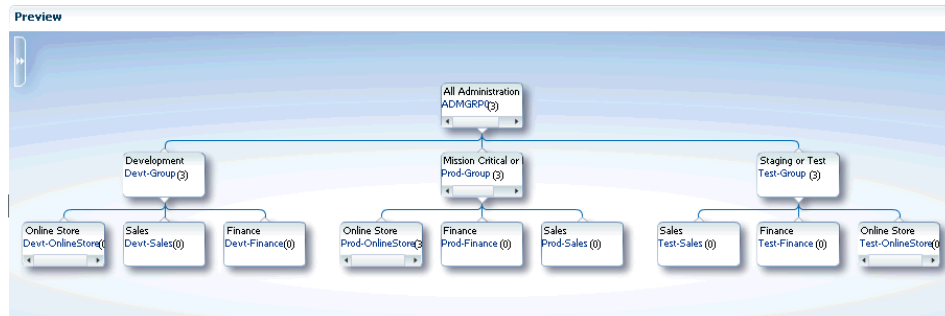
Hierarchy Levels

- Lifecycle Status
- Line of Business

Hierarchy Nodes

- *Lifecycle Status*
 - Development
 - Mission Critical or Production
 - Staging or Test
- *Line of Business*
 - Online Store
 - Sales
 - Finance

The following graphic shows the administration group hierarchy.



You decide that you want to merge the *Mission Critical or Production* group with the *Staging or Test* group.

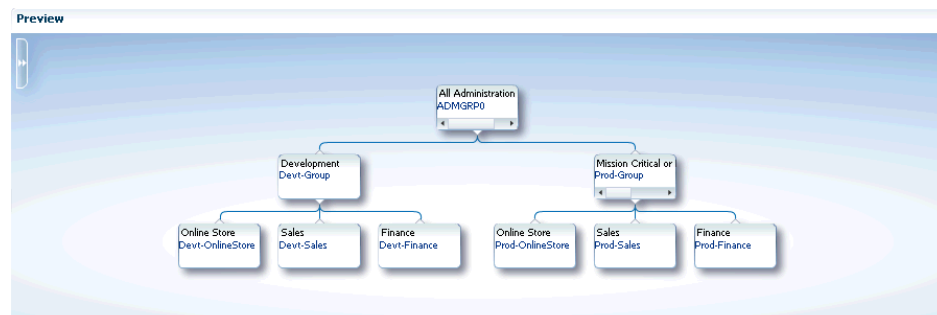
Choose Lifecycle Status from the **Hierarchy Levels** table.

From the **Hierarchy Nodes** table, choose both *Mission Critical or Production* and *Staging or Test*.

Select **Merge** from the **Hierarchy Nodes** menu.

The **Merge Values** dialog displays. In this case, you want to keep original name (*Mission Critical or Production*) of the retained group.

After clicking **OK** to complete the merge, the resulting administration group hierarchy is displayed. All targets from *Test-Sales* group moved to the *Prod-Sales* group. The *Test-Sales* group was deleted. All targets from the *Test-Finance* group moved to the *Prod-Finance* group. The *Test-Finance* group got deleted.



Click **Update** to save the changes.

6.4.4 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** 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** 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.
6. Click **Update**.

Using Monitoring Templates

Monitoring templates simplify the task of setting up monitoring for large numbers of targets by allowing you to specify the monitoring and Metric and Collection Settings once and applying them to many groups of targets as often as needed.

This chapter covers the following topics:

- [About Monitoring Templates](#)
- [Definition of a Monitoring Template](#)
- [Default Templates \(Auto Apply Templates\)](#)
- [Viewing a List of Monitoring Templates](#)
- [Creating a Monitoring Template](#)
- [Editing a Monitoring Template](#)
- [Applying Monitoring Templates to Targets](#)
- [Comparing Monitoring Templates with Targets](#)
- [Comparing Metric Settings Using Information Publisher](#)

7.1 About Monitoring Templates

Monitoring templates let you standardize monitoring settings across your enterprise by allowing you to specify the monitoring settings once and apply them to your monitored targets. You can save, edit, and apply these templates across one or more targets or groups. A monitoring template is specified for a particular target type and can only be applied to targets of the same type. 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 metric extensions), thresholds, metric collection schedules, and corrective actions.

Once a monitoring template is defined, it can be applied to your targets. This can be done either manually through the Enterprise Manager console, via the command line interface (EM CLI), or automatically using template collections. See "[Defining Template Collections](#)" on page 6-15 for more information. For any target, you can preserve custom monitoring settings by specifying metric settings that can never be overwritten by a template.

Oracle Certified Templates and Oracle Provided Templates

In addition to templates that you create, there are also Oracle-supplied templates.

- *Oracle-provided templates:* Contains the out-of-box monitoring settings for the target type. So, if you updated a target's monitoring settings, and later wanted to revert to the original out-of-box settings, then you can use and apply the Oracle-provided templates.
- *Oracle Certified templates:* Contains a specific set of metrics for a specific purpose and the purpose is indicated in the description associated with the template.

Example: The template called *Oracle Certified - Enable AQ Metrics for SI Database* contains AQ-related metrics for single instance databases. You would use this template (or copy the metric settings into your own template) if you want to use the AQ metrics.

7.2 Definition of a Monitoring Template

A monitoring template defines all Enterprise Manager parameters you would normally set to monitor a target. A template specifies:

- **Name:** A unique identifier for the template. The template name must be globally unique across all templates defined within Enterprise Manager.
- **Description:** Optional text describing the purpose of the template.
- **Target Type:** Target type to which the template applies.
- **Owner:** Enterprise Manager administrator who created the template.
- **Metrics:** Metrics for the target type. A monitoring template allows you to specify a subset of all metrics for a target type. With these metrics, you can specify thresholds, collection schedules and corrective actions.
- **Other Collected Items:** Additional collected information (non-metric) about your environment.

7.3 Default Templates (Auto Apply Templates)

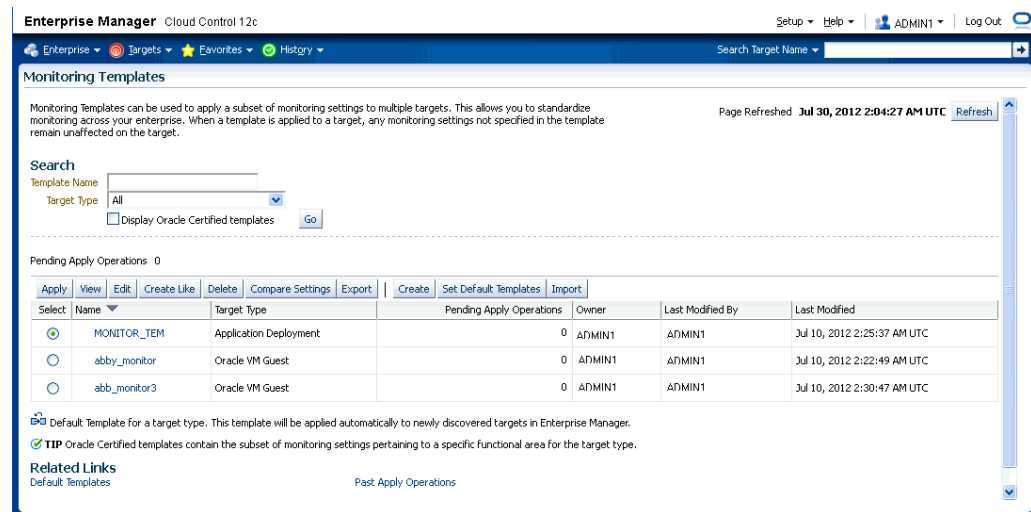
Under certain circumstances, Oracle's out-of-box monitoring settings may not be appropriate for targets in your monitored environment. Incompatible Metric and Collection Settings for specific target types can result in unwanted/unintended alert notifications. Enterprise Manager allows you to set default monitoring templates that are automatically applied to newly added targets, thus allowing you to apply monitoring settings that are appropriate for your monitored environment.

Note: Super Administrator privileges are required to define default monitoring templates.

7.4 Viewing a List of Monitoring Templates

To view a list of all Monitoring Templates, from the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.

The Monitoring Templates page displays all the out-of-box templates and the templates for which you have at least VIEW privilege on. Enterprise Manager Super Administrators can view all templates.

Figure 7–1 Monitoring Templates

You can begin the monitoring template creation process from this page.

7.5 Creating a Monitoring Template

Monitoring template allow you to define and save monitoring settings for specific target types. To define a new template:


1. From the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.
2. Click **Create**. Enterprise Manager gives you the option of selecting either a specific target or a target type. Template monitoring settings are populated according to the selected target or target type. Click **Continue**.

Note: If the selected target type is either Web Application or Service, you will only be able to select those targets for which you have Operator privilege.

3. Enter requisite template information on the General, Metric Thresholds, and Other Collected Items tabs.

On the Metric Thresholds tab, you can delete or add monitoring template metrics. To delete existing metrics, select one or more metrics and click **Remove Metrics from Template**.

To add metrics, click **Add Metrics to Template**. The Add Metrics to Template page displays as shown in the following graphic.



On this page, you can select a source from which you can copy metrics to the template. Sources include specific targets, other monitoring templates, or metric extensions. Click **Continue** once you have finished modifying the template metrics.

4. Once you have finished entering requisite information, click **OK**.

7.6 Editing a Monitoring Template

The Monitoring Templates page lists all viewable templates. To edit a template, you must have FULL access privileges.

To edit a Monitoring Template:

1. From the **Enterprise** menu, select **Monitoring**, and then **Monitoring Templates**.
2. Choose the desired template from the table.
3. Click **Edit**.
4. Once you have finished making changes, click **OK**.

Sharing Access with Other Users

By default, template owners (creators) have FULL access privileges on the template and Enterprise Manager Super Administrators have FULL access privileges on all templates. Only the template owner can change access to the template. You, as owner, can grant VIEW (view the template) or FULL (edit or delete the template) on the template to a user or role.

7.7 Applying Monitoring Templates to Targets

As mentioned earlier, a monitoring template can be applied to one or more targets of the same target type, or to composite targets such as groups. For composite targets, the template is applied to all member targets that are of the appropriate type. If you applied the template manually or via EM CLI, once a template is applied, future changes made to the template will not be automatically propagated to the targets: You must reapply the template to all affected targets

Administration Groups and Template Collections: Applying Monitoring Templates Automatically

Monitoring templates can be automatically applied whenever a new targets are added to your Enterprise Manager environment. Automation is carried out through Administration Groups and Template Collections Administration Groups are a special type of group used to automate application of monitoring settings to targets upon joining the group. When a target is added to the administration group Enterprise

Manager applies monitoring settings from the associated template collection consisting of monitoring templates, compliance standards, and cloud policies. If changes are later made to the monitoring template, Enterprise Manager automatically applies the changes to the relevant targets based on the synchronization schedule. For more information, see ["Using Administration Groups"](#) on page 6-1.

7.7.1 Applying a Monitoring Template

To apply a template, you must have at least *Manage Target Metrics* target privileges on the destination target(s).

1. From the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.
2. Select the desired template from the table.
3. Click **Apply**.
4. Select the desired apply options and the target(s) to which you want the templates applied. See [Section 7.7.2, "Monitoring Template Application Options"](#) for additional information.
5. Click **OK**.

7.7.2 Monitoring Template Application Options

You can choose aggregate targets such groups, systems or clusters as destination targets. The templates will apply to the appropriate members of the group/system/cluster as they currently exist. If new members are later added to the group, you will need to re-apply the template to those new members. Template application is performed in the background as asynchronous jobs, so after the apply operation is performed, you can click on the link under the Pending Apply Operations column in the main templates table to see any apply operations that still are pending.

When applying a Monitoring Template, metric settings such as thresholds, comparison operators, and corrective actions are copied to the destination target. In addition, metric collection schedules including collection frequency and upload interval are also copied to the target. You determine how Enterprise Manager applies the metric settings from the template to the target by choosing an apply option.

7.7.2.1 Apply Options

Template apply options control how template metric and policy settings are applied to a target. Two template apply options are available:

- **Template will completely replace all metric settings in the target:** When the template is applied, all metrics and policies defined in the template will be applied to the target. Pre-existing target monitoring settings not defined in the template will be disabled: Metric thresholds will be set to NULL or blank. Policies will be disabled. This effectively eliminates alerts from these metrics and policies by clearing current severities and violations.
- **Template will only override metrics that are common to both template and target:** When the template is applied, only metrics and policies common to both the template and target are updated. Existing target metric and policies that do not exist in the template will remain unaffected. When this option is selected, additional template apply options are made available for metrics with key value settings.

7.7.2.2 Metrics with Key Value Settings

A metric with key value settings is one that can monitor multiple objects at different thresholds. For example, the Filesystem Space Available(%) metric can monitor different mount points using different warning and critical thresholds for each mount point.

When the template contains a metric that has key value settings, you can choose one of three options when applying this template to a target. As an example, consider the case where the template has the following metric:

Filesystem Space Available(%)

Mount Point	Operator	Warning Threshold	Critical Threshold
/	£	40	20
/private	£	30	20
/u1	£	30	20

And a host target has the same metric at different settings:

Mount Point	Operator	Warning Threshold	Critical Threshold
/	£	30	10
/private	£	25	15
/private2	£	20	20
All Others	£	25	15

These are the results for each option:

1) All key value settings in the template will be applied to the target, any additional key values settings on the target will not be removed

When the template is applied to the target using this copy option, all the template settings for the mount points, /, /private, and /U1 will be applied. Existing target settings for mount points not covered by the template remain unaffected. Thus, the resulting settings on the target for this metric will be:

Mount Point	Operator	Warning Threshold	Critical Threshold
/	£	40	20
/private	£	30	20
/u1	£	30	20

2) All key value settings in the template will be applied to target, any additional key value settings on the target will be removed.

When the template is applied to the target using this copy option, all template settings will be applied to the target. Any object-specific threshold settings that exist only on the target will be removed, any object-specific thresholds that are only in the template will be added to the target. Thus, the final settings on the target will be:

Mount Point	Operator	Warning Threshold	Critical Threshold
/	£	40	20

Mount Point	Operator	Warning Threshold	Critical Threshold
/private	£	30	20
/u1	£	30	20
All Others	£	25	15

3) Only settings for key values common to both template and target will be applied to the target

When the template is applied to the target using this copy option, only the settings for the common mount points, / and /private will be applied. Thus, the resulting settings on the target for this metric will be:

Mount Point	Operator	Warning Threshold	Critical Threshold
/	£	40	20
/private	£	30	20
/private2	£	20	10
All Others	£	25	15

7.8 Comparing Monitoring Templates with Targets

The intended effect of applying Monitoring Templates to destination targets is not always clear. Deciding how and when to apply a template is simplified by using the Compare Monitoring Template feature of Enterprise Manager. This allows you to see at a glance how metric and collection settings defined in the template differ from those defined on the destination target. You can easily determine whether your targets are still compliant with the monitoring settings you have applied in the past. This template comparison capability is especially useful when used with aggregate targets such as groups and systems. For example, you can quickly compare the metric and collection settings of group members with those of a template, and then apply the template as appropriate.

Performing a Monitoring Template-Target comparison:

1. From the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.
2. Choose the desired template from the table.
3. Click **Compare Settings**. The Compare Monitoring Template page displays.
4. Click **Add** to add one or more destination targets. The Search and Select dialog displays.
5. Select a one or more destination targets and then click **Select**. The selected targets are added to the list of destination targets.
6. Select the newly added destination targets and then click **Continue**. A confirmation message displays indicating the *Compare Template Settings* job was successfully submitted.
7. Click **OK** to view the job results. Note: Depending on the complexity of the job run, it may take time for the job to complete.

7.8.1 When is a metric between a template and a target considered "different"?

The metric is said to be different when any or all the following conditions are true (provided the target does not have "Template Override" set for that metric):

- The Warning Threshold settings are different
- The Critical Threshold settings are different.
- The Collection Schedules are different.
- The Upload Intervals are different.
- The number of occurrences (for which the metric has to remain at a value above the threshold before an alert is raised) are different.)
- For user-defined metrics, in addition to the above, the OS Command/SQL statement used to evaluate the user-defined metric is different. Note that this applies only if the user-defined metric name and the return type are the same.
- The metric extension marked for delete will be shown as "different" on the destination target and the template only if:
 - A metric extension with the same name exists on both the destination target and template.
 - The return type (String, Numeric) of the metric extension is the same on both the destination target and template.
 - The metric type is the same on both the destination target and the template.

7.9 Comparing Metric Settings Using Information Publisher

In addition to viewing metric differences between Monitoring Templates and destination targets using the Compare Monitoring Template user-interface, you can also use Information Publisher to generate reports containing the target-template differences. Using Information Publisher's reporting capabilities gives you more flexibility for displaying and distributing metric comparison data. For more information, see ["Using Information Publisher"](#) on page 14-1.

Create a Report Definition

1. From the **Enterprise** menu, select **Reports** and then **Information Publisher Reports**.
2. Click **Create**. The Create Report Definition user interface is displayed.
3. On the General page, specify the report name, how targets should be included, target privileges, report time period, and display options.
4. On the Elements page, click **Add** to access the Add Element page.
5. Select the Monitoring Template Comparison element and click **Continue** to return to the Element page.
6. Once you have added the report element, click the **Set Parameter** icon to specify requisite operational parameters. On this page, you specify a report header, select a monitoring template, destination targets, and template application settings for multiple threshold metrics. Click **Continue** to return to the Elements page.
7. Click **Layout** to specify how information should be arranged in the report.
8. Click **Preview** to validate that you are satisfied with the data and presentation of the report.

9. On the Schedule page, define when reports should be generated, and whether copies should be saved and/or sent via e-mail, and how stored copies should be purged.
10. On the Access page, click **Add** to specify which Enterprise Manager administrators and/or roles will be permitted to view this generated report. Additionally, if you have GRANT_ANY_REPORT_VIEWER system privilege, you can make this report definition accessible to non-credential users via the Enterprise Manager Reports Website
11. Click **OK** when you are finished.
12. Validate the report definition. If the parameters provided conflict, validation errors or warnings will appear and let you know what needs attention.
13. Once the report definition has been saved successfully, it appears in the Report Definition list under the Category and Subcategory you specified on the General page.

Viewing the Report

1. Find the template comparison report definition in the Report Definition list. You can use the Search function to find or filter the list of report definitions.
2. Click on the report definition title. If the report has a specified target, the report will be generated immediately. If the report does not have a specified target, you will be prompted to select a target.

Scheduling Reports for Automatic Generation

1. Create or edit a report definition.
2. On the Schedule page, choose the **Schedule Report** option.
3. Specify a schedule type. The schedule parameters on this page change according to the selected schedule type.

When reports are scheduled for automatic generation, you have the option of saving copies to the Management Repository and/or sending an e-mail version of the report to designated recipients.

If a report has been scheduled to save copies, a copy of the report is saved each time a scheduled report completes. When a user views a report with saved copies by clicking on the report title, the most recently saved copy of the report is rendered. To see the complete list of saved copies click on the Saved Copies link at the top of the report. Enterprise Manager administrators can generate a copy of the report on-demand by clicking on the Refresh icon on the report.

7.10 Exporting and Importing Monitoring Templates

For portability, monitoring templates can be exported to an XML file and then imported into another Enterprise Manager installation as an active template.

Important: You can export templates from Enterprise Manager 10g release 2 or higher and import them into Enterprise Manager 12c. .

Exporting a Monitoring Template

To export a template to an XML file:

1. From the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.

2. Select the desired monitoring template from the table.
3. Click **Export**.

Note: If you are running an Enterprise Manager 11g or earlier release, use the EM CLI *export_template* verb to perform the export operation. Note that if the monitoring template contains *policy rules* from earlier Enterprise Manager releases (pre-12c), these will not be imported into Enterprise Manager 12c as policy rules no longer exist in this release.

Importing a Monitoring Template

To import a template from an XML file:

1. From the **Enterprise** menu, select **Monitoring** and then **Monitoring Templates**.
2. Click **Import**. The Import Template page displays.
3. Specify the monitoring template XML file you want to import.
4. Click **Import**.

7.11 Changing the Monitoring Template Apply History Retention Period

You can view monitoring template apply history using the predefined report in Information Publisher. From the **Enterprise** menu, select **Reports** and then **Information Publisher**. On the Information Publisher page, you can enter "template" in the **Title** text entry field and click **Go**. The predefined report *Monitoring Template Apply History (last 7 days)* appears in the report list.

By default, Enterprise Manager retains the monitoring template apply history for a period of 180 days. If required, you can change the retention period to a value suitable for your monitoring needs. Although the retention period cannot be indefinite, it can be set to an extremely long period of time. Enterprise Manager provides the following PL/SQL API to change the retention period.

```
mgmt_template_ui.modify_purge_policy(p_retention_days=><num_days>)
```

This procedure takes a NUMBER as input (*num_days*).

Using Metric Extensions

Metric extensions provide 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:

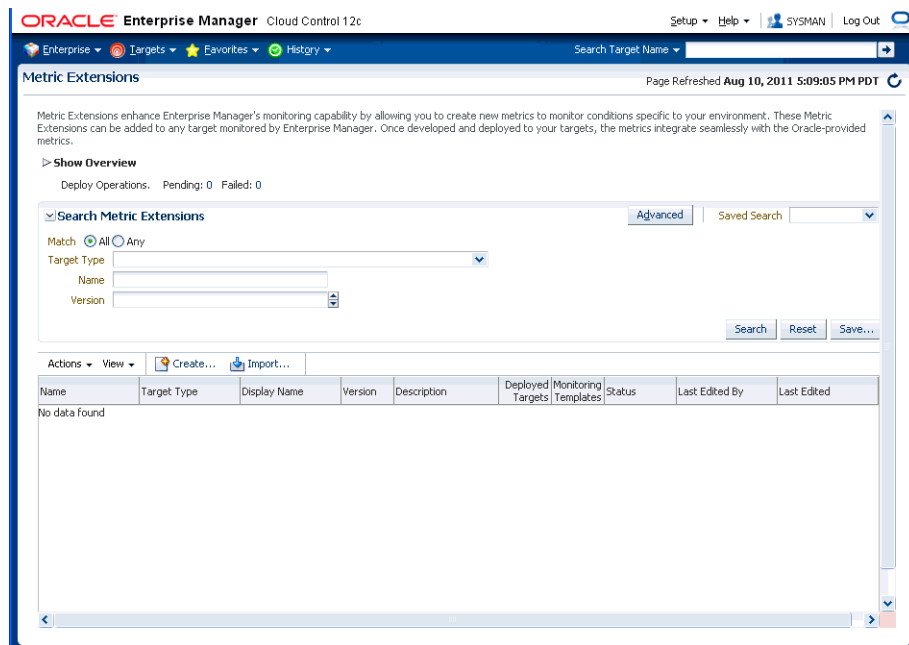
- [What are Metric Extensions?](#)
- [Metric Extension Lifecycle](#)
- [Working with Metric Extensions](#)
- [Adapters](#)
- [Converting User-defined Metrics to Metric Extensions](#)
- [Metric Extension Command Line Verbs](#)

8.1 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:

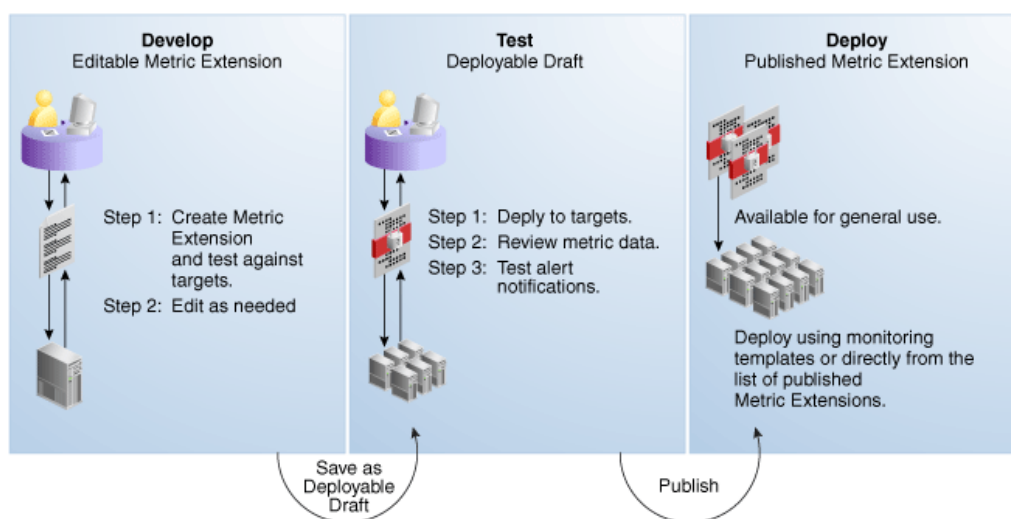
- **Limited Integration:** If the OS or SQL user-defined metric executed custom scripts, or required atonal dependent files, the user needed to manually transfer these files to the target's file system.
- **Limited Application of Query Protocols:** OS user-defined metrics cannot model child objects of 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:** User-defined metrics only allow OS user-defined metrics against host targets and SQL user-defined metrics against database targets. No other target types are permitted. If, for example, you want to deploy a user-defined metric against WebLogic instances in your environment, you will not be able 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 extensions 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.

8.2 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.

ORACLE Enterprise Manager Cloud Control 12c Help

Metric Extensions

General Properties Adapter Columns Credentials Test Review

Create New : General Properties Back Step 1 of 6 Next Finish Cancel

Specify the basic properties for the metric extension.
The default collection can be overridden on a target instance basis in the Metric and Policies Settings page.

General Properties

* Target Type: Host

* Name:

A Metric Extension Name can only contain alpha-numeric characters, _, -, and . (non leading)

* Display Name:

* Adapter: OS Command - Multiple Columns

Tokenizes OS command output using user-specified delimiter

Description:

Collection Schedule

Data Collection ☐ Disabled ☒ Enabled

Collection Frequency: By Minutes

Repeat Every: 15 Minutes

Use of Metric Data: ☐ Alerting Only ☒ Alerting and Historical Trending

Upload Interval: 1 Collections

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.

ORACLE Enterprise Manager Cloud Control 12c Help

Metric Extensions

General Properties Adapter Columns Credentials Test Review

Edit command (ME\$ME_Host) v1 : Test Back Step 5 of 6 Next Finish Cancel

You can perform real-time metric evaluations here on specified test targets.
It is recommended that you test your metric extension here first before deploying to targets. Targets that you select need to be up in order for your test to succeed.

Test Targets

Target Name	Target Type	Hostname	Current Status	Agent
dadvmn0630.us.oracle.com	Host	dadvmn0630.us.oracle.com		https://dadvmn0630.us.oracle.com:11852/emd/main/

Test Results

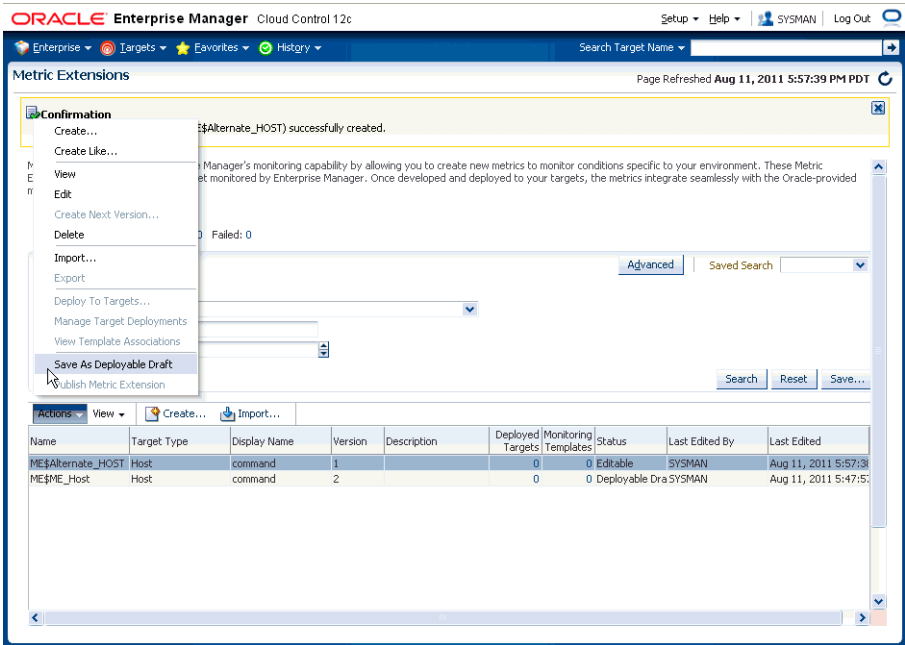
Target Name	Error Message
-------------	---------------

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

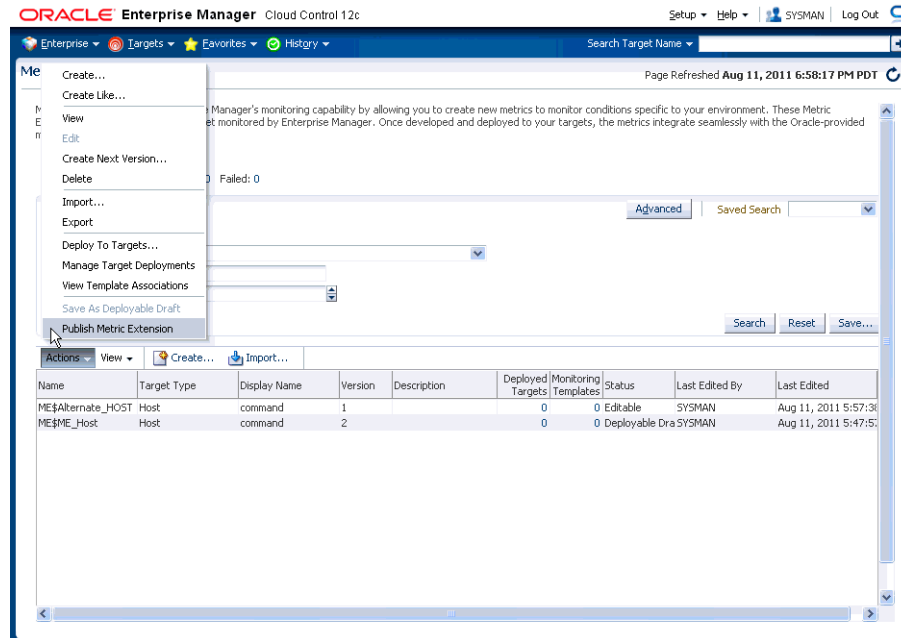
Once your metric extension returns the expected data during real-time target testing, you are ready to test its robustness and actual behavior in Enterprise Manager by deploying it against targets and start collecting data. At this point, the metric extension is still private (only the developer can 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.

8.3 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 Manager 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.

8.3.1 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. This privilege must be granted on a per-metric extension basis.
 - **Full Metric Extension:** Allows users with 'Create Metric Extension' privilege to edit and create new versions 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.

8.3.2 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.

8.3.3 Managing Administrator Privileges

Before an Enterprise Manager administrator can edit or delete a metric extension created by another administrator, that administrator must have been granted requisite access privileges. *Edit* privilege allows editing and creating next versions of the extension. *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 **Resource Privileges** and click Manage Privilege Grants (pencil icon) for the Metric Extensions 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.

On this page, you can also grant an administrator the *Create Metric Extension* privilege, which will allow them to manage metric extension access. See ["Managing Administrator Access to Metric Extensions"](#) for more information.

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.

8.3.4 Managing Administrator Access to Metric Extensions

Administrators commonly share the responsibility of monitoring and managing targets within the IT environment. Consequently, creating and maintaining metric extensions becomes a collaborative effort involving multiple administrators. Metric extension owners can control access directly from the metric extension UI.

8.3.4.1 Granting Full/Edit Privileges on a Metric Extension

As metric extension owner or Super Administrator, perform the following actions to assign full/edit privileges on a metric extension to another administrator:

1. From the **Enterprise** menu, select **Monitoring**, then select **Metric Extensions**.
2. Choose a metric extension requiring update.
3. From the **Actions** menu, select **Manage Access**.
4. Click **Add**. The administrator selection dialog box appears. You can filter the list by administrator, role, or both.
5. Choose one or more administrators/roles from the list.
6. Click **Select**. The chosen administrators/roles appear in the access list.

In the **Privilege** column, **Edit** is set by default. Choose **Full** from the drop-down menu to assign **Full** privileges on the metric extension.

Edit Privilege: Allows an administrator to make changes to the metric extension but not delete it.

Full Privilege: Allows an administrator to edit and also delete the metric extension. The privilege granted to a user or role applies to all versions of the metric extension.

7. Click **OK**.

8.3.4.2 Revoking Access Privileges on a Metric Extension

As metric extension owner or Super Administrator, perform the following actions to revoke metric extension privileges assigned to another administrator:

1. From the **Enterprise** menu, select **Monitoring**, then select **Metric Extensions**.
2. Choose a metric extension requiring update.
3. From the **Actions** menu, select **Manage Access**.
4. Choose one or more administrators/roles from the list.
5. Click **Remove**. The chosen administrators/roles is deleted from the access list.
6. Click **OK**.

Enterprise Manager allows metric extension ownership to be transferred from the current owner of the metric extension to another administrator as long as that administrator has been granted the *Create Metric Extension* privilege.

Note: The Enterprise Manager Super Administrator has full managerial access to all metric extensions (view, edit, and ownership transfer).

As mentioned above, *manage access* is only enabled for the owner of the extension or an Enterprise Manager Super User. Once the ownership is transferred, the previous

owner does not have any management privileges on the metric extension unless explicitly granted before ownership transfer. The **Change Owner** option is only available to users and not roles.

Manage access allows the metric extension owner or Super Administrator to grant other Enterprise Manager users or roles the ability to edit, modify, or delete metric extensions.

8.3.4.3 Transferring Metric Extension Ownership

Enterprise Manager allows metric extension ownership to be transferred from the current owner of the metric extension to another administrator as long as that administrator has been granted the *Create Metric Extension* privilege.

Note: The Enterprise Manager Super Administrator has full managerial access to all metric extensions (view, edit, and ownership transfer).

As mentioned above, *manage access* is only enabled for the owner of the extension or an Enterprise Manager Super User. Once the ownership is transferred, the previous owner does not have any management privileges on the metric extension unless explicitly granted before ownership transfer. The **Change Owner** option is only available to users and not roles.

Manage access allows the metric extension owner or Super Administrator to grant other Enterprise Manager users or roles the ability to edit, modify, or delete metric extensions.

8.3.5 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 multi-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.

Collection Schedule

You defined the frequency with which metric data is collected and how it is used (Alerting Only or Alerting and Historical Trending) by specifying collection schedule properties.

Depending on the target type selected, an *Advanced* option region may appear. This region may (depending on the selected target type) contain one or two options that determine whether metric data continues to be collected under certain target availability/alert conditions. The options are:

- **Option 1:** Continue metric data collection even if the target is down. This option is visible for all target types except for *Host* target types as it is not possible to collect metric data when the host is down.
- **Option 2:** Continue metric data collection when an alert severity is raised for a specific target metric. This metric is defined in such a way (AltSkipCondition element is defined on this metric) that when a severity is generated on this metric, the metric collections for other target metrics are stopped. The explanatory text above the checkbox for this option varies depending on the selected target type.

The Management Agent has logic to skip evaluation of metrics for targets that are known to be down to reduce generation of metric errors due to connection failures. If the AltSkipCondition element is defined for that target metric, other metrics are skipped whenever there is an error in evaluating the Response metric or there is a non-clear severity on the Response:Status metric. There are two situations where a metric collection will be skipped or not happen:

- When a target is down (option 1). This is same as the Severity on Response/Status metric.
- When a target is UP, but there is a severity on any other metric. Such conditions are called Alt Skip (Alternate Skip) conditions.

Option 2 is only visible if an AltSkipCondition defined for one of the target's metrics. For example, this option will not be visible if the selected target type is *Oracle Weblogic Domain*, but will be visible if the selected target type is *Database Instance*.

The following graphic shows the Advanced collection schedule options.

The screenshot shows a configuration panel for 'Collection Schedule'. At the top, 'Data Collection' is set to 'Enabled' (radio button). Below it, 'Use of Metric Data' is set to 'Alerting and Historical Trending' (radio button). The 'Upload Interval' is set to '1' 'Collections'. A 'Frequency' section shows 'Collection Frequency' set to 'By Minutes' (dropdown) and 'Repeat Every' set to '15' 'Minutes'. Below this is an 'Advanced' section with two checkboxes: 'Select to collect this Metric Extension on above condition.' (unchecked) and 'Select to collect this Metric Extension on above condition.' (checked). The text for the checked option reads: 'By default, this metric will not be collected when target has a severity on Response:State metric. Override default behavior and collect metric.'

5. From the Columns page, add metric columns defining the data returned from the adapter. Note that the column order should match the order with which the adapter returns the data.

- **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. You can also create rate and delta metric columns based on an existing data column. See *Rate and Delta Metric Columns* below.

- **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 `processor1, core1`, and a threshold of 60% for `processor2, core2`, you would specify:
`processor1,core1=50;processor2,core2=60`

- **Manually Clearable Alert**

Note: You must expand the Advanced region in order to view the Manually Clearable Alert option.

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 an 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.

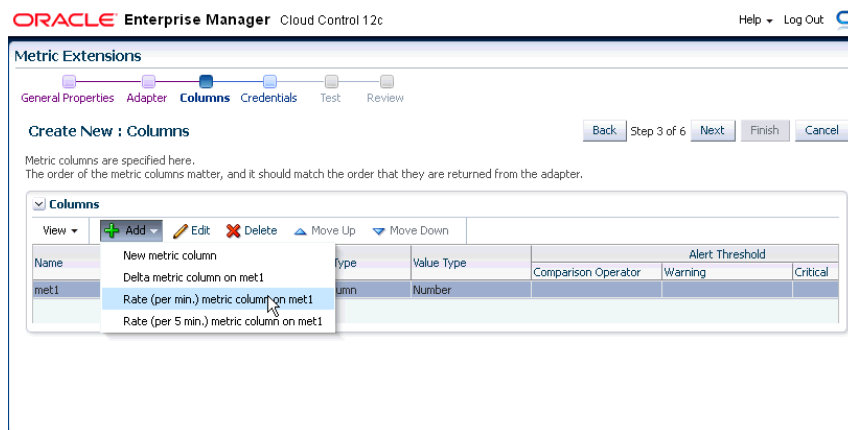
- **Metric Category**

The metric category this column belongs to.

Rate and Delta Metric Columns

You can create additional metric columns based on an existing data column that measures the rate at which data changes or the difference in value (delta) since the last metric collection. The rate/delta metric definition will be allowed when a metric's collection frequency is periodic. For example, collected every 10 minutes. Conversely, a metric that is computed every Monday and Tuesday only cannot have a rate/delta metric as data sampling is too infrequent.

After at least one data column has been created, three additional options appear in the **Add** menu as shown in the following graphic.



- Add Delta metric columns based on another metric column

Example: You want to know the difference in the table space used since the last collection.

Delta Calculation:

current metric value - previous metric value

- Add Rate Per Minute metric column based on another metric column

Example: You want to know the average table space usage per minute based on the table space column metric which is collected every 1 hr.

Rate Per Minute Calculation:

(current metric value - previous metric value) / collection schedule

where the *collection schedule* is in minutes.

- Add Rate Per Five Minutes metric column based on another metric column

Example: You want to know the average table space usage every five minutes based on the table space column which is collected say every 1 hour]

Rate Per Five Minute Calculation:

*[(current metric value - previous metric value) / collection schedule] * 5*

where the *collection schedule* is in minutes.

To create a rate/delta metric column, click on an existing data column in the table and then select one of the rate/delta column options from the **Add** menu.

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 dbnmp 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**.

8.3.6 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**.

8.3.7 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**.

8.3.8 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**.

8.3.9 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**.

8.3.10 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 Production, Deployable Draft and Published metric extension versions.

6. Confirm the export file is downloaded.

8.3.11 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.

8.3.12 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.

8.3.13 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.

8.4 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](#)

8.4.1 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 its 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.

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'
```

8.4.2 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=out_x
em_result=out_y
```

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.

8.4.3 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 multi-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 8–1 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
out_x
out_y
```

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 its 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.

8.4.4 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 specified.

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

Overriding default monitoring credentials by creating and using a custom monitoring credential set for database target.

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'
```

8.4.5 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.

8.4.6 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.

8.5 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.

8.5.1 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 task within that session. The status of the session and tasks is viewable throughout the workflow.

8.5.2 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 udmnig_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 udmnig_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 udm mig_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 udm mig_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 udm mig_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 udm mig_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 udm mig_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 udmrig_list_matches to list available matches

```

The session details command shows that the migration process is indeed complete.

8.6 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 udmnig_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 udmnig_summary

```

```
emcli udmnig_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 Extensions

Options:

```
-session_id=<id of the session>
    Specify the id that was returned at time of session created,
    or from the output of udmig_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 udmig_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 udmig_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 udmig_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 udmig_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 udmig_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 udmig_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 udmig_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>

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](#)

9.1 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 9-1](#)) 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 9–1 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.

9.1.1 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.

9.1.2 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.

9.2 Preliminary Considerations

Before proceeding to the procedural information presented in [Section 9.3, "Creating Jobs"](#) on page 9-4, it is advisable to read the topics presented in the sections below:

- [Creating Scripts](#)
- [Sharing Job Responsibilities](#)
- [Submitting Jobs for Groups](#)

9.2.1 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.

9.2.2 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.

9.2.3 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 Groups"](#)

9.3 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](#)

9.3.1 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.

9.3.2 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 9.3.2.1, "Specifying a Single Operation"](#) or [Section 9.3.2.2, "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 9.3.2.3, "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, "Using 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.

9.3.2.1 Specifying a Single Operation

Note: The following information applies to step 2 in [Task 3, "Specify Parameters"](#) on page 9-6.

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
```

- The job status depends on the exit code returned by the command. If the command execution returns 0, the job returns a status of Succeeded. If it returns any other value, it returns a job status of Failed.

9.3.2.2 Specifying a Script

Note: The following information applies to step 2 in [Task 3, "Specify Parameters"](#) on page 9-6.

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.

The job status depends on the exit code returned by the last command in the script. If the last command execution returns 0, the job returns a status of Succeeded. If it returns any other value, it returns a job status of Failed. You should implement proper exception handling in the script and return non-zero exit codes when appropriate. This will avoid situations in which the script failed, but the job reports the status as Succeeded.

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
```

9.3.2.3 Access Level Rules

Note: The following rules apply to [Task 6, "Specify Who Can Access the Job - \(optional\)"](#) on page 9-8.

- 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.

9.3.3 Creating a SQL Script Job

The basic process for creating a SQL script job is the same as described in [Section 9.3.2, "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](#)

9.3.3.1 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.

9.3.3.2 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.

9.3.3.3 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.

9.3.3.4 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;
```

9.3.4 Creating a Multi-task Job

The basic process for creating a multi-task job is the same as described in [Section 9.3.2, "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](#)

9.3.4.1 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.

9.3.4.2 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.

9.3.4.3 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:

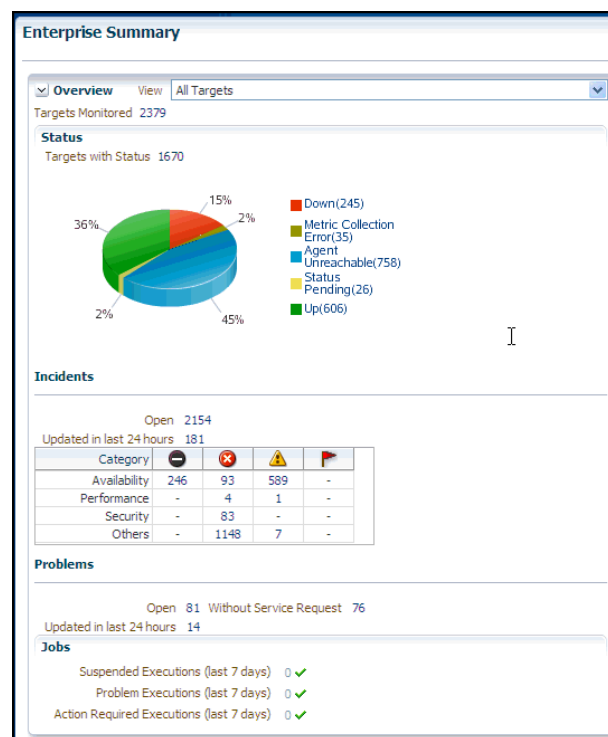
- **Always** — Task is executed each time the job is run.
- **On Success** — Task execution **Depends On** the successful execution of another task.
- **On Failure** — Task execution **Depends On** the execution failure of another task.

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.

9.4 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 9-2 shows the Jobs section at the bottom of the Enterprise Summary page.

Figure 9-2 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.

9.5 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 9-3) 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.

To access this page, from the Setup menu, select **Incidents**, then **Job Events**.

Figure 9-3 Job Event Generation Criteria Page

Job Event Generation Criteria Page Refreshed Jun 14, 2013 3:28:1

In Enterprise Manager, all changes to the status of a job are treevents. You can either create incidents based on these events, and view and manage them in the Incident Manager, or notifications that can act as triggers for an workflow in a Change Management System. To avoid overloading them, it is advisable to enable events only on those targets and only for job status changes that are critical to your data. Events are enabled for job status Action Required, Problems but not for targets.

Step 1: Events For Job Status And Severity

Match status and severity: ☒ Either ☐ Both

Enable Events for Job Status:

- ☐ All
- ☐ Scheduled
- ☐ Running
- ☒ Action Required
- ☐ Suspended
- ☐ Succeeded
- ☒ Problems

Enable Events for Status Severity:

- ☐ All
- ☐ Critical
- ☐ Informational

Severity of Events for Job Status

Job Status	Severity	Description
> Scheduled		The event is generated with Informational severity.
> Running		The event is generated with Informational severity.
> Action Required		The event is generated with Critical severity.
> Suspended		The event is generated with Critical or Informational severity.
> Succeeded		The event is generated with Informational severity.
> Problems		The event is generated with Critical or Informational severity.

Overview

- Super Administrator: Select the Job your enterprise that is allowed to generate events. Optionally, add targets that generate these job events.
- Any Administrator: Add targets that generate events for the Job Status by the Super Admin.

Step 2: Events For Jobs Without Target(s)

Enable Events for jobs without target(s): ☐ Yes ☒ No

Tip Only users with Super Administrator privileges can modify settings in step 1 and step 2.

Step 3: Events For Targets

Add individual targets to enable events.

+ Add... - Remove

Name	Type
No Targets are currently selected.	

9.5.1 Enabling Events For Job Status, Status Severity, 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 "Enable Events for Job Status" region, select the statuses for which you want to publish events.
 - b. In the "Enable Events for Status Severity" region, select whether you want to enable events for a critical status, informational status, for for both.
 - c. In the "Enable Events for Jobs Without Target(s)" section, select **Yes** if you want to create events for jobs that are not associated with any target.
 - d. In the "Events for Targets" section, click **Add** to add targets for which you want the job events to be enabled.

6. Click **Apply**.

9.5.2 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**.

9.6 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](#)

9.6.1 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 appears.
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**.

Incident Rules - All Enterprise Rules

Create Rule Set [Save] [Cancel]

A rule set is a collection of rules that applies to a common set of objects, for example, targets, jobs, and templates. A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets.

* Name: Sample Job Rule Set

Description: [Text Area]

Enabled: ☒ Enabled

Owner: SYSMAN [How is this used?]

Type: ☒ Enterprise ☐ Private

Applies To: Job [v]

Job Rules

A pre-requisite 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.

+ Add... Edit... Remove

Name	Type	Owner
SAME JOB RULE	All Job Types	SYSMAN

In the **Job** tab, click **Add** to add jobs for which you want to create event rules.

7. In the Add Jobs dialog box, if you select the job **By Pattern**, provide **Job name like** and select the **Job Type**. Specify **Job owner like**. For the **Specific jobs** choice, select the job. Click **OK**.
8. In the **Rules** tab, click **Create**.

In the Select Type of Rule to Create dialog box that appears, you can 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.
9. Select **Incoming events and updates to events**, and in the Create New Rule: Select Events page, do the following:
 - a. **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.

- b. Set the other criteria for which you want to set the rule as displayed in the graphic below.

Create Rule Set - Sample Job Rule Set

Step 1 of 4: **Create New Rule : Select Events**

This rule acts on events that meet the criteria you specify. Type must be specified. Most of the event rules can be defined using the event type. The options in 'Advanced Selection Options' apply to advanced scenarios and should only be used after careful consideration.

Select By:

- ☒ Type = Job Status Change
 - ☐ All events of type Job Status Change
 - ☒ Specific events of type Job Status Change

Selected events of type Job Status Change

The job status values listed are based on those enabled to generate events. To change this, go to Setup->Incidents->Job Events.

Job Status in:

- ☐ All
- ☒ Problems
- ☐ Action Required

Advanced Selection Options

- ☒ Severity In Warning
- ☒ Target type In Database Instance
- ☐ Target Lifecycle Status
- ☒ Category In Configuration
- ☒ Associated with incident Equals Yes
- ☐ Associated incident acknowledged
- ☐ Event name
- ☐ Total occurrence count
- ☐ Causal analysis update
- ☐ Comment added

Expand the entries below to get more details on the event severity associated with each job status.

Job Status	Severity	Description
Problems	Warning	The event is generated with Critical or Informational severity.
Error	Error	
Failed	Error	
Stopped	Warning	
Slipped	Warning	
Reassigned	Error	

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

Create Rule Set - Sample Job Rule Set

Step 2 of 4: **Create New Rule : Select Incidents**

Select the specific incidents to which this rule should apply.

- ☐ All new incidents and updated incidents
- ☐ All new incidents
- ☒ Specific incidents

This rule applies to all newly created incidents or incidents that match the criteria specified below.

- ☐ Rules that created the incidents
- ☐ Category
- ☐ Target type
- ☐ Severity
- ☐ Acknowledged
- ☐ Owner
- ☐ Priority
- ☐ Status
- ☐ Escalation level
- ☐ Associated with ticket
- ☐ Diagnostic Incident
- ☐ Unassigned

[> Explain the options...](#)

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 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.

9.6.2 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**.

Incident Rules - All Enterprise Rules

Create Rule Set Save Cancel

A rule set is a collection of rules that applies to a common set of objects, for example, targets, jobs, and templates. A rule contains a set of automated actions to be taken on specific events, incidents or problems. For example, individual rules can respond to incoming or updated events, incidents, or problems, and then take actions such as sending e-mails, creating incidents, updating incidents, and creating tickets.

* Name: Sample Job Rule Set

Description:

Enabled: ☒ Owner: SYSMAN How is this used?

Type: ☒ Enterprise ☐ Private

Applies To: Targets

Targets Rules

Select targets to which this rule set applies. You can exclude specific targets from the scope - e.g. all database targets except 'MyDevDB'.

☐ All targets

☐ All targets of types

☒ Specific targets

Add Groups + Add - Remove

Name	Type
No target selected	

Excluded targets: (None)

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.

9.7 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](#)

9.7.1 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.

9.7.2 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 9.7.3, "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 9.7.6, "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.

9.7.3 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 9.3, "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 9.7.4, "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 9.7.2, "Creating Corrective Actions for Metrics,"](#) for more information.

9.7.4 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.

9.7.4.1 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.

9.7.4.2 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.

9.7.5 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, then 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.

9.7.6 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 9.7.2, "Creating Corrective Actions for Metrics."](#)

9.7.6.1 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>
```

9.7.6.2 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.

9.7.6.3 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>
```

Configuring and Using Services

This chapter provides an overview of services and describes the procedures to configure and monitor services with Enterprise Manager. It contains the following sections:

- [Introduction to Services](#)
- [Creating and Configuring a Service](#)
- [Creating a Test Based Service](#)
- [Creating a System Based Service](#)
- [Setting Up and Using Service Level Agreements](#)
- [Using the Services Dashboard](#)
- [Monitoring Settings](#)
- [Setting Up Monitoring Templates](#)
- [Configuring Service Levels](#)
- [Configuring a Service Using the Command Line Interface](#)
- [Troubleshooting Service Tests](#)

10.1 Introduction to Services

The critical and complex nature of today's business applications has made it very important for IT organizations to monitor and manage application service levels at high standards of availability. Problems faced in an enterprise include service failures and performance degradation. Since these services form an important type of business delivery, monitoring these services and quickly correcting problems before they can impact business operations is crucial in any enterprise.

Enterprise Manager provides a comprehensive monitoring solution that helps you to effectively manage services from the overview level to the individual component level. When a service fails or performs poorly, Enterprise Manager provides diagnostics tools that help to resolve problems quickly and efficiently, significantly reducing administrative costs spent on problem identification and resolution. Finally, customized reports offer a valuable mechanism to analyze the behavior of the applications over time.

Enterprise Manager monitors not only individual components in the IT infrastructure, but also the applications hosted by those components, allowing you to model and monitor business functions using a top-down approach, or from an end-user perspective. If modeled correctly, services can provide an accurate measure of the availability, performance, and usage of the function or application they are modeling.

10.1.1 Defining Services in Enterprise Manager

A **service** is defined as an entity that provides a useful function to its users. Some examples of services include CRM applications, online banking, and e-mail services. Some simpler forms of services are business functions that are supported by protocols such as DNS, LDAP, POP, FTP or SMTP.

Enterprise Manager allows you to define one or more services that represent the business functions or applications that run in your enterprise. You can define these services by creating one or more tests that simulate common end-user functionality. You can also define services based on system targets, or on both system and service tests.

You can create service tests to proactively monitor your services. Using these tests, you can measure the performance and availability of critical business functions, receive notifications when there is a problem, identify common issues, and diagnose causes of failures.

Services can be of two types:

- **Generic Service:** A Generic Service is the simple service model you can create in Enterprise Manager. You can define one or more service models by associating service tests and/or associating relevant system targets that represent a critical business function.
- this section should replace: defining service tests
- **Aggregate Service:** A number of services can be combined together to form an Aggregate Service. Within the context of an Aggregate Service, the individual services are referred to as **sub-services**. An Aggregate Service can also be used as a sub-service to create other Aggregate Services

10.1.2 Modeling Services

You can define a service target type and use it to model and monitor your business applications from within Enterprise Manager. While creating a service, you can define availability, performance and usage metrics, as well as the Service Level Agreement (SLA). This section describes the following:

- [Availability of a Service](#)
- [Performance and Usage](#)
- [Service Level Agreements \(SLA\)](#)

10.1.2.1 Availability of a Service

The **availability** of a service is a measure of the end users' ability to access the service at a given point in time. However, the rules of what constitutes availability may differ from one application to another. For example, for a Customer Relationship Management (CRM) application, availability may mean that a user can successfully log on to the application and access a sales report. For an online store, availability may be monitored based on whether the user can successfully log in, browse the store, and make an online purchase.

Availability of a test can be based on the following:

- **Service Tests:** Choose this option if the availability of your service is determined by the availability of a critical functionality to your end users. Examples of critical functions include accessing e-mail, generating a sales report, performing online banking transactions, and so on. While defining a service test, choose the protocol that most closely matches the critical functionality of your business process, and

beacon locations that match the locations of your user communities. You can define one or more service tests using standard protocols and designate one or more service tests as **Key Tests**. These key tests can be executed by one or more **Key Beacons** in different user communities. A service is considered available if one or all key tests can be executed successfully, depending on your availability definition. See [Deploying and Using Beacons](#) for details on beacons and how to create them.

- **System:** The availability of a service can alternatively be based on the underlying system that hosts the service or selected components of the system. If availability is based on selected system components, you must select the components that are critical to running your service and designate one or more components as **Key Components**, which are used to determine the availability of the service. The service is considered available as long as at least one or all key components are up and running, depending on your availability definition.

10.1.2.2 Performance and Usage

You can define metrics to measure the performance and usage of the service. Performance indicates the response time of the service as experienced by the end user.

Usage metrics are based on user demand for the service, or load on the system. Note that usage metrics can be collected for a service only if the service contains a system.

Performance metrics are collected for service tests when the tests are run by beacons. You can calculate the minimum, maximum, and average response data collected by two or more beacons. For example, you can monitor the time required to retrieve e-mails from your e-mail service in San Francisco, Tokyo, and London, then compare results.

You can also collect performance metrics for system components, then calculate the minimum, maximum, and average values across all components. For example, you can monitor average CPU utilization, memory utilization, and disk I/O utilization across several hosts.

Usage metrics are collected based on the usage of the system components on which the service is hosted. For example, if you are defining an e-mail service that depends on an IMAP server, you can use the Total Client Connections metric of the IMAP server to represent the usage of this e-mail service. You can monitor the usage of a specific component or statistically calculate the minimum, maximum, and average values from a set of components. You can also set thresholds on the above metrics and receive notifications and alerts.

10.1.2.3 Service Level Agreements (SLA)

An SLA is a contract between the service provider and the customer which defines the expected quality of service over a specified period. It describes the nature and standard of service that the service provider will be providing to the customers.

In Enterprise Manager, an SLA describes a set of Service Level Objectives (SLOs) that define the service levels to be provided. An SLA can contain multiple SLOs for different business calendars and different service periods.

SLOs define the service level objectives to be provided. An SLO is a logical grouping of individual measurable Service Level Indicators (SLIs). For example, an SLO can define the percentage of time a service is available to the user, how well the service is performing in terms of response time or volume, and so on.

Service Level Indicators (SLIs) are quantifiable performance and usage metrics that can be used to evaluate the quality of a service. For more details on configuring and setting up an SLA, see [Setting Up and Using Service Level Agreements](#).

10.2 Creating and Configuring a Service

This section covers the following topics:

- [Creating a Generic Service](#)
- [Configuring a Service](#)
- [Configuring an Aggregate Service](#)

10.2.1 Creating a Generic Service

Before you create a service, you must be familiar with the concepts of service management. You must also perform the following tasks:

- Identify the locations where the Management Agents must be available to monitor the services using the appropriate service tests and protocols. For example, if your service includes HTTP based service tests or IMAP based service tests, ensure that the location of the Management Agent within your network architecture allows these tests. You must ensure that the Management Agents are installed at appropriate locations according to the network security (firewalls) and network routing guidelines.

Note that the beacon targets must already be created on the Management Agents before creating the service.

- Discover all the components for your service so that they can be listed as Enterprise Manager targets.
- Define systems on which the service is based.

To create a service, from the **Targets** menu, select **Services**. The Services main page is displayed. Select the service type, Generic or Aggregate Service from the Add drop-down list and click **Go**. The following screen is displayed:

Figure 10–1 Create Service - General Page

Create Generic Service: General

Define a service to model and monitor a business process or application.

* Name

Time Zone Use System Time Zone
Select the time zone for this service. Monitored data will be displayed using the selected time zone.

System

Select a system target that hosts this service, then mark the system's key components -- the targets critical for running this service.

System None

Component	Type	Key Component
No system selected.		

Tip
A "system" is the infrastructure used to host one or more services. A system consists of components such as hosts, databases, and other targets. The system components that you mark as "Key Components" may be used to determine service availability, or, in case of service failure, to perform root cause analysis. Click **Help** for details.

Step 1 of 7

Follow the steps in the wizard to create your service. This involves the following:

- Identifying the type of service to be created. This can be a Generic Service or an Aggregate Service.

- Specifying the name and time zone for the service. The availability of the service and the SLA computation are based on the time zone you select here. You can select any time zone in which the service has to be monitored.

Note: If you select **System Time Zone**, you must select a system on which the service is to be hosted.

- Selecting a system target that hosts this service, and then marking the key components of the system that are critical for running the service. These key components are used to determine the availability of the service and identify possible causes of service failure.

Note that this step is necessary only for system-based services. It is not needed for test-based services.

- Setting up the availability definition for the service. This can be service test-based or system-based. If you select service test, the service's availability is based on the execution of the service test by one or more key beacons, or on the system target directly. If availability is based on system, availability is determined based on the status of one or more key components of the system.
- Adding one or more beacons to monitor service tests. Click **Add** to add one or more beacons for monitoring the service. It is recommended that you use beacons that are strategically located in your key user communities in order for them to proactively test the availability of the service from those locations. If no beacons exist, you must create a new beacon. See <Creating a Beacon> for details.

Notes:

- Only a single beacon should be added from a Management Agent to monitor service tests. Adding multiple beacons from the same Management Agent to a service test is not recommended.

Beacons are targets that are used to monitor service tests, primarily to measure performance of the service or business function from a different geographic location. Thus, adding multiple beacons from the same Management Agent does not add any value.

- Beacons marked as key beacons will be used to determine the availability of the service. The service is available if one or more service tests can be successfully executed from at least one key beacon.
 - It is recommended that you create the beacons before you create the service.
-

- Defining the metrics that will be used to measure the performance of the service. Performance metrics can be based on service tests or system components. After defining the metrics, you can specify the critical and warning thresholds. You can also specify the metric that is to be displayed in a graphical format on the Service Home page.
- Defining the metrics that will be used to measure the user demand for the service. Usage metrics can be based on one or more system components. After defining the metrics, you can specify the critical and warning thresholds. You can also specify the metric that is to be displayed in a graphical format on the Service Home page.

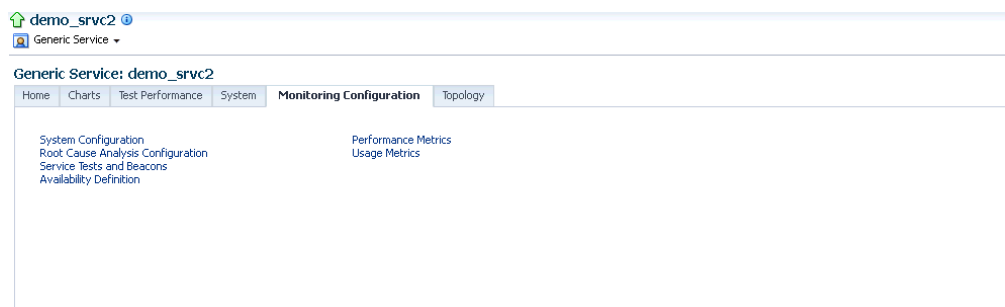
Note: You can define usage metrics for system based services only. See [Creating a System Based Service](#) for details on system based services.

- After you have completed all the steps in the wizard, click **Finish** to create your service. Refer to the Enterprise Manager Online Help for more details on these pages.

10.2.2 Configuring a Service

After you have created the service, you can configure it further by selecting an option from the Monitoring Configuration page. To configure a service, select a service from the Services main page and click **Configure** to go to the Monitoring Configuration page. The following screen is displayed.

Figure 10–2 Monitoring Configuration Page



The following options are available:

- [Root Cause Analysis Configuration](#)
- [Service Tests and Beacons](#)
- [Availability Definition](#)
- [Performance Metrics](#)
- [Usage Metrics](#)

10.2.2.1 Root Cause Analysis Configuration

You can use Root Cause Analysis (RCA) to filter a set of events to determine the cause of a higher level system, service, or application problem. RCA can help you to eliminate apparent performance problems that may otherwise appear to be root causes but which are only side effects or symptoms of the actual root cause of the problem, allowing you to quickly identify problem areas. You can view the RCA results on the Home page or Topology page of any service that is currently down. The Topology page gives you a graphical representation of the service, along with the system and component dependencies. Targets that have caused the service failure are highlighted in the Topology page.

Before running RCA, you can choose to:

- Configure the tool to run automatically whenever a service fails.
- Disable RCA by changing the default Analysis Mode to Manual.

- Define component tests for the service and thresholds for individual tests.

To configure Root Cause Analysis, follow these steps:

1. From the Service Home page, click **Monitoring Configuration**.
2. From the Monitoring Configuration page, click **Root Cause Analysis Configuration**.
3. If the current mode is set to Automatic, click **Set Mode Manual** to disable RCA. If you choose to perform the analysis manually, you can perform the analysis from the Service home page at anytime by choosing **Perform Analysis** if the service is down. If the current mode is set for Manual, click **Set Mode Automatic** to enable RCA when the state of the service and its components change.
4. Click the link in the **Component Tests** column of the table for the key component you want to manage. You can then manage the key components for the service on the Component Tests page by adding, removing, or editing component tests. When a service is down, you can drill down to the key components to verify the underlying issue. Refer to the Enterprise Manager Online Help for details on defining component tests.

Note: When you disable RCA and set it back to automatic mode, RCA does not store the previous history results for you, thus providing no history for later reference.

10.2.2.1.1 Getting the Most From Root Cause Analysis Root Cause Analysis (RCA) can provide you with great value by filtering through large amounts of data related to your services and identifying the most significant events that have occurred that are affecting your service's availability. If you are constructing your own services to manage in Enterprise Manager it is important that the services are defined with some thought and planning in order to get the most out of RCA.

The first item to consider in getting the most from RCA is the set of dependencies that your service has on other services or system components. Be sure to identify all of the system components that your service utilizes in order to accomplish its task. If you omit a key component and the service fails, RCA will not be able to identify that component as a possible cause. Conversely, if you include components in the service definition that the service does not actually depend on, RCA may erroneously identify the component as a cause of service failures.

When building service dependencies, keep in mind that you can take advantage of the aggregate service concept that is supported by Enterprise Manager. This allows you to break your service into smaller sub-services, each with its own set of dependencies. RCA considers the status of a sub-service (a service that you depend on) as well the system components or service on which the sub-service depends.

The second item to consider in getting the most from RCA is the use of component tests. As you define the system components that your service depends on, consider that there may be aspects of these components that may result in your service failure without the component itself failing. Component tests allow RCA to test the status not only of the target itself but also the status of its key aspects.

The RCA system allows you to create component tests based on any metric that is available for the key component. Remember, this includes any user-defined metrics that you have created for the component, allowing you great flexibility in how RCA tests the aspects of that component should your service fail. RCA can be configured to run in two modes. It can run automatically based on the failure of a service or can be

configured to run manually. You can decide the mode based on the Expected Service Level Agreement % of the service being monitored. If the Expected Service Level Agreement % is high, you must select the automatic mode to ensure that that possible errors and the root cause of the failure is easily detected.

10.2.2.2 Service Tests and Beacons

You can add additional service tests and specify one or more beacons that will execute these service tests. To add a service test or modify an existing service test, click the **Service Test and Beacons** link in the **Monitoring Configuration** page. The Service Tests and Beacons page appears. You can select a test type from the drop down list and create a service test. For more details, see [Creating a Test Based Service](#).

10.2.2.2.1 Deploying and Using Beacons A beacon is a target that allows the Management Agent to remotely monitor services. A beacon can monitor one or more services at any point in time.

Note: Before you create a beacon, you must ensure that the Oracle Beacon 12.1.0.2 plug-in has been deployed.

To create a beacon to run one or more service tests, follow these steps:

1. From the **Setup** menu, select **Add Targets**, then select **Add Targets Manually**.
2. In the Add Targets Manually page, select **Add Non-Host Targets by Specifying Target Monitoring Properties** option.
3. Select **Beacon** from the Target Type drop down list, select a Monitoring Agent and click **Add Manually**.
4. The Create Beacon page appears.

Figure 10–3 Create Beacon Page

Create Beacon
Select a Management Agent to enable its beacon functionality.

* Agent

* Name
You can name the beacon based on its geographic location.

Proxy Information
If the beacon will access the service through a firewall, you must specify proxy server settings.

Proxy Host

Proxy Port

Don't use Proxy for

Proxy Authentication Realm

Proxy Authentication Username

Proxy Authentication Password

Miscellaneous Information

HTTP Request Headers
Enable this feature to include an additional header in HTTP requests issued during execution of Web Transaction and HTTP Ping service tests. This HTTP request header supplies information to identify the test.

Enable Message ID Request Header ☒ No ☒ Yes
This feature must be enabled to allow Real User Experience Insight (RUEI) monitoring of Web Transaction and HTTP Ping service tests.

Web Transaction (Browser)
For Windows agents, specify the account credentials to be used when launching a browser for web transaction playback.

Username

Password

5. Enter the following details:
 - Name: Name of the beacon being created.
 - Agent: Select the Management Agent on which the beacon will be running.

- **Proxy Information:** If the beacon is accessing the service through a firewall, you must specify the proxy server settings as follows:
 - **Proxy Host and Port:** The name of the proxy server host and through which the beacon communicates.
 - **Proxy Authentication Realm:** The authentication realm (used for Basic and Digest authentication schemes) that is used to verify the credentials on the proxy server.
 - **Proxy Authentication Username:** The (fully qualified) username to be used for proxy server authentication.
 - **Proxy Authentication Password:** The accompanying password to be used for proxy server authentication.
- **Enable Message ID Request Header:** Select the checkbox to include an additional header in HTTP requests issued when Web Transactions and HTTP Ping service tests are executed. This allows Real User Experience Insight (RUEI) monitoring of Web Transactions and HTTP Ping tests.
- **Web Transaction:** For Windows agents, specify the account credentials to be used when launching a browser to playback a Web transaction.

Note: You can launch the Create Beacon option from the Services menu. From the **Services Features** menu, select **Beacons** and click **Add** to launch the Create Beacon page.

6. Click **Create** to create the beacon and return to the Beacon Home page. You can now use the beacon to monitor service tests.
7. From the Generic Service menu, select **Administration**, then select **Service Tests and Beacons**. You will see a list of service tests that have been enabled along with a list of beacons.
8. Select the service test to be monitored, then from the Beacons table, select the beacon that you have created. Indicate if it is a key beacon.
9. Click **Verify Service Test** to execute the service test by the selected beacon.

10.2.2.2 Configuring the Beacons This section lists additional beacon related configuration tasks.

- **Configuring SSL Certificates for the Beacon:** When a beacon is used to monitor a URL over Secure Sockets Layer (SSL) HTTPS URL, the beacon must be configured to recognize the Certificate Authority that has been used by the Website where that URL resides.

To use the SSL option with the Port Checker test, you may need to add additional certificates to the Management Agent's monitoring wallet. To add an additional certificate, follow these steps:

1. Obtain the certificate, which is in Base64encoded X.509 (.CER) format, in the `b64SiteCertificate.txt` file. (The file name may be different in your configuration.) An example of the contents of the file is given below:

```
-----BEGIN CERTIFICATE-----
MIIDBzCCAnCgAw...
..... base 64 certificate content .....
-----END CERTIFICATE-----
```

This file is stored in the Home directory of the Management Agent as `agent_home/sysman/config/b64InternetCertificate.txt` file.

2. In the Oracle Home of the Management Agent monitoring the wallet, run the following commands to add the certificate to the Management Agent:

```
{ORACLE_HOME}/bin/mkwallet -i welcome
{ORACLE_HOME}/sysman/config/monwallet
{ORACLE_HOME}/sysman/config/b64SiteCertificate.txt NZDST_CLEAR_PTP
```

3. Start the Management Agent.

- **Configuring Dedicated Beacons:** Beacon functionality on an agent requires the use of an internal Java VM. The use of a Java VM can increase the virtual memory size of the agent by several hundred megabytes. Because of memory constraints, it is preferable to create beacons only on agents that run on dedicated hosts. If you are running large numbers of tests (e.g., several hundred per minute) on a given beacon, you may also wish to install that beacon's agent on a dedicated host. To take full advantage of dedicated hardware, edit the agent's `{ORACLE_HOME}/sysman/config/emd.properties` file, as follows:
 - Set the property, `ThreadPoolModel=LARGE`. This allows the agent to simultaneously run many threads.
 - Set the property, `useAllCPUs=TRUE`. This allows the agent to run on multiple CPUs simultaneously.
 - Append `-Xms512m -Xmx1024m` to the `agentJavaDefines` property. This increases the Java VM heap size to 1024 MB.
- **Configuring a Web Proxy for a Beacon:** Depending on your network configuration, the beacon may need to be configured to use a Web proxy. To configure the Web proxy for a beacon, search for the beacon in the All Targets page. Select the beacon you wish to configure and click **Configure**. Enter the properties for the Web proxy. For example, to set up `www-proxy.us.oracle.com` as the beacon's Web proxy, specify the values as the following:

```
Proxy Host: www-proxy.us.oracle.com
Proxy Port: 80
Don't use Proxy for: .us.oracle.com,.oraclecorp.com
```

Note: You cannot play Siebel service tests and Web Transaction (Browser) service tests on the same machine.

10.2.2.3 Availability Definition

You can modify the availability definition (service test-based or system-based) for the selected service. If availability is based on service tests, you can specify whether the service should be available when:

- All key service tests are successful (Default)
- At least one key service test is successful

Note: A service test is considered available if it can be executed by at least one key beacon. If there are no key beacons, the service test will have an **unknown** status.

If availability is based on the key system components, you can specify whether the service should be available when:

- All key components are up (Default)
- At least one key component is up

You can also mark one or more components as key system components that will be used to compute the availability of the service. Key system components are used to determine the possible root cause of a service failure. For more information, refer to ["Root Cause Analysis Configuration"](#) on page 10-6.

You can also indicate whether the service test is a key test by enabling the Key Service Test checkbox. Only key service tests are used to compute the availability of the service. You can then select the beacons that will be used to execute the key tests and determine the availability of the service.

10.2.2.4 Performance Metrics

Performance metrics are used to measure the performance of the service. If a service test has been defined for this service, then the response time measurements as a result of executing that service test can be used as a basis for the service's performance metrics. Alternatively, performance metrics from the underlying system components can also be used to determine the performance of the service. You can do the following:

- Add a performance metric for a service test. After selecting a metric, you can choose to:
 - Use the metric values from one beacon. Choose this option if you want the performance of the service to be based on the performance of one specific location.
 - Aggregate the metric across multiple beacons. Choose this option if you want to consider the performance from different locations. If you choose this option, you need to select the appropriate aggregation function:

Table 10–1 Beacon Aggregation Functions

Function	Description
Maximum	The maximum value of the metric from data collected across all beacons will be used. Use this function if you want to measure the worst performance across all beacons.
Minimum	The minimum value of the metric from data collected across all beacons will be used. Use this function if you want to measure the best performance across all beacons.
Average	The average value of the metric will be used. Use this function if you want to measure the 'average performance' across all beacons.
Sum	The sum of the metric values will be calculated. Use this function if you want to measure the sum of all response times across each beacon.

Note: If you are configuring a Web transaction, you can specify the **Source** which can be transaction, step group or step. Based on this selection, the metric you add will be applicable at the transaction, step group, or step level.

- Add a performance metric for the underlying system components on which the service is hosted. After selecting a metric for a target, you can choose to:
 - Use the metric from a specific component. Choose this option if you want the performance of the service to be based on the performance of one specific system component.
 - Aggregate the metric across multiple components. Choose this option if you want to consider the performance from multiple components. If you choose this option, you need to select the appropriate aggregation function.

Table 10–2 System Aggregation Functions

Function	Description
Maximum	The maximum value of the metric across all components will be used as the value of this performance metric for the service.
Minimum	The minimum value of the metric across all components will be used as the value of this performance metric for the service.
Average	The average value of the metric across all components will be used.
Sum	The sum of values of metrics across all components will be calculated.

Note: When a system is deleted, performance metrics associated with the system will not be collected.

- Edit a performance metric that has been defined. For service test-based performance metrics, you can modify the beacon function that should be used to calculate the metric values. For system-based performance metrics, you can modify the target type, metric, and whether the aggregation function should be used. You can also modify the Critical and Warning thresholds for the metric.
- Delete a performance metric that has been defined.

10.2.2.5 Usage Metrics

Usage metrics are used to measure the user demand for the service. Usage metrics are collected based on the usage of the underlying system components on which the service is hosted. You can do the following:

- Add a usage metric. After selecting a metric for a target, you can choose to:
 - Use the metric from a specific component. Use this option if you want to monitor the usage of a specific component.
 - Aggregate the metric across multiple components. Use this option if you want to statistically calculate the usage across multiple components. If you choose this option, you need select the appropriate aggregation function.

Table 10–3 Aggregation Functions - Usage Metrics

Function	Description
Maximum	The maximum value of the metric across all components will be used as the value of this usage metric for the service.
Minimum	The minimum value of the metric across all components will be used as the value of this usage metric for the service.
Average	The average value of the metric across all components will be used.

Table 10–3 (Cont.) Aggregation Functions - Usage Metrics

Function	Description
Sum	The sum of the values of metrics across all components will be calculated.

- Edit a usage metric that has been defined.
- Delete a usage metric that has been defined.

Note that only metrics from system targets can be added as usage metrics. Metrics from tests are not indicative of usage, and therefore cannot be added as usage metrics.

10.2.3 Configuring an Aggregate Service

Aggregate services consist of one or more services, called subservices. A subservice is any service created in Enterprise Manager. The availability, performance, business criteria, and usage for the aggregate service depend on the availability, performance, business criteria, and usage for the individual subservices comprising the service. To create an aggregate service, navigate to the Services main page, select Aggregate Service from the Add drop-down list and click **Go**. The Add Aggregate Service page appears. Creating an Aggregate Service involves the following:

- Specifying the name and time zone for the service.
- Adding the services that make up this aggregate service.
- Establishing the availability definition for the aggregate service. Availability of an aggregate service depends on the availability of its constituent subservices. The availability for a subservice may depend on the successful execution of a service test or on the availability of the system components on which the subservice runs, depending how the subservice was defined.
- Defining the metrics used to measure the performance of your aggregate service. You can add performance metrics from single subservices, or based on statistical aggregations of more than one metric. Once you have selected the performance metrics, you can set the thresholds used to trigger incidents, or remove metrics you no longer want.
- Defining the metrics used to measure the usage of your aggregate service. Usage metrics can be based on the metrics of one or more system components. You can add usage metrics from single subservices, or based on statistical aggregations of more than one metric. Once you have selected the usage metrics, you can set the thresholds used to trigger incidents, or remove metrics you no longer want.
- Defining the metrics that are used to measure of the performance of business in the organization. These metrics are based on business indicators that can assess the business performance. You can add business metrics from single subservices, or based on statistical aggregations of more than one metric. Once you have selected the business metrics, you can set the thresholds to trigger incidents, or remove metrics you no longer want.

After you have created an aggregate service, you can add or remove its constituent subservices, modify the availability definition and add or delete performance or usage metrics. Refer to the Enterprise Manager Online Help for details on these operations.

WARNING: If you delete or remove a subservice from an aggregate service, the aggregate service performance, usage, and business metrics may be affected if they are based on a deleted subservice's metrics.

10.3 Creating a Test Based Service

You can create different types of service tests based on the protocol and the location of the beacons. From the Service Tests and Beacons page, you can do the following:

- Add one or more service tests for your service. Select the Test Type and click **Add**. Some of the test types that can be defined are ATS, FTP, Web Transaction, DNS, SOAP and others.
- After you have created the service test, you must enable it. If your service test is not enabled, it will not be executed by any of the beacons. You can define one or more service tests as key tests. These key tests are used to monitor the availability and performance of your service. Only service tests that are enabled can be designated as key tests. To set up a service test as a key test, click the **Availability Definition** link at the bottom of the page.
- Create, add, or remove a beacon. When you identify the beacon locations, select locations on your internal network or on the Internet that are important to your e-business. These are typical locations where your end users are located. For example, if your business is hosted in Canada and you have customers in the United States, use a beacon installed on a host computer in the United States to measure the availability and performance of your applications.
- After you have created the service test, you can verify it by clicking **Verify Service Test**.

Note:

- While defining a SOAP (Simple Object Access Protocol) service test, if the WSDL URL to be accessed is outside the company's intranet, proxy settings need to be added to the \$OMS_HOME/sysman/config/emoms.properties file.

For example, to set up `www-proxy.us.oracle.com` as proxy, specify the values as follows:

```
proxyHost=www-proxy.us.oracle.com
```

```
proxyPort=80
```

```
dontProxyFor=us.oracle.com,oraclecorp.com
```

The `proxyUser`, `proxyPwd`, `proxyRealm`, and `proxyPropsEncrypted` properties are used to configure an authenticated proxy. After you have modified the proxy settings, you must restart all the OMSes for the changes to be effective.

- The Forms Transaction test type has been deprecated in Enterprise Manager 12c. Forms transactions created in earlier releases can still be used but you cannot create new Forms Transaction test types. You must create a Generic Service target and create an ATS Transaction using OATS EBS/Forms Load test scripts. This ATS test type is used to monitor Oracle Forms applications.
- The Web Transaction test type is in maintenance mode only. To monitor Web applications, we recommend that you create an ATS load script and use the ATS Transaction test type to monitor Web applications. See [Creating an ATS Service Test Using OATS Load Script](#) for details.

The creation of different types of service tests is covered in detail in the Enterprise Manager Online Help. In this chapter, we have covered the creation of the ATS test type as an example.

10.3.1 Creating an ATS Service Test Using OATS Load Script

You can use the Oracle Application Test Suite (OATS) to define an Openscript Transaction Service Test. This test is used to enable beacon application transaction monitoring using Openscript load testing scripts. Openscript is a component of OATS and provides advanced capabilities to record and play back various types of Web transactions, such as web/HTTP, Oracle EBS/Forms, Oracle Fusion/ADF, Siebel, Adobe Flex, and so on.

By using ATS load scripts, you can:

- Reuse ATS testing scripts for production application transaction monitoring as part of application lifecycle management.
- Expand the beacon capabilities by:
 - Supporting complex application flows with mixed application types and protocols such as HTTP and Oracle Forms application in one flow.
 - Supporting protocol based Siebel application monitoring.
 - Providing Databank support.

- Incorporating enhanced scripting and debugging features from Openscript.
- Adding the latest script modules and features without updating Enterprise Manager.

Creating an ATS service test involves the following:

- Visit <http://www.oracle.com/technetwork/oem/app-test/index-084446.html> to download Openscript.
- Launch the installer and follow the steps to install Openscript.
- Record a new ATS transaction script by following these steps:
 - Launch Openscript and from the **File** menu, select **New**.
 - Select the type of script to be created and click **Next**. Some of the script types you can create are Adobe Flex, Oracle Fusion / ADF, Siebel, Database, Java Code Script, Web/HTTP and so on.
 - Select the location where the script is to be stored, enter a script name, and click **Finish**.
 - From the **View** menu, select **Openscript Preferences** to set the recording preferences.
 - Select **Record Category** and then select **HTTP Preferences**. Change the **Record Mode** from **Web** to **HTTP**. This ensures that scripts are played back correctly.
 - Select **Record** from the **Script** menu. The browser automatically opens when you start recording.
 - Load the web page where you want to start recording into the browser. Navigate the web site to record page objects, actions, and navigations. When you have finished navigating pages, close the browser and click **Stop**.
 - Select **Playback** from the **Script** menu to verify that the script has been properly recorded. Watch the application flow being played back in the play pane. Make sure the message log pane does not have any errors or failures. Save the script.
- After the script has been recorded, from the **File** menu, select **Export Script**.
 - Specify the location in which the script is to be saved. You can save the script in a repository or workspace.
 - Enter a name for the script and click **Finish**. A script bundle (.zip) is created. Make sure that the script bundle is self contained. See [Creating a Self Contained Zip File](#).

Note: If the script file is very large, uncheck the **Recorded Data** option.

- Log into Enterprise Manager, upload the script bundle, and create an ATS service test. See [Creating an ATS Service Test](#) for details.

For more details on OATS, please refer to *Oracle® Functional Testing OpenScript User's Guide*.

10.3.1.1 Creating a Self Contained Zip File

You must ensure that the zip file is self contained and contains the following:

- **<txn name>.jwg**: The archive file that contains the compiled script executable to be run by Execution Engine.
- **script.java**: The actual script Java source file.
- **<script name>-descriptor.xml** – Describes the step hierarchy
- **script.xml**: Describes the variables in the databank.
- **modules.properties**: Describes which internal modules are required for the engine.
- **Assets.xml**: Describes the dependent resources used by the root scripts including databank files, sub scripts, object libraries, and so on.
- **Databank Files**: The databank files used by the script while substituting different variable values.
- **Object Libraries**: The libraries that contain user-defined object identification rules and names. This is only applicable for functional testing scripts
- **Dependent Scripts**: A script can call out to other script.

10.3.1.2 Creating an ATS Service Test

To create an ATS service test, follow these steps:

1. From the **Targets** menu, select **Services**.
2. Select Generic Service from the **Type** drop down list and click **Add**.
3. Enter a name for the service and select a time zone.
4. Click **Next**. In the Create Generic Service: Availability page, select **Service Test**.
5. In the Create Generic Service: Service Test, select the **Test Type** as ATS Transaction.
6. Enter a name, description, and collection frequency for the service test.
7. In the ATS Zip Archive region, click **Upload**. Click **Browse** and upload the script bundle that has been exported from Openscript. Click **Continue**.
8. You will return to the ATS Service Test page where you can specify the following:
 - **Usage Options**: You can configure the script variable values by selecting the required usage option. You can either use the values recorded during the transaction or use the databank. A databank is an external CSV file that ATS scripts can refer to supply different input values over multiple iterations of the same script. For example, a login script can use a databank file, named `login_credential.csv`, to supply different login credentials during iteration.

You can select:

- **Use Recorded Values**: While playing back the transaction, the beacon uses the recorded values in the script.
- **Use Values From EM Test Property**: You can specify values in the databank columns. These values are used by the beacon while playing back the script. This is useful if the same value is to be used for each variable. If variables defined as test properties, the value can be easily modified without having to modify script bundle or databank files.
- **Loop Through All Databank Records**: While playing back the transaction, the beacon will go through each row in the script. For

example, the first iteration will use the first rows of all the databanks. The second iteration will use the second rows of data and so on.

- **Encryption Password:** If you have configured ATS Openscript to encrypt script data (using the **Openscript View File > Openscript Preferences > Generic > Encryption**) option, when you create the scripts, you need to enter the same encryption password as specified in the ATS openscript, so that beacon can play back the script properly.
- **Default Playback Options:** The default playback options the beacon uses to play back the ATS Script.
- **Additional Playback Options:** If additional playback options have to be specified, you can specify them here.

10.3.1.3 Databanking and Parameterization

You can parameterize recorded script inputs to perform data driven testing. Examples of inputs that can be parameterized include user name, password on the login page, data entered in the search field, recorded navigations or user actions, and so on.

You can use databanks as the data source for parameterizing script inputs. Databanks are one or more external comma-separated value (CSV) or text (TXT) files that provide inputs to script parameters. Multiple Databank files can be attached to a single script and users can specify how OpenScript assigns data during script playback.

You can select data input values from an `external.csv` file and substitute the variable values with the values from the Databank. The field names are on the first line of the file separated by commas (no spaces). The field data is on subsequent lines separated by commas (different line for each record, no spaces around commas). An example is shown below:

```
FirstName,LastName,Mail,Phone
John,Smith,JohnS@company.com,x993
Mary,Ellen,MaryE@company.com,x742
```

To use the databank records, follow these steps:

1. Open or create a script project.
2. Configure the Databank to use with a script in the Assets Script Properties.
3. Select the script node where you want to use the Databank record.
4. Select the Script menu and then select **Other** from the Add sub menu.
5. Expand the General node and select **Get Next Databank Record**. Click **OK**.
6. Select Databank avitek alias to specify the Databank file from which the records are to be retrieved.
7. Click **OK**. A `GetNextDatabankRecord`: databank alias node will be added to the script. In the Java Code view, the `getDatabank("databank alias").getNextDataBankRecord()` method will be added to the script code as follows:

```
getDatabank("avitek").getNextDatabankRecord();
```

After you have configured the databank for use with a script, you can map the databank files to specific script parameters. To map databank fields to script parameters, follow these steps:

1. Expand the [4] Oracle WebLogic Server - Medical Record Sample Application script tree node.

2. Right-click the usernameInput parameter and select **Substitute Variable**. The Substitute Variable window opens with the databank field names listed.
3. Select the Username field and click **Finish**. The parameter value changes to a databank variable in double braces
`{{db.avitek.Username,fred#@golf.com}}`
4. Right-click the passwordInput parameter and select **Substitute Variable**. The Substitute Variable window opens with the databank field names listed.
5. Select the Password field and click **Finish**. The parameter value changes to a databank variable in double braces `{{db.avitek.Password,weblogic}}`
6. Save the script.

For more details on setting up the Databank, see the *Oracle® Functional Testing OpenScript User's Guide*.

10.3.1.4 Parameterizing URLs

You can create variables to use for URLs in a script. If you need to change the base URL of a script, parametrizing the URLs provides a quick way to re-baseline a script to use a new URL.

To parameterize URLs, follow these steps:

1. After you have recorded your script, from the **Tools** menu, select **Parameterize URLs**.
2. Select the URL to parameterize and enter a variable name to use for the URL. Click **Next**.
3. Select or unselect the checkbox to specify the instances of the URL that are to be changed. Click **Finish**.
4. In the Java Code view, the `getVariables().set("variable name", "value",scope);` method will be added to the script code in the `initialize()` section as follows:

```
getVariables().set("myServerVar", "http://myServer.com",
Variables.Scope.GLOBAL);
```

5. Repeat these steps to parameterize other URLs in the script.

10.3.1.5 Success or Failure Validation

You can perform text matching from TreeView and CodeView. For example, if you enter a string "hello" in a Google search window, the text matching is as follows:

```
web.document("/web:window[@index='0' or
@title='Google']/web:document[@index='0']").assertText("MatchText", "hello", Source.DisplayContent, TextPresence.PassIfPresent, MatchOption.Exact);
```

Source - enum - (Html,Display Content)

Html - Raw Html including Tags

Display Content - Html without tags

MatchOption - Not Case Sensitive

- **Exact - sensitive:** Matches any part of source string. For example, if the text entered is **abcdef**, if you enter **abc**, the string will match.
- **ExactEntireString:** Matches the exact source string.

- **RegEx - Not Case Sensitive:** Matches source string and sub-string. For example, if the text entered is **abcdef**, you can enter **a.*d**.
- **RegExEntireString:** Matches the entire source string only. For example, if the text entered is **abcdef**, you can enter **a.*f**.
- **Wildcard - wildcard pattern:** Matches source string and substring. For example, if the text entered is **abcdefghijklm**, you can enter **a?c*f**.
- **WildcardEntireString:** Matches the entire source string. For example, if the text entered is **abcdefghijklm**, you can enter **a*m**.

10.3.1.6 Using Beacon Override

You can use the beacon override feature to specify different variable values for a test running on a different beacon. To do so, follow these steps:

1. Databank a script.
2. Select **Use EM Test Property** option.
3. Define beacon by specifying the sensitive and non-sensitive values as follows:

```
Databank_Alias>."<COLUMN_NAME>"="VALUE", <Databank_Alias>."<COLUMN_NAME>"="VALUE", ...
```

For example:

```
FusionCredentials."host"="fs-aufsn4x0cxf", FusionCredentials."hostlogin"="login-aufsn4x0cxf", FusionCredentials."username"="faadmin", FusionCredentials."password"="fusionfal"
```

10.3.1.7 Updating the Databank File

To update the ATS test script, follow these steps:

1. From the **Generic Service** menu, select **Administration**, then select **Service Tests and Beacons**.
2. From the Service Tests table, select ATS Transaction test type, and click **Edit**.
3. In the ATS Zip Archive region, click **Download**. Select the location where it is to be downloaded and click OK.
4. Edit the .csv file using a spreadsheet editor and save the changes.
5. Log into Enterprise Manager and navigate to the ATS Transaction Test page.
6. In the ATS Zip Archive region, click **Upload** to upload the updated file.

10.3.1.8 Using SLM Header for RUEI Integration

If the ATS service test data is to be monitored by RUEI, you must specify the `x-oracle-slm-message-id` request header in the **Additional Playback Options** field. The format is in the form: `name1:value1;name2:value2;name3:value3`.

```
For example, x-oracle-slm-message-id: bcn=<beacon_name>;  
svc=<service_name>;test=<test_name>;step={{@getTopLevelStepName()}}
```

10.3.2 Using the Transaction Recorder

You can record a transaction using an intuitive playback recorder that automatically records a series of user actions and navigation paths. You can play back transactions interactively, know whether it is internal or external to the data center, and understand

the in-depth break-out of response times across all tiers of the Web application for quick diagnosis.

You must install the transaction recorder in your computer to record transactions. The transaction recorder is also used for playing back and tracing transactions. The transaction recorder is downloaded from the Enterprise Manager Cloud Control server the first time any of these actions is performed. The transaction recorder requires some Microsoft libraries to be installed in your computer. If these libraries are not present during installation, they are automatically downloaded and installed from the Microsoft site. Make sure that your computer has access to the Internet to download these files. After the installation has been completed, you may need to restart your computer to make the changes effective.

10.4 Creating a System Based Service

A system is the set of infrastructure components, for example hosts, databases, and application servers that work together to host your applications. Before you create a service, you must specify the system that will be used to host your service.

After you have selected the system, you must mark one or more components as key components that are critical to running your service. These key components are used to determine the availability of the service and identify possible causes of service failure for root cause analysis.

To create a system based service, follow these steps:

- Identify the type of service being created (Generic or Aggregate Service).
- Enter a name for the service and specify the time zone.
- In the Availability page, select **System** in the **Define Availability based on** list.

Figure 10–4 Availability Based on System

Create Generic Service: Availability

You can define availability of the service based on either the execution of the service test by the key beacons, or the availability of key components of the service, or the availability of the system.

Define availability based on: **System**

☐ Availability of System Target directly Oemrep_Database_sys
☒ Availability of selected components of the system
☐ All key components are up
☐ At least one key component is up

System

Mark the system's key components -- the targets critical for running this service.

System: Oemrep_Database_sys
Time Zone: (UTC-08:00) US Pacific Time

Component	Type	Key Component
Oemrep_Database	Database Instance	<input checked="" type="checkbox"/>

Tip

Select the method of availability testing that best models your business process or application.

Oracle recommends using:

- **Service Test** to define availability for most cases. A service test emulates the way clients access the service.
- **System** when a suitable service test is unavailable.

Tip

A "system" is the infrastructure used to host one or more services. A system consists of components such as hosts, databases and other targets. The system components that you mark as "Key Components" may be used to determine service availability, or, in case of service failure, to perform root cause analysis. Click **Help** for details.

- Specify whether the **Availability is based on the System Target directly** or **Availability of selected components of the system**. If you select the first option, availability is based on the system target on which the service is based. If you select the second option, availability is based on the key components of the system that you have selected.
- Follow the rest of the steps in the wizard and create a system based service.

10.5 Setting Up and Using Service Level Agreements

A service level agreement (SLA) is a contract between a service provider and a customer on the expected quality of service for a specified business period. An SLA consists of one or more service level objectives (SLOs) which define the service levels to be provided. Whether an SLA is satisfied or not is based on the evaluation of the underlying SLOs. Service level indicators (SLIs) allow SLOs to be quantified and measured. An SLO can have one or more SLIs.

To create an SLA, follow these steps:

1. Log in to Enterprise Manager as a user with `EM_ADMINISTRATOR` role.
2. From the **Targets** menu, select **Services**.
3. Click on a Generic Service target on the list. The Service Home page is displayed.
4. From the **Generic Service** menu, select **Service Level Agreement**, then select **Configuration**. The Service Level Agreement Configuration page appears.
5. On this page, you will see a list of all the SLAs defined for the selected service. Select an SLA from the list to view the details in the Service Level Agreement Details table. You can create an SLA or make a copy of an existing SLA (Create Like).
6. In the Service Level Agreement region, click **Create**. The Configure Service Level Agreement page appears.

Figure 10–5 Create Service Level Agreement: Configure Service Level Agreement

Enter the following details:

- Name and description of the SLA.
- Name of the customer for whom the SLA is being created.
- The Lifecycle Status of the SLA. When an SLA is being created, it will be in the Definition Stage. For more details on the Lifecycle Status, see [Lifecycle of an SLA](#).
- Specify the SLA Period. This is the contractual time period for which the SLA is determined and/or evaluated for compliance. (ie. quarterly, monthly, weekly SLA). Click the **Select** icon and select Monthly, Weekly, or Daily. Enter the Frequency which the SLA is to be evaluated and the date from which the SLA is to be evaluated. The SLA goals are reset when the SLA is evaluated.

For example, if you specify the SLA Evaluation Period as Monthly, Frequency as 12 and the date as 09/01/12, the SLA will be evaluated on that date followed 11 consecutive evaluations in the months of October, November, and so on.

- Specify the SLA Agreement Period. This is the **From** and **To Date** for which the recurring SLA periods are in effect. If you do not specify the **To Date** here, the SLA will have an Indefinite expiry date.
- An SLO may sometimes not be evaluated due to planned downtime or blackouts that have been scheduled for a service. In the Service Level Agreement Evaluation Options region, select the **Include blackout times (planned downtimes) in Service Level Objective evaluation** checkbox and specify whether the blackout times are to be included in the SLO evaluation. You can choose to:
 - Include time as met
 - Include time as not met
 - Exclude the blackout time during the overall computation of the SLO.

For example, if the blackout or planned downtime for the week is 1 day, then the weekly availability is $(7-1) / (7-1)$ days which is still 100% availability.

By default, the **Include blackout times (planned downtimes) in Service Level Objective evaluation** option is not selected.

7. Click **Next**. In the Service Level Objectives page, define one or more SLOs that are to be part of the SLA. You can select the Evaluation Condition for the SLA which can be:
 - All Service Level Objectives must be met.
 - At least one Service Level Objective must be met.

An SLA must have at least one SLO. More than one SLO can be active at any given time. You can either specify if all SLOs or at least one SLO should be met.
8. Click **Create** to define a new SLO. See [Creating a Service Level Objective](#) for details.
9. You can add more SLOs or edit the SLO you have defined. Click **Next**. In the Enable Service Level Agreement page, you can specify when the SLA is to be enabled. You can select:
 - Do Not Enable: If the SLA is not enabled, it will be in the Definition state and can be modified if required.
 - Enable Now: If the SLA is enabled, it cannot be modified as it will be in an Active state.
 - Enable Later: The SLA can be enabled later on a specified date.
10. Click **Next**, review details of the SLA, and click **Submit**. The SLA will be enabled on the specified date and you will return to the Service Level Agreement Configuration page.

10.5.1 Actionable Item Rules for SLAs

The table below shows a list of actions that can be performed on an SLA based on its status.

Status of SLA	Create Like	Edit	Enable	Disable	Delete
Definition	Yes	Yes	Yes	No	Yes
Scheduled	Yes	Yes	No	Yes	No
Active	Yes	No	No	Yes	No

Status of SLA	Create Like	Edit	Enable	Disable	Delete
Retired	Yes	No	No	No	Yes

- An SLA in a **Scheduled** or **Active** state cannot be directly deleted. You have to disable the SLA before you can delete it.
- When you edit an SLA in a **Scheduled** state, the status of the SLA changes to **Definition**.

10.5.2 Creating a Service Level Objective

A Service Level Objective measures the service level of one or more indicators for a specified measurement window. Service Level Objectives (SLOs) define the service levels to be provided. You can specify if the SLA is considered to be satisfied if:

- All Service Level Objectives are met.
- At least one Service Level Objective is met.

To create an SLO, follow these steps:

1. Click **Create** in the Configure Service Level Objective page. The Create Service Level Objective page appears.

Figure 10–6 Create Service Level Objective

Configure Service Level Objective : Create Service Level Objective

This is the first step in defining a new SLO. This is a sub-wizard of the overall create Service Level Agreement wizard. On completion or cancel of this, the flow will return to the Create Service Level Agreement wizard.

Service Level Objective Conditions

A Service Level Objective measures the service level of one or more indicators for a specified measurement window. Service Level Objectives (SLOs) define the service levels to be provided. Specify SLOs for an SLA. The Service Level Agreement defined is logical grouping (AND / OR) of all the SLOs.

* Name: SLO1 Type: Performance

Service Level Percentage

* Expected Service Level (%): 85.000 ☒ Generate service level warning when below warning alert level

* Warning Alert Level (%): 90.000

Measurement Window

Select the time periods for the SLO to be tracked and measured. One or more calendars can be chosen and configured as either includes and excludes. Including a calendar will stretch the measurement window and excluding a calendar will constrict the measurement window.

View	Add	Remove	Description	Include/Exclude
Business Calendar			All Day Monitoring	<input checked="" type="radio"/> Include <input type="radio"/> Exclude

2. Enter the following details:
 - Name of the SLO being defined.
 - Type of SLO: The SLO can be based on Availability or Performance metrics.
 - Expected Service Level %: This indicates the percentage of time the SLO conditions are met to ensure that the SLA is satisfied.
 - Warning Alert Level %: If the SLO conditions do not meet the specified threshold, a critical alert is generated.

For example, if the Expected Service Level% is 90% and the Actual Service Level% is in the range of 90 to 99%, a Warning Alert is generated. If the Actual Service Level% is lesser than 90%, a Critical Alert is generated. This indicates that the SLA has been breached. If the Actual Service Level% is greater than 99%, it indicates that the SLA conditions have been satisfactorily met.

- Measurement Window: The time periods during which the SLO is in effect. A measurement window can have more than one time period assigned. For example, a measurement window can be configured as weekday peak hours

which are Monday to Friday, from 9AM to 6PM and the weekend peak hours as 10AM to 2PM.

While creating an SLO, you can choose more than one Business Calendar for an SLO. For example, suppose you want to evaluate each SLO from 8AM to 5PM except at lunch time (12PM to 1PM). You can create two measurement windows and exclude the lunch time from being measured.

Another example of merging two measurement windows is when you want to combine weekly evaluation with calendar evaluation. If you want to evaluate an SLO every Monday and on the 15th of every month, you can create two monitoring windows and include these conditions in both the windows.

By default, there are 3 predefined business calendars. You can also create your own calendar. See [Defining Custom SLA Business Calendars](#) for details.

3. Click **Next**. In the Create Service Level Indicators page, you can add one or more SLIs or conditions that allow the SLO to be measured.

Figure 10–7 Create Service Level Indicators

Configure Service Level Objective : Create Service Level Indicators

This step allows definition of one or more Service Level Indicators (SLI). A SLI definition requires selecting a service metric followed by the condition against which it will be measured.

Service Level Indicators

Service Level Indicators (SLIs) allow Service Level Objectives to be measured and quantified. A SLI Metric expression describes something that must evaluate to TRUE in order for the Indicator to be in the Green state. Example : In the case of availability , the expression might say that a target is in the UP state or in the BLACKOUT state. In the case of performance, the expression would say that a metric must be less than some critical threshold. The Service Level Indicator is considered to be violated if the rule specified below evaluates to false.

Evaluation Option

☒ All Service Level Indicators must be met.
☐ At least one Service Level Indicator must be met.

View

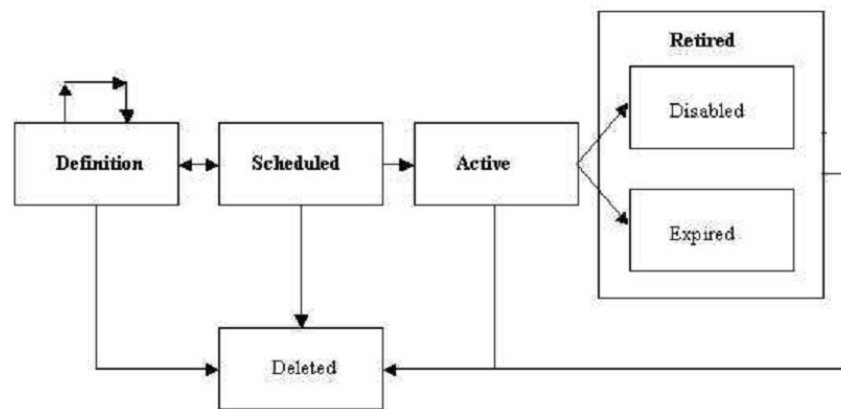
Metric Name	Comparison Operator	Value
Nursery Size (MB)	>=	20.0

For example, if you are adding a performance SLI, you can specify that the Page Load Time should be less than or equal to 3 seconds. If this condition is not met, the SLI is considered to be violated. Specify the Evaluation Condition for the SLI:

- All Service Level Indicators must be met.
 - At least one Service Level Indicator must be met.
4. Click **Add** to add one or more metrics and specify the value and the evaluation condition. Click **Submit** to return to the Configure Service Level Objective page.

10.5.3 Lifecycle of an SLA

The following diagram shows the lifecycle of an SLA.

Figure 10–8 SLA Lifecycle

The SLA lifecycle consists of the following phases:

- **Definition**: This is the stage where the SLA is created and the SLOs are defined. You can configure or edit the SLA definition till the SLA is activated.
- **Scheduled**: This stage represents the period before the SLA is scheduled to go into effect at a future date.
- **Active**: This is the stage where the start date of a scheduled SLA is reached, or when the SLA is manually enabled.
- **Retired**: This is the stage when the SLA reaches the Expiry Date or the SLA is manually disabled.
- **Disabled**: An SLA can be manually disabled before it reaches the Expiry Date. Once an SLA is disabled, it cannot be reactivated. You must use the Create Like option to create a similar SLA and enable it.
- **Expired**: This is the stage where the SLA has reached the Expiry Date and is no longer active.
- **Deleted**: An SLA can be deleted if it is the Definition or Retired stage. An SLA that is an Active or Scheduled stage cannot be deleted.

10.5.4 Viewing the Status of SLAs for a Service

You can view the status of all SLAs for a service. To view the current status of the SLAs for a service, follow these steps:

1. From the **Targets** menu, select **Services**.
2. Click on a Generic Service target on the list. The Service Home page is displayed.
3. From the **Generic Service** menu, select **Service Level Agreement**, then select **Current Status**. The Service Level Agreement Current Status page appears.
4. This page shows a list of all the active SLAs that have been defined for this service. For each SLA, the SLA Status, SLA Evaluation Period, and the Service Level Objectives are displayed.
5. Select an SLA to view detailed information in the SLA. The following details are displayed:

- **Tracking Status:** This is the instant status of the SLI. For an Availability SLO, it is the status of the target. For a Performance SLO, it is the value of the Performance or Usage metric at a specific point in time.
- **Service Level (%) :** The percentage of time (from the beginning of the current evaluation period till the current date) the SLO conditions are met or the Tracking Status is **true**. If the Actual Service Level % is lesser than the Expected Service Level %, or the SLO conditions are met, the Service Level % graph is green.
- **Type:** This is the type of SLOs that have been defined for the SLA. This can be based on Availability or Performance metrics. An Availability SLO is based on the Response Metric [Service Target Availability]. It is specified in terms of the amount or percentage of time when the availability objective should be met. A Performance SLO gauges how well a service is performing. It includes measurements of speed and/or volume such as throughput or workload (ie. response times, transactions/hour). A Performance SLO can either be specified in terms of a set of SLIs, SLO conditions, and the amount or percentage of time when the objective should be met.
- **SLO Violation:** The violation allowances for each SLA evaluation period.
 - **Total:** The duration of the Evaluation Period * (Expected Service Level).
 - **Actual:** The time when the SLO is not met during the Evaluation Period.
 - **Remaining:** The time when the SLO could not be met without breaching the SLA. If the SLO is always met during the Evaluation Period, it indicates that there are no used allowances and the value in the Actual field will be 0.

10.5.5 Defining Custom SLA Business Calendars

Business Calendars are measurement windows that define a specific window of time in which the Service Level Objectives (SLO) are being measured. Out-of-the-box predefined business calendars are available. Apart from these, you can create custom business calendars. To create a custom business calendar, from the **Targets** menu, select **Services**. From the **Services Features** menu, select **SLA Business Calendars**.

A list of business calendars that have been defined is displayed here. You can:

- **Create:** Click **Create** to set up a business calendar. The Add / Edit Business Calendar page is appears.
- **Create Like:** Select a calendar and click **Create Like** to make a copy of this calendar.
- **Edit:** Select a calendar, click **Edit** and make the necessary changes in the Add / Edit Business Calendar page.
- **Delete:** Select a calendar and click **Delete** to delete it. You cannot edit or delete a business calendar that is associated with one or more SLAs.
- **View Associated Service Level Agreements:** A business calendar can be used by one or more SLAs. Select a business calendar and click **View Associated Service Level Agreements** to view the SLAs that are associated with this calendar.

10.6 Using the Services Dashboard

The services dashboard provides a brief summary of all service related information in a single place. It provides a consolidated view of critical aspects of a service such as

availability, performance, SLAs associated with the service, status of key system components, and so on.

10.6.1 Viewing the All Dashboards Page

To view the All Dashboards page, follow these steps:

1. From the **Targets** menu, select **Services**.
2. From the **Services Features** menu, select **Dashboards**.
3. The All Dashboards page appears where you can see a list of all dashboards that have been created.
4. From the All Dashboards page, you can do the following:
 - **Create Dashboard:** Enter a unique name in the Dashboard Name field and a description, and click **Create Dashboard**. The newly added dashboard appears in the table. To create a dashboard, you must have an EM_ADMINISTRATOR role with **Create Services Dashboard** privilege.
 - **Customize Dashboard:** Select a dashboard from the list and click **Customize** and make the changes in the Edit Services Dashboard page. The dashboard can be customized only by the user who has created it.
 - **Delete:** Select a dashboard from the list and click **Delete**. The selected dashboard is deleted. The dashboard can be deleted only by the user who has created it.
5. Click on a Dashboard Name link to drill down to the Dashboard Details page.

10.6.2 Viewing the Dashboard Details Page

This page displays the following details:

- **Service Name:** Click on the link to drill down to the Service Details page.
- **Incidents:** Any incidents that have occurred.
- **Performance / Usage Metric:** The name of the performance and usage metrics available for the service and the latest value of each metric is displayed. The Trend charts show the metric trend over the last 24 hours. Click on the Trend chart to see a detailed view over the trend.
- **SLA:** Shows the number of enabled SLAs that are in Active, Critical or Warning state.
- **Key Components:** Shows the key targets that are up or available for this service.
- **System Incidents:** Any incidents that have occurred for the underlying systems of the service are displayed.

Figure 10–9 Services Dashboard

Services > All Dashboards > Dashboard (My Dashboard)

Services Dashboard

View Filter by **Filter** Remove Filter **Email**

Service Name	Type	Status	Incidents	Performance Metric			Usage Metric			SLA
				Metric Name	Latest Value	Trend (last 24 hours)	Metric Name	Latest Value	Trend (last 24 hours)	
sys svc 1	Generic Service	↑	-	Nursery Size (MB)	-	-	Status	1	-	-
svc 1	Generic Service	↑	✖ 3	Perceived Time per Page (ms)	392		Not Configured			✖ 1
				HTML Time (ms)	65					
				Connect Time (ms)	212					
EM Console Service	EM Service	↑	-	Perceived Time per Page (ms)	129		Page Hits (per minute)	0	-	-
EM Jobs Service	EM Service	↑	-	Throughput - Job Steps/sec	0		Job Dispatcher Processing Time (% of last hour)	89.04		-
				Backlog - Jobs Steps	0					

You can filter the list of services that are listed in the dashboard. Specify a value in the Filter By field and click **Filter**. The filter will be applied on each row in all the services and the resulting list is displayed.

You can email a dashboard to one or more email addresses. Click **Email** and enter the email address and the subject of the dashboard. Click **Send**. This feature works only the http mode.

10.6.3 Customizing and Personalizing the Dashboard

You can customize a dashboard and make the changes available to all users. To customize a dashboard, select a row on the All Dashboards page and click **Customize**.

Note: The following privileges are required to cr

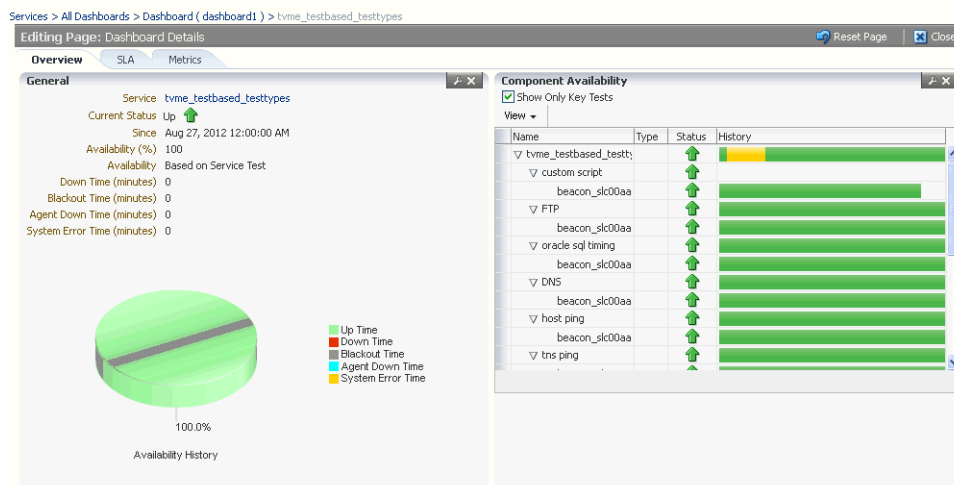
To add one or more services to the dashboard, click the Wrench icon. The Component Properties: Services Dashboard window appears. Select the type of service that you want to add to the dashboard and click **Search**. A list of services is displayed in the Available Targets table. Select one or more services that you want to add, move them to the Selected Targets table and click **Apply**. To add metrics to the respective services click on the Metrics tab and select the respective services to add metrics and click **OK**. The selected services and metrics now appear in the Services Dashboard table.

To delete a service target from the dashboard, select the row and click the Wrench icon. The Component Properties: Services Dashboard window appears. Deselect the services and metrics that are to be removed from the dashboard, click **Apply** and then click **OK**. To reset the changes you have made to the dashboard, click **Reset Page**. Any changes that have been made to the dashboard will be removed permanently. Click **Close** to exit the Edit mode.

You can make specific changes to a dashboard to suit your requirements. Click the **Personalize** icon and add or delete one or more services from the dashboard. The changes that you make will be visible only to you and the other users cannot see the changes.

10.6.4 Viewing the Dashboard Service Details Page

This page shows detailed information for the selected service.

Figure 10–10 Dashboard Details Page

It contains the following tabs:

- **Overview:** This tab provides a brief overview of the selected service. Click on the Service link to drill down to the Service Home page. It contains the following regions.
 - **General:** This region shows the name of the service, status, date from which the service is available, availability percentage, type of service (test or system based), down time, and error time.
 - **Component Availability:** This region shows the status of the components in the service. It shows the status of the component and the date from which the service has been Up. Select the Show Only Key Tests check box to view only the key service tests
- **SLA:** Shows a list of SLAs that have been enabled for this service. The name, status and the date from which the SLA is applicable is displayed. The SLA history over the last 7 days is also displayed.
- **Metrics:** This tab shows the performance and metrics charts that have been defined for this service. It also shows the incidents that have occurred for the service and the underlying systems on which the service is based.

10.7 Monitoring Settings

For each service, you can define the frequency (which determines how often the service will be triggered against your application) and the performance thresholds. When a service exceeds its performance thresholds, an alert is generated.

To define metrics and thresholds, from the Service Tests and Beacons menu, select **Monitoring Settings for Tests**. The Metric and Policy Settings page is displayed. Click the **Monitoring Settings** link. The Monitoring Settings - Thresholds page appears.

- **View By Metric, Beacon** - In this view, you can click **Add Beacon Overrides** to override the default threshold values for one or more beacons. In this case, the default thresholds will only be used for beacons without any overrides. Any new beacons added to the service will use the default thresholds. Click **Add Metric** to add one or more metrics.

- **View By Beacon, Metric** - In this view, you can click on the **Default** icon to toggle between Edit and View modes for a specific metric. In the Edit mode, you can modify the parameters of the metric. You can also modify the parameters of the metric for a specific beacon. In the View mode, the default parameters of the metric will be used.

Apart from these procedures, you can also define metrics at the step, and step group level for Web transactions. You can choose either of the following views:

- **View By Step, Metric, Beacon:** In this view, you can click **Add Beacon Overrides** to override the default threshold values for one or more beacons. In this case, the default thresholds will only be used for beacons without any overrides. Any new beacons added to the Web transaction will use the default thresholds. Click **Add Metric** to define thresholds for one or more metrics. Incidents are generated only if the value of the Data Granularity property is set to 'Transaction' for the service tests. For more information on the Web transaction properties, refer to the Create / Edit Service Test - Web Transaction help page in the Enterprise Manager Online Help.
- **View By Step, Beacon, Metric:** In this view, you can click on the **Default** icon to toggle between Edit and View modes for a specific metric. In the Edit mode, you can modify the parameters of the metric for a specific beacon. In the View mode, the default parameters of the metric will be used. Incidents are generated only if the value of the Data Granularity property is set to 'Step'.

To define the default collection frequency and collection properties, click the **Collection Settings** tab on the Monitoring Settings page. You can do the following:

- Specify the default collection frequency for all the beacons. To override the collection frequency for a specific beacon, click **Add Beacon Overrides**.
- Specify the collection properties and their corresponding values for one or more beacons.

Refer to the Enterprise Manager Online Help for more details on the defining the collection intervals and performance thresholds.

10.8 Setting Up Monitoring Templates

A monitoring template for a service contains definitions of one or more service tests, as well as a list of monitoring beacons. A monitoring template can be used to create service tests on any number of service targets, and specify a list of monitoring beacons.

A monitoring template must be created from a service target. Once the template is created, the user can edit the template, create copies, or delete it. Finally, the user can apply the template to other targets, which creates the service tests on the other targets and adds the monitoring beacons.

To create a Monitoring Template, follow the steps given below:

1. Click **Setup** to navigate to the main Setup page in Enterprise Manager.
2. Click the **Monitoring Templates** link in the left panel.
3. Click **Create** to create a monitoring template.
4. In the target selection box, enter or select a service target and click **Continue**.
5. In the Monitoring Template General Page, enter the name of the template that you wish to create.

6. Click **Tests** to add / remove or configure service tests associated with the selected service target. Make the required changes to this page and click **OK** to save the template to the repository.

After you have created the Monitoring Template, use the **Apply** option to apply this template to a service test. You can click **Edit** to modify the template. For more details on these operations, refer to the Online Help.

10.8.1 Configuring Service Tests and Beacons for the Monitoring Template

You can configure the service tests and beacons associated with the template by using the options in the **Tests** page. A service test-based template contains the following elements:

- **Variables:** A variable may occur at multiple locations in the service tests. The Variables table allows you to specify default values for all the variables. These default values will be stored in the template along with the variables. You can specify values other than the default while applying the template to a target. You can perform the following operations:
 - **Add** a variable. The variable can consist of letters, numbers and underscores only.
 - **Rename** a variable. When you rename a variable, all variable references in the service tests will be replaced with the new name.
 - **Remove** variables for properties within service tests. If you remove a non-password variable, all references to the variable in test properties will be replaced with the variable's default value
 - **Replace Text** in test properties with a variable definition.
- **Service Tests:** You can edit the test definition and define variables for various properties. You can select the tests from the original target that are to be part of the template by clicking the **Add / Remove** button. You can specify whether the service test is a key test and if it should be enabled. You can also click **Monitoring Settings** to drill down to this page and define metrics and thresholds for the service tests.
- **Beacons:** Use the **Add / Remove** button to specify which beacons are to be included in the template. You can also specify whether each beacon is a key beacon.

Refer to the Enterprise Manager Online Help for detailed instructions on these operations.

10.9 Configuring Service Levels

A service level rule is defined as an assessment criteria used to determine service quality. It allows you to specify availability and performance criteria that your service must meet during business hours as defined in your Service Level Agreement. For example, e-mail service must be 99.99% available between 8am and 8pm, Monday through Friday.

A service level rule specifies the percentage of time a service meets the performance and availability criteria as defined in the Service Level Rule. By default, a service is expected to meet the specified criteria 85% of the time during defined business hours. You may raise or lower this percentage level according to service expectations. A service level rule is based on the following:

- **Business Hours:** Time range during which the service level should be calculated as specified in your Service Level Agreement.
- **Availability:** Allows you to specify when the service should be considered available. This will only affect the service level calculations and not the actual availability state displayed in the console. You can choose a service to be considered up when it is one or more of the following states:
 - Up: By default the service is considered to be Up or available.
 - Under Blackout: This option allows you to specify service blackout time (planned activity that renders the service as technically unavailable) as available service time.
 - Unknown: This option allows you to specify time that a service is unmonitored because the Management Agent is unavailable be counted as available service time.
- **Performance Criteria:** You can optionally designate poor performance of a service as a Service Level violation. For example, if your Website is up, but it takes 10 seconds to load a single page, your service may be considered unavailable.
- **Business Criteria:** Business criteria are useful in determining in the health of the business processes for a particular service. You can optionally define business metrics that can affect the Service Level. A Service Level violation occurs when a critical alert is generated for a specified business metric.

Note: The **Business Criteria** column is displayed only if one or more key business indicators are associated with the service. Refer to *Oracle Enterprise Manager Integration Guide*.

- **Actual Service Level:** This is calculated as percentage of time during business hours that your service meets the specified availability, performance, and business criteria.
- **Expected Service Level:** Denotes a minimum acceptable service level that your service must meet over any relevant evaluation period.

You can define only one service level rule for each service. The service level rule will be used to evaluate the **Actual Service Level** over a time period and compare it against the **Expected Service Level**.

10.9.1 Defining Service Level Rules

A Service Level Rule is defined as assessment criteria to measure Service quality. A Service Level Rule is based on the following:

- Time range for which the rule is applicable.
- Metrics that define the rule.
- The user expectation on these metrics values

The Expected Service Level is the expected quality for the service and is defined based on the time range and metrics of the Service Level Rule. For example, the Expected Service Level can be that the service is available 99% of the time during business hours.

When you create a service, the default service rule is applied to the service. However, you must edit the service level rule for each service to accurately define the assessment criteria that is appropriate for your service. To define a service level rule:

1. Click the **Targets** tab and **Services** subtab. The Services main page is displayed.
2. Click the service name link to go to the Service Home page.
3. In the Related Links section, click **Edit Service Level Rule**.
4. On the Edit Service Level Rule page, specify the expected service level and the actual service level and click **OK**. The expected service level specifies the percentage of time a service meets the performance, usage, availability, and business criteria defined in the Service Level Rule. The actual service level defines the baseline criteria used to define service quality and includes business hours, availability, performance criteria, usage criteria, and business criteria.

Note: Any Super Administrator, owner of the service, or Enterprise Manager administrator with OPERATOR_TARGET target privileges can define or update the Service Level Rule.

10.9.2 Viewing Service Level Details

You can view service level information directly from the either of the following:

- **Enterprise Manager Cloud Control Console** -From any Service Home page, you can click on the Actual Service Level to drill down to the Service Level Details page. This page displays what Actual Service Level is achieved by the service over the last 24 hours / 7 days / 31 days, compared to the Expected Service Level. In addition, details on service violation and time of each violation are presented in both graphical and textual formats.
- **Information Publisher** - Information Publisher provides an out-of-box report definition called the Services Dashboard that provides a comprehensive view of any service. From the Report Definition page, click on the **Services Monitoring Dashboard** report definition to generate a comprehensive view of an existing service. By default, the availability, performance, status, usage, business, and Service Level of the service are displayed. The Information Publisher also provides service-specific report elements that allow you to create your own custom report definitions. The following report elements are available:
 - **Service Level Details:** Displays **Actual Service Level** achieved over a time-period and violations that affected it.
 - **Service Level Summary:** Displays service level violations that occurred over selected time-period for a set of services.
 - **Services Monitoring Dashboard:** Displays status, performance, usage, business, and service level information for a set of services.
 - **Services Status Summary:** Information on one or more services' current status, performance, usage, business, and component statuses.

Refer to the Online Help for more details on the report elements.

10.10 Configuring a Service Using the Command Line Interface

Using the Command Line Interface, you can define service targets, templates and set up incidents. EMCLI is intended for use by enterprise or system administrators writing scripts (shell/batch file, perl, tcl, php, etc.) that provide workflow in the

customer's business process. EMCLI can also be used by administrators interactively, and directly from an operating system console. Refer to *Enterprise Manager Command Line Interface Guide* for details.

Samples EMCLI templates to create an ATS service test and a Web transaction are shown below.

Example 10–1 ATS Service Test Template

```
<?xml version = '1.0' encoding = 'UTF-8'?>
<transaction-template template_type="generic_service" xmlns="template">
  <variables/>
  <transactions>
    <mgmt_bcn_transaction>
      <mgmt_bcn_txn_with_props>
        <mgmt_bcn_txn is_representative="true" name="ATS Page"
          monitoring="true" txn_type="OATS"/>
        <properties>
          <property name="Collection Interval" num_value="5.0" prop_type="2"
            encrypt="false"/>
          <property name="scriptDescription" string_value="[1] Sign
In&#xA;[2] Welcome&#xA;[3] Single Sign-Off&#xA;[4] Sign In" prop_type="1"
            encrypt="false"/>
          <property name="fileUploadTime" string_value="2012-08-09
08:47:22.0" prop_type="1" encrypt="false"/>
          <property name="OpenScriptJwgName" string_value="ATKHomepage.zip"
            prop_type="1" encrypt="false"/>
          <property name="usageOptions" string_value="userDefined"
            prop_type="1" encrypt="false"/>
          <property name="fileSize" string_value="41368" prop_type="1"
            encrypt="false"/>
          <property name="beaconDistributionOverride"
            string_value="AtsCredentials=1" prop_type="1" encrypt="false"/>
          <property name="FilePropertyValue" prop_type="7" encrypt="false"/>
          <property name="databankFilesJar" prop_type="7" encrypt="false"/>
          <property name="databankFiles"
            string_value="AtsCredentials,AtsCredentials.csv,3;" prop_type="1"
            encrypt="false"/>
          <property name="granularity" string_value="transaction"
            prop_type="1" encrypt="false"/>
          <property name="databankValues" string
            _value="AtsCredentials.&quot;test-id&quot;=&quot;ATKHomepage&quot;;AtsCredentials.
&quot;login-host&quot;=&quot;login-aufsn4x0cxf.oracleoutsourcing.com&quot;;
AtsCredentials.&quot;host&quot;=&quot;fs-aufsn4x0cxf.oracleoutsourcing.com&quot;;
AtsCredentials.&quot;username&quot;=&quot;faadmin&quot;;AtsCredentials.&quot;
password&quot;=&quot;fusionfal&quot;" prop_type="1" encrypt="false"/>
          <property name="modules"
            string_value="oracle.oats.scripting.modules.http;version=2.6.1&#xA;
oracle.oats.scripting.modules.adfload;version=2.6.1&#xA;oracle.oats.scripting.
modules.basic;version=2.6.1" prop_type="1" encrypt="false"/>
          <property name="databankAliasMapping"
            string_value="AtsCredentials=AtsCredentials.csv" prop_type="1" encrypt="false"/>
          <property name="defaultCLIOptions" string_value="-dboptions
*:1:FIRST_RECORD_ONLY -jwg ATKHomepage.jwg -noReport true" prop_type="1"
            encrypt="false"/>
        </properties>
      </mgmt_bcn_txn_with_props>
    </mgmt_bcn_transaction>
  </transactions>
</transaction-template>
```

```

        <mgmt_bcn_step step_number="1" name="[1] Sign In"
step_type="OATS"/>
        <properties>
            <property name="url" string
_value="aHR0cHM6Ly9mcy1hdWZzbjR4MGRwZi5vcnFjbGVvdXRzb3VyY2luZy5jb20vaG9tZVBh
Z2UvZmFjZXMvQXRrSG9tZVBhZ2VXZWxjb211" prop_type="1" encrypt="false"/>
        </properties>
    </mgmt_bcn_step_with_props>
    <mgmt_bcn_step_with_props>
        <mgmt_bcn_step step_number="2" name="[2] Welcome"
step_type="OATS"/>
        <properties>
            <property name="url" string
_value="aHR0cHM6Ly9mcy1hdWZzbjR4MGRwZi5vcnFjbGVvdXRzb3VyY2luZy5jb20vaG9
tZVBhZ2UvZmFjZXMvQXRrSG9tZVBhZ2VXZWxjb211" prop_type="1" encrypt="false"/>
        </properties>
    </mgmt_bcn_step_with_props>
    <mgmt_bcn_step_with_props>
        <mgmt_bcn_step step_number="3" name="[3] Single Sign-Off"
step_type="OATS"/>
        <properties>
            <property name="url" string
_value="aHR0cHM6Ly9sb2dpbi1hdWZzbjR4MGRwZi5vcnFjbGVvdXRzb3VyY2luZy5jb20vb
2FtL3NlcnZlci9sb2dvdXQ=" prop_type="1" encrypt="false"/>
        </properties>
    </mgmt_bcn_step_with_props>
    <mgmt_bcn_step_with_props>
        <mgmt_bcn_step step_number="4" name="[4] Sign In"
step_type="OATS"/>
        <properties>
            <property name="url" string
_value="aHR0cHM6Ly9mcy1hdWZzbjR4MGRwZi5vcnFjbGVvdXRzb3VyY2luZy5jb20vaG
9tZVBhZ2UvZmFjZXMvQXRrSG9tZVBhZ2VXZWxjb211" prop_type="1" encrypt="false"/>
        </properties>
    </mgmt_bcn_step_with_props>
</steps_defn_with_props>
<stepgroups_defn/>
<txn_thresholds>
    <mgmt_bcn_threshold warning_threshold="0.0" warning_operator="1"
critical_threshold="0.0" critical_operator="1" num_occurrences="1">
        <mgmt_bcn_threshold_key metric_name="openscript_response"
metric_column="status"/>
    </mgmt_bcn_threshold>
</txn_thresholds>
<step_thresholds/>
<stepgroup_thresholds/>
</mgmt_bcn_transaction>
</transactions>
</transaction-template>

```

Example 10–2 Web Transaction Service Test Template

```

<?xml version = '1.0' encoding = 'UTF-8'?> <transaction-template
template_type="generic_service" xmlns="template">
    <variables>
        <variable name="HOST1" value="smpperf-linux26.us.oracle.com"/>
        <variable name="PORT1" value="5416"/>
        <variable name="PROTOCOL1" value="https"/>
    </variables>
    <transactions>

```

```

    <mgmt_bcn_transaction>
      <mgmt_bcn_txn_with_props>
        <mgmt_bcn_txn description="Test for checking the availability of
EM Console/Website" is_representative="true" name="EM Console Service Test"
monitoring="true" txn_type="HTTP" />
        <properties>
          <property name="readTimeout" num_value="120000.0" prop_type="2"
encrypt="false" />
          <property name="Collection Interval" num_value="5.0" prop_type="2"
encrypt="false" />
          <property name="certValidationMode" string_value="1" prop_type="1"
encrypt="false" />
          <property name="maxDownloadSize" num_value="1.0E8" prop_type="2"
encrypt="false" />
          <property name="sensitiveValuesProtection" string_value="0"
prop_type="1" encrypt="false" />
          <property name="failureStringModes" string_value="regularText"
prop_type="1" encrypt="false" />
          <property name="UserAgent" string_value="Mozilla/4.0 (compatible;
MSIE 6.0; Windows NT 5.1) OracleEMAgentURLTiming/3.0" prop_type="1"
encrypt="false" />
          <property name="successStringModes" string_value="regularText"
prop_type="1" encrypt="false" />
          <property name="variablesModes" string_value="urlEncode"
prop_type="1" encrypt="false" />
          <property name="content" string_value="0" prop_type="1"
encrypt="false" />
          <property name="AcceptLanguage" string_value="en"
prop_type="1" encrypt="false" />
          <property name="connectionTimeout" num_value="120000.0"
prop_type="2" encrypt="false" />
          <property name="useCache" string_value="yes" prop_type="1"
encrypt="false" />
          <property name="stringValidationMode" string_value="1"
prop_type="1" encrypt="false" />
          <property name="granularity" string_value="transaction"
prop_type="1" encrypt="false" />
          <property name="numThreads" num_value="4.0" prop_type="2"
encrypt="false" />
          <property name="retries" num_value="1.0" prop_type="2"
encrypt="false" />
          <property name="timeout" num_value="300000.0" prop_type="2"
encrypt="false" />
          <property name="retryInterval" num_value="5000.0" prop_type="2"
encrypt="false" />
        </properties>
      </mgmt_bcn_txn_with_props>
    </mgmt_bcn_transaction>
    <steps_defn_with_props>
      <mgmt_bcn_step_with_props>
        <mgmt_bcn_step step_number="1" name="1.Access Logout page"
step_type="HTTP" />
        <properties>
          <property name="req_mode" num_value="1.0" prop_type="2"
encrypt="false" />
          <property name="http_method" string_value="G" prop_type="1"
encrypt="false" />
          <property name="url"
string_value="{PROTOCOL1}:{HOST1}:{PORT1}/em/console/logon/logoff?event=load"
prop_type="1" encrypt="false" />

```

```

        </properties>
    </mgmt_bcn_step_with_props>
</steps_defn_with_props>
<stepgroups_defn/>
<txn_thresholds>
    <mgmt_bcn_threshold warning_threshold="6000.0" warning_operator="0"
critical_threshold="12000.0" critical_operator="0" num_occurrences="1">
        <mgmt_bcn_threshold_key metric_name="http_response"
metric_column="avg_response_time"/>
    </mgmt_bcn_threshold>
    <mgmt_bcn_threshold warning_threshold="0.0" warning_operator="1"
critical_threshold="0.0" critical_operator="1" num_occurrences="1">
        <mgmt_bcn_threshold_key metric_name="http_response"
metric_column="status"/>
    </mgmt_bcn_threshold>
</txn_thresholds>
<step_thresholds/>
<stepgroup_thresholds/>
</mgmt_bcn_transaction>
</transactions>
</transaction-template>

```

10.11 Troubleshooting Service Tests

This section lists some of the common errors you may encounter while using the Forms and the Web Transaction test type. The following topics are covered here:

- [Verifying and Troubleshooting Forms Transactions](#)
- [Verifying and Troubleshooting Web Transactions](#)

10.11.1 Verifying and Troubleshooting Forms Transactions

The section covers the following:

- [Troubleshooting Forms Transaction Playback](#)
- [Troubleshooting Forms Transaction Recording](#)

10.11.1.1 Troubleshooting Forms Transaction Playback

This section lists some of the common errors you may encounter while playing back a Forms transaction and provides suggestions to resolve these errors.

1. **Error Message:** Connection to Forms Server is lost. Possible version mismatch between agentjars and formsjars.

Possible Cause: The transaction was recorded using an out-of-the-box Forms version.

Solution: Verify the version of the Forms Application that you are running by checking the version number in the About Oracle Forms Online Help. If this version is not supported, follow the steps listed under Error Message 2.

2. **Error Message:** Version Not Supported <forms_version>

Possible Cause: The machine on which the beacon has been installed does not contain the necessary forms jar files.

Solution: To resolve this error, follow these steps:

1. Login to the system on which the Forms server has been installed. Locate the `frmall.jar` (if you are using Forms 10.1 or later) or `f90all.jar` (if are using Forms 9.0.4 or later) under the `$FORMS_HOME/forms/java` directory.
2. Login to the system on which the beacon has been deployed and copy the jar file to the `$ORACLE_HOME/jlib/forms/<version>/` directory. The version you specify here should be the same as the version string in the error message. Make sure that the directory is empty before you copy over the jar file.

If you are using Oracle Applications R12 and you encounter this error, follow these steps to resolve the error:

1. Login to the system in which the Oracle Application server has been deployed. Locate the following files:

`$JAVA_TOP/oracle/apps/fnd/jar/fndforms.jar`
`$JAVA_TOP/oracle/apps/fnd/jar/fndewt.jar`
2. Login to the system on which the beacon has been deployed and copy these files to the `$ORACLE_HOME/jlib/forms/apps/` directory. Make sure that the directory is empty before you copy over the jar files.

Note: You cannot monitor two deployments of Oracle Applications from the same beacon if different versions of Oracle Applications have been used.

3. **Error Message:** Forms URL is not pointing to the forms servlet.

Possible Cause: When the Forms transaction was recorded, the location of the forms servlet could not be determined.

Solution: Make sure that the Forms URL Parameter is pointing to the forms servlet. It should be `http://<hostname>:<port>/forms/frmservlet` for Forms10g or `http://<hostname>:<port>/forms/f90servlet` for Forms 9i. This parameter is automatically set by the Forms Transaction Recorder. But if it has not been set, you can locate the URL by following these steps:

- Launch the Forms application.
- View the source HTML file in the Forms launcher window.
- Locate the `xsurl` variable. The URL is stored in this variable.

4. **Error Message:** Could not connect to <machine name>.

Possible Cause: The machine on which the beacon has been installed cannot access the Forms Application.

Solution: Make sure the machine on which the beacon has been installed can access the Forms Application and firewalls have been properly configured. Support for playing back Forms transactions through proxy server is not available in this release.

5. **Error Message:** Invalid module path in the initial message.

Possible Cause: The transaction may have been incorrectly recorded or may be corrupt.

Solution: Try to record the transaction again.

6. **Error Message:** Cannot connect to login server.

Possible Cause: This error may occur due the following reasons:

- The Login URL that you have specified may be incorrect.
- An invalid HTTPS certificate may have been provided for the login server.

Solution:

- Verify that the Login URL is correct.
- If you are using HTTPS to connect to login server, make sure the certificate on the server is written for the login server machine itself. Make sure the SSL Certificate is imported into Agent and the CN of the certificate matches the host name of the login Server URL.

10.11.1.2 Troubleshooting Forms Transaction Recording

This section lists some troubleshooting steps that you can use when the Forms transaction cannot be recorded successfully.

1. Make sure that all your Internet Explorer instances are closed and no java runtime programs are open.
2. Start recording again with the java console open. You can view any exceptions or error messages displayed on the console.
3. You should now see the text "Forms Transaction Recorder Version: <version number>" on the console. If this text is displayed, proceed to step 5. If you do not see the text, check if the `formsRecorder.jar` has been copied to the Forms archive directory. You can perform this check using either of the following methods:
 1. Navigate to the Forms archive directory and check if the `formsRecorder.jar` file is present in the directory.
 2. Navigate to the **Enable Forms Transaction Monitoring** page, select the corresponding Forms server target and click **Configure**. Enter the host credentials to see if the Forms Transaction Recorder has already been configured on this Forms server. If the `formsRecorder.jar` is not present in the Forms archive directory, you need to configure your Forms server for transaction monitoring. After ensuring that the `formsRecorder.jar` is present in the archive directory of the Forms server, go back to **Step 1** and try recording again.
4. If you see an exception related to the `java.policy` file displayed on the java console, check the file to ensure that it has the required content and is in the right location. If any errors are found, you must fix these errors and try recording again.
5. If the recording still fails, check if the Enterprise Manager Certificate has been imported to the secure site. If the certificate has not been imported, you must import it and try recording again.

10.11.2 Verifying and Troubleshooting Web Transactions

This section lists some of the common errors you may encounter while recording and playing back Web Transactions.

1. **Scenario:** Verify Service Test displays: `Connection establishment timed out -- http://....`

Possible Cause: The beacon can only access that URL via a proxy server and it has not been configured.

Solution: From the All Targets page, select the beacon, click **Configure** and set the beacon proxy setting.

2. **Scenario:** Verify Service Test displays: Authorization Required --
https://...../

Possible Cause: The Basic Authentication information is not recorded automatically.

Solution: To resolve this error, follow these steps:

1. From the Service Tests and Beacons page, select the service test, click Edit.
2. Make sure you enter all the Basic Authentication information: Username, Password, and Realm.

Note: Realm usually appears above the Username label in the Browser's authorization dialog box.

3. **Scenario:** Verify Service Test displays
sun.security.validator.ValidatorException:No trusted
certificate found -- https://...../.

Possible Cause: The beacon does not know about this SSL Certificate.

Possible Solution: From the Service Tests and Beacons page, select the service test, and click **Edit**. Under **Advanced Properties**, and set **Authenticate SSL Certificates** to **No**.

4. **Scenario:** Verify Service Test displays: Timeout of 300000 exceeded for
https://...../ Response time = 3000000

Possible Cause: The test may be too complex to complete within the allotted time. Or, this may be an actual performance issue with the server.

Possible Solution: From the Service Tests and Beacons page, select the service test, and click **Edit**. If this is not a server performance issue, under **Advanced Properties**, increase the **Timeout Value**.

5. **Scenario:** The Verify Service Test option reports that the service as down, but the Web application is up and you can successfully play back the Web transaction.

Possible Cause: The Web application is only compatible with Internet Explorer or Mozilla-based browsers.

Possible Solution: From the Service Tests and Beacons page, select the service test, and click **Edit**. Under **Advanced Properties**, set the **User Agent Header** as
Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1)
OracleEAgentURLTiming/3.0.

Note: For Enterprise Manager 10.2.0.4 and beyond, this User Agent Header is set automatically during Web transaction recording.

6. **Scenario:** Test Performance Page does not show any step metrics.

Possible Cause: By default, only transaction-level metrics are collected.

Possible Solution: From the Service Tests and Beacons page, select the service test, click **Edit**, and set **Data Granularity** to **Step**.

Configuring a 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](#)
- [What's New in Software Library](#)
- [Performing Software Library Tasks Using EM CLI Verbs or in Graphical Mode](#)
- [Software Library Storage](#)
- [Prerequisites for Configuring Software Library](#)
- [Configuring Software Library Storage Location](#)
- [Configuring Software Library on a Multi-OMS System](#)
- [Using Software Library Entities](#)
- [Tasks Performed Using the Software Library Home Page](#)
- [Maintaining Software Library](#)

11.1 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 11-1](#), 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 11–1 Software Library Console

Software Library Page Refreshed Aug 11, 2011 7:07:32 AM PDT

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

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

Name	Type	Subtype	Revision	Status	Maturity	Owner	Description
Software Library						ORACLE	Root Folder for Software Library entities
Application Server Provisioning Utilities						ORACLE	Entities belonging to AS Provisioning
Bare Metal Provisioning						ORACLE	Bare Metal Provisioning directory
BPEL Provisioning						ORACLE	BPEL Provisioning Entities
Cloud						ORACLE	Cloud
Coherence Node Provisioning						ORACLE	Coherence Node Provisioning Entities
Common Provisioning Utilities						ORACLE	Directives belonging to Common Provisioning (SIDB and RACPRO)
Components						SYSMAN	Components Folder
Directives						SYSMAN	Directives Folder
Images						SYSMAN	Images Folder
Networks						SYSMAN	Networks Folder
Suites						SYSMAN	Suites Folder
CompositeDeploy						ORACLE	CompositeDeploy Entities
CVU Prerequisite-fixup components						ORACLE	CVU Prerequisite-fixup components belonging to DB Provisioning
DB Provisioning						ORACLE	Directives and Components belonging to DB Provisioning
Fusion Middleware Provisioning Utilities						ORACLE	Directives belonging to FMW Provisioning
Java EE Provisioning						ORACLE	Java EE Application Provisioning Entities
MultiOMS						ORACLE	List of Oracle shipped Directives
Oracle VM Server Provisioning						ORACLE	Oracle VM Server Provisioning directory
OSB Provisioning						ORACLE	OSB Provisioning Entities
Patching						ORACLE	Patching directory
Prerequisite-fixup components						ORACLE	Prerequisite-fixup components Components belonging to DB Prov
Soa Provisioning						ORACLE	SOA Provisioning Entities

The Software Library Page facilitates storage of 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 supports patching and provisioning in Online mode and Offline mode. For example, if database patches cannot be downloaded directly from *My Oracle Support*, you can download them separately, and stage them from 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 11.7, "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.

11.2 Users, Roles, and Privileges

By default, all the Software Library folders and entities that ship with the product 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. Administrators by default do not have any Software Library privileges, it is for the Super Administrator to grant access, privileges to an Administrator.

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

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

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

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 deleting entities, changing the maturity status, 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 until the Super Administrator or the owner of the entity grants them at least a view privilege on the entity.

Note: All the folders and entities that ship with the product also known 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 11-1](#) 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 11–1 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
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 11–2 describes all the primary users of Software Library, and their associated privileges:

Table 11–2 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 11–1 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.

11.3 What's New in Software Library

For Enterprise Manager 12c Software Library enhancements, refer to *Enterprise Manager Cloud Control Introduction Guide*.

11.4 Performing Software Library Tasks Using EM CLI Verbs or in Graphical Mode

Starting with Oracle Enterprise Manager Cloud Control 12c Release 2 (12.1.0.2), command line utility has been introduced for Software Library users in Oracle Enterprise Manager Cloud Control that enables you to perform some of the console-based Software Library operations using the text-based consoles.

The following table describes both approaches to perform some of the Software Library tasks:

- Enterprise Manager Command Line Interface (EM CLI)
- Enterprise Manager Graphical User Interface (EM GUI)

Note: For more information about the syntax and usage of the EM CLI verbs described in [Table 11-3](#), along with workflow examples, refer to *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

Table 11-3 Software Library EMCLI Verbs

Description	Approach A: Using EM CLI Verb	Approach B: Using Enterprise Manager Cloud Control Console
Adding a Software Library storage location	<code>add_swlib_storage_location</code>	Configuring an OMS Shared Filesystem Location Configuring an OMS Agent Filesystem Location Configuring a Referenced File Location
Creating a Software Library entity	<code>create_swlib_entity</code>	Creating Generic Components Creating Directives
Creating a Software Library folder	<code>create_swlib_folder</code>	Organizing Entities
Listing the Software Library entities	<code>list_swlib_entities</code>	Accessing Software Library Home Page Searching Entities
Listing Software Library entity types	<code>list_swlib_entity_types</code>	NA
Listing Software Library entity subtypes	<code>list_swlib_entity_subtypes</code>	NA
Listing Software Library folders	<code>list_swlib_folders</code>	NA

Table 11–3 (Cont.) Software Library EMCLI Verbs

Description	Approach A: Using EM CLI Verb	Approach B: Using Enterprise Manager Cloud Control Console
List Software Library storage locations	list_swlib_storage_locations	Accessing Software Library Administration Page
Referring files from a Software Library entity	refer_swlib_entity_files	Creating Entities Viewing, Editing, and Deleting Entities
Re-Importing Software Library metadata	reimport_swlib_metadata	Re-Importing Oracle Owned Entity Files
Removing a Software Library storage location	remove_swlib_storage_location	Removing (and Migrating) Software Library Storage Location
Modifying a Software Library entity	update_swlib_entity	Viewing, Editing, and Deleting Entities
Uploading files to a Software Library entity	upload_swlib_entity_files	Creating Entities Viewing, Editing, and Deleting Entities
Modify a Software Library OMS Agent storage location to change the associated OMS Host and the credential for accessing the location.	switch_swlib_oms_agent_storage	NA
Verify and report Software Library State	verify_swlib	NA

Note: For more information on the usage of EM CLI verbs, refer to *Oracle Enterprise Manager Command Line Interface Guide*.

11.5 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 File System, or OMS Agent File System) on the host where the OMS is running. A storage location in Software Library represents a repository of files that are either uploaded to Software Library or referenced by it.

Note: If you choose to newly configure an OMS Shared Storage Location, then ensure that the file system path that you specify for the location is either a shared path or a mounted path. By doing so, the newly configured location can be made accessible in a multiple OMS environment in the future. If the new location is being added in a multiple OMS environment, then the file system path should be accessible from all the OMS hosts.

However, if you have configured the OMS Shared Storage Location on a local file system, then perform the steps listed in the [Section 11.8](#) to migrate this location to another OMS Shared Storage Location that has a shared or mounted path.

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 11–2 Software Library Administration

Software Library: Administration Page Refreshed Aug 11, 2011 7:11:34 AM PDT

Software Library > Software Library: Administration

The administration console allows for configuring and administering Software Library storage locations.

Upload File Locations Referenced File Locations

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

Storage Type: OMS Shared Filesystem

Configure filesystem locations on OMS Host(s). These locations must be locally accessible by all the OMS instances, typically a mounted/shared location. You can optionally configure the common credential to be used by Software Library for reading/writing from/to a location.

Actions: View + Add... ✎ Edit... ✕ Migrate and Remove

Name	Status	Location	Associated Entities	Total Space	Available Space	Last Refreshed
Testing	Active	/scratch/nbhaktha/swlib/	Show	96,462 GB	61,662 GB	Thu Aug 11 07:11:34 PDT 2011

The Software Library Administration Page as shown in [Figure 11–2](#) 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 12c, 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, and 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: For production environments, Oracle recommends allocating a minimum 100GB of storage for your software library. Also, ensure that this storage can easily be extended in future, if it starts running out of space.

Once the storage location starts running out of space, 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 11.11](#)

The following types of storage locations are available:

- [Upload File Locations](#)
- [Referenced File Location](#)

11.5.1 Upload File Locations

Upload File Locations are locations configured for storing files uploaded by Software Library as part of creating or updating an entity.

For Software Library to become usable, at least one upload file location must be configured. On adding the first upload file location, a job is submitted to import the Software Library metadata from the Oracle home of each of the installed Enterprise Manager plug-in. Ensure that you wait for this job to complete successfully, before performing other patching or provisioning operations.

Note: To physically delete a file system configured as an Upload storage location with Software Library, you must ensure that you first configure an alternate storage location where you can migrate the existing contents (entities). If you fail to perform this migration, then the entities dependent on the files from this location will be rendered unusable. For more information about deleting a storage location, and migrating the contents, see [Section 11.11.3](#).

Prerequisites

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.

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 an OMS that will be provisioned using the Add Management Service 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 as follows:

OMS Shared File System (*Recommended Storage Option*)

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.

Note: Oracle recommends using OMS Shared File System option for storing files uploaded to Software Library.

In Windows, for sharing, the OCFS2 cluster file system is recommended. However, if you can not set up a shared file system because of some constraints, then you may use the OMS Agent File System. For more information see "[OMS Agent File System](#)".

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. To use OMS Agent Filesystem 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.

Note: If you can not set up an OMS Shared File System for storage because of some constraints, then you may use the OMS Agent File System.

11.5.2 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.

Note: To use entities referring files of a location, you must have view privilege on the credentials associated with the locations.

11.6 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.
- 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.
- 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.

11.7 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](#)

Note: Starting with Oracle Enterprise Manager Cloud Control 12c Release 2 (12.1.0.2), in addition to using the GUI as described in this section, you can alternatively use the command line interface tool to Configure the Software Library. To do so, refer to *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

11.7.1 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, click **Show Detailed Results**. Typically, the name of the job starts with SWLIBREGISTERMETADATA_.*.

If the Import job fails, see [Section 11.11](#) 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.

11.7.2 Configuring an OMS Agent Filesystem Location

Note: The OMS Agent File system must be set up only when the recommended storage option, which is the OMS Shared File System cannot be setup because of some constraints. For more information, see [Section 11.5.1](#).

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:

Add OMS Agent Filesystem Location

Provide a name, a host and file system location on the host. The host can only be one of the hosts where the Oracle Management Server instance is running. The credential used for selecting the file system location on host will be saved along with this location. Ensure that this credential is accessible to all Software Library storage administrators and designers.

* Name: AgilLoc1
 Host: stnss.xyz.com
 Location: /vol/stnss_swlib

OK Cancel

- 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.

Note: For a user to access and leverage an OMS Agent Filesystem upload location successfully, the owner of the Named Credential (basically, the credential used to connect to the host machine), must grant a View Privilege on the credential chosen to all the Administrators (or users) accessing this OMS Agent Filesystem location.

For more information about granting privileges on a Named Credential, refer to *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

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.

11.7.3 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**.



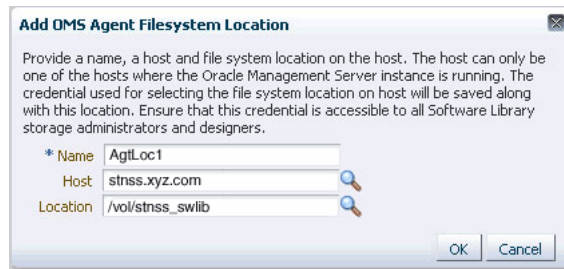
In the Add NFS Location dialog box, do the following:

- 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**.

Note: While creating a procedure, if you have a component step or a directive step that refers to an NFS file location, then you must ensure that you set the preferred privileged credentials for the target host before the procedure is submitted for execution.

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.

11.8 Configuring Software Library on a Multi-OMS System

Oracle recommends that you configure each OMS Shared Storage Location to use a shared or mounted file system path. Doing this will ensure that this newly configured location remains accessible from any OMS host as and when they are added. All upload and stage requests for the files will happen through the Management Agent monitoring the OMS host.

Note: Starting with Enterprise Manager 12c, you can use the EM CLI utility for migrating files across upload locations of different storage types. Additionally, the EM CLI verb (`remove_swlib_storage_location`) also supports migration from an OMS Shared storage location to an OMS Agent storage location, and vice versa. This can be performed only through the command line using EM CLI utility, you can not perform the same task using the Cloud Control UI.

If however, you have configured the OMS Shared storage location to use a local file system path, then you must migrate it to another OMS Shared Storage Location that uses a shared or mounted path. 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, add a new OMS Shared storage location by specifying a name (for example: NewShared), and a shared file system path.

3. On successful completion, select the location you want to migrate, (For example: OldNonShared), and click **Migrate and Remove**.
4. In the popup dialog box, select the new OMS Shared storage location (NewShared) as the destination location for the migration of files.
5. Click **Remove** to submit a job, which on successful completion deletes the storage location entry from the table.

11.9 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:

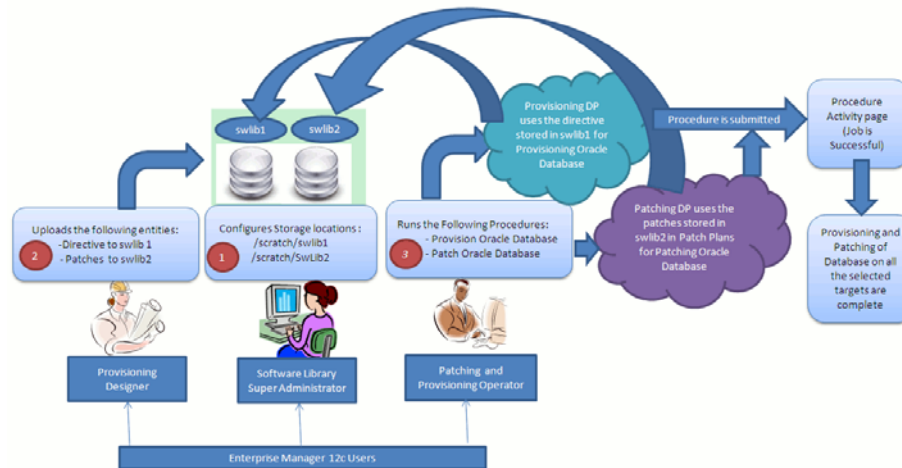
Types	Description
Oracle-owned Entities	These entities are available by default on the Software Library Home page, once the Software Library is configured. In the following graphic, all the entities that are owned by Oracle , qualify as Oracle-owned entities, and all the folders that appear with a lock 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. For example, in the following graphic you can see a custom folder called My Entities, and entities called os2 and os1 created by the owner of the entity. These entities are called User-owned entities.

Name	Type	Subtype	Revision	Status	Maturity	Owner	Description
Software Library						ORACLE	Root Folder for Software Library entities
Application Server F						ORACLE	Entities belonging to AS Provisioning
Bare Metal Provisionir						ORACLE	Bare Metal Provisioning directory
BPELProvisioning						ORACLE	BPEL Provisioning Entities
Cloud						ORACLE	Cloud
Coherence Node Pr						ORACLE	Coherence Node Provisioning Entities
Common Provisionir						ORACLE	Directives belonging to Common Provisionir
Components						SYSMAN	Components Folder
Directives						SYSMAN	Directives Folder
Images						SYSMAN	Images Folder
Networks						SYSMAN	Networks Folder
Suites						SYSMAN	Suites Folder

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 11-3 Using Software Library Entities for Provisioning and Patching Tasks



11.10 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

11.10.1 Organizing Entities

Only designers who have the privilege to create any Software Library entity, can create folders.

Note: Starting with Oracle Enterprise Manager Cloud Control 12c Release 2 (12.1.0.2), in addition to using the GUI as described in this section, you can alternatively use the command line interface tool to Create Folders. To do so, refer to *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

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.

11.10.2 Creating Entities

From the Software Library Home page, you can create the following entities:

- [Creating Generic Components](#)
- [Creating Directives](#)

Note: Starting with Oracle Enterprise Manager Cloud Control 12c Release 2 (12.1.0.2), in addition to using the GUI as described in this section, you can alternatively use the command line interface tool to Create Entities. To do so, refer to *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

11.10.2.1 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 11.10.1](#).

3. From the **Actions** menu, select **Create Entity** and click **Component**. Alternately, right click the custom folder, and from the menu, 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 because there could be an entity with the same name in the folder that is not visible to you, as you do not have view privilege on it.

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

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

6. On the Configure page, 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 either 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 11.7.3](#)

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.

11.10.2.2 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 11.10.1](#).

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 because there could be an entity with the same name in the folder that is not visible to you, as you do not have view privilege on it.

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

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

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.

For example, if the command line argument you want to pass is as follows:

```
./test.sh -user={username}
```

Then,

Argument Prefix is: `-user`

Property Name is: `username`

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 either 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 11.7.3](#)

7. On the Review page, review all the details, and click **Finish** to create the component and upload it on the Software Library.

11.10.3 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 11.10.4.10](#).

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 11.10.4.8](#).

For more information on Oracle Owned Entities and User Owned Entities, see [Section 11.9](#).

11.10.4 Managing Entities

From the Software Library Home page you can perform the following maintenance tasks on the existing entities:

- [Accessing Software Library Home Page](#)
- [Accessing Software Library Administration Page](#)
- [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](#)
- [Staging Files Associated With an Entity](#)

Note: Starting with Oracle Enterprise Manager Cloud Control 12c Release 2 (12.1.0.2), you can either use the GUI or use the command line interface tool for performing all the tasks listed in [Table 11-3](#).

11.10.4.1 Accessing Software Library Home Page

To access the Software Library Home page, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.

11.10.4.2 Accessing Software Library Administration Page

To access the Software Library Administration page, from the **Setup** menu, select **Provisioning and Patching**, then select **Software Library**.

11.10.4.3 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.

11.10.4.4 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 11.10.4.10](#).

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.

11.10.4.5 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 11.10.4.10](#).

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.

11.10.4.6 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 11.10.4.10](#).

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.

11.10.4.7 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 11.10.4.10](#).

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.

11.10.4.8 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 11.10.4.10](#).

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 11.10.4.9](#).

11.10.4.9 Purging Deleted Entities

To purge the deleted entities from all the configured Agent Storage locations, you can run a purge job. To do so, follow these steps:

1. In Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. On the Software Library home page, from **Actions** menu, select **Deleted Entities**. A list of entities that are deleted from Software Library are displayed.

Note: The **Space Used** attribute is displayed only for the deleted entities that had uploaded files to Software Library.

3. On the Deleted Entities page, click **Purge** to permanently remove these entities from Oracle Management Repository, and the associated files from upload storage locations.
4. A Confirmation Message dialog box is displayed. Click **Job Details** to view the status of the purge job submitted.

Note: A periodic job named SWLIBPURGE runs daily to purge the deleted entities from the Software Library.

11.10.4.10 Searching Entities

This section contains the following topics:

- [Performing Basic and Advanced Search](#)
- [Saving Searches](#)
- [Retrieving Saved Searches](#)
- [Managing Saved Searches](#)

Performing Basic and Advanced Search

To perform a basic or an advanced search for an entity, follow these 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 **Type**, **Name**, **Description**, **Revision**, **Maturity**, **Status**, and **File Name** 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**.

Saving Searches

Optionally, search criteria on the Advanced Search screen of the console, can be saved. Saved searches can be retrieved and executed again. They can also be edited and deleted.

1. Search for entities using the steps listed in [Section 11.10.4.10](#).
2. Click **Save Search**.
3. Enter the preferred name for the search in the text box, and click **Ok**.

Retrieving Saved Searches

To retrieve saved searches, follow these steps:

1. Search for entities using the steps listed in [Section 11.10.4.10](#).
2. Click **Saved Searches**, and select the preferred saved search from the list.

Alternatively, you can also select the preferred saved search from the Favorites menu. To do so, from the **Favorites** menu, select **Saved Software Library Searches**, and select the preferred saved search.

Managing Saved Searches

Using the Manage Saved Searches option, you can edit the name of the saved search, or delete the saved search. To do so, follow these steps:

- To manage saved searches, you can perform one of the following steps:
 - From the **Favorites** menu, select **Manage Favorites**.
 - Click **Saved Searches**, and select **Manage Saved Searches**.
- To edit the name of the saved search, select the preferred saved search, and in the **Name** text field, enter the new name. Click **Ok** to save changes.
- To delete or remove a saved search, select the preferred saved search, and click **Remove Selected**. Click **Ok** to save changes.

11.10.4.11 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 11.10.4.12](#).

 - 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 12c can not be imported into Enterprise Manager 12c.

11.10.4.12 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 12c can not be imported into Enterprise Manager 12c.

11.10.4.13 Staging Files Associated With an Entity

For transferring the files associated with an entity to a specific target host follow the steps outlined in this section.

Prerequisites

Ensure that you meet the following prerequisites before staging the files associated with an entity:

1. Only hosts that are monitored by the Enterprise Manager can be specified as the destination for staging the files associated with an entity.
2. Only files that have been successfully uploaded to the entity (hence, in *Ready* status) can be selected for staging.

Note: To verify if the entity has any files in the **Ready** state, follow these steps:

1. Select the entity, and click **View**.
 2. On the View Entity page, select **Select Files** tab, to verify the files associated with the entity.
 3. Unless there is at least one file with a *Ready* status, you cannot proceed with the staging process.
-

3. Only users with Create Job Privileges can stage the files associated with an entity.
4. If the source files to be staged are on NFS, then the credentials used for browsing the destination target should have `root` permissions to be able to mount the NFS location.

Staging Procedure

Log in to Enterprise Manager Cloud Control with designer privileges, and perform the following steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. Select the entity, and from **Actions** menu, select **Stage Entity**.
3. On the Stage Entity page, the entity details, the source upload location, and the selected entity files are displayed.

You can optionally select the **Overwrite files on the staging location** to overwrite an existing version of the same file. If not, ignore this option and proceed.

4. In the Staging Destination section, select the destination host target (Management Agent) that is reporting to the OMS, where you want to stage these entity files.

To select a location on the Management Agent target to host the files, you must log in to the host target and select a location to host the files.

5. Click **Submit**.
6. To verify the status of the submitted job, from **Enterprise** menu, select **Job**, then click **Activity**.

11.11 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](#)
- [Removing \(and Migrating\) Software Library Storage Location](#)
- [Scheduling Purge Job](#)
- [Backing Up Software Library](#)

11.11.1 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. To do so, on the Administration page, in the upload file locations tab, select the storage location. From the **Actions** menu, select **Refresh**. On successful refresh, a confirmation is displayed. Alternately, you can search for the periodic refresh job `SWLIBREFRESHLOCSTATS`, and edit the schedule and other attributes to suit your requirements. By default, this job is scheduled to run every 6 hours.
- Purge deleted entities to conserve disk space. To do so, see [Section 11.11.4](#). Alternatively, you can search for the periodic purge job `SWLIBPURGE`, and edit the schedule and attributes to suit your requirements. By default, this job is scheduled to run every 24 hours.
- Check accessibility of the configured Software Library locations. To do so, on the Administration page, in the upload file locations tab, select the storage location. From the **Actions** menu, select **Check Accessibility**.

11.11.2 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.

11.11.3 Removing (and Migrating) 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**.

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 Shared Filesystem (upload) storage location to OMS Agent storage location using Enterprise Manager UI.

For migrating files across upload locations of different storage types, you must use the Enterprise Manager Command Line option. For example, you can migrate an OMS Shared Filesystem storage location to an OMS Agent Filesystem storage location. Even the reverse operation is supported using EM CLI. However, note that this type of migration, across storage types, is supported specifically for "upload" storage types, and is not applicable for the reference storage types.

To migrate the files from one upload location to another, you can use the EM CLI verb `emcli remove_swlib_storage_location`. For more information about this command, see [Section 11.4](#).

3. In the confirmation dialog box, click **Remove** to submit a job, which on successful completion deletes the storage entry from the table.

Note: When one storage location is migrated to another location, for example, from `/vol/swlib1` to `/vol/swlib2`, the file system contents of the source location (`/vol/swlib1`) are not deleted during the migration. However, going forward, the source location and the files are never referenced by Software Library.

11.11.4 Scheduling Purge Job

Starting with Enterprise Manager 12c 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 all the details and click **OK** to submit the job, on successful completion of the job all the deleted entities are removed from the storage location.

11.11.5 Backing Up Software Library

For information about backing up your Software Library, see [Chapter 29](#).

Introducing Enterprise Manager Support for SNMP

This chapter provides a brief overview of Enterprise Manager support for SNMP. It includes the following sections:

- [Benefits of SNMP Support](#)
- [About the SNMP Management Station](#)
- [How Enterprise Manager Supports SNMP](#)
- [Sending SNMP Trap Notifications](#)
- [Monitoring External Devices Using SNMP](#)
- [About the Management Information Base \(MIB\)](#)
- [About Metric Extensions](#)

The Simple Network Management Protocol (SNMP) is a protocol used for managing or monitoring devices, where many of these devices are network-type devices such as routers, switches, and so on. SNMP enables a single application to first retrieve information, then push new information between a wide range of systems independent of the underlying hardware.

Designed primarily for database, network, and system administrators, SNMP support integrates Enterprise Manager into a number of existing, widely-used management systems. Also, Enterprise Manager can extend its monitoring scope to devices that can be monitored using SNMP.

12.1 Benefits of SNMP Support

The primary benefits of SNMP support include the following:

- The monitoring of key Oracle products is quickly integrated into any management framework based upon SNMP.
- These Oracle products are located, identified, and monitored in real time across enterprise networks of any size.
- Administrators see standard Oracle icons that represent Oracle products in a network map. You can dynamically customize this map.
- Administrators see the current status of Oracle products, as shown by several status variables that are defined for each product in a management information base (MIB), or they can select which elements to view by their status.

- Administrators can anticipate exceptional conditions by defining thresholds and alerts, to respond to special situations as soon as they occur or to enable automatic responses.
- Administrators can store and analyze historical data that has been obtained through SNMP.
- Providers of management applications can easily build customized solutions for Oracle customers because SNMP is an open standard.

Strictly speaking, SNMP support is intended more for monitoring Oracle products than for managing them. SNMP support is invaluable for tracking the status of an entire network of Oracle applications — first, to verify normal operations, and second, to spot and react to potential problems as soon as they are detected. However, for purposes of investigating and solving some problems, other Oracle tools such as Oracle SQL *Plus Worksheet may be more appropriate. This is because SNMP support is designed to query status, but not to change system parameters, whereas other tools are designed to set or tune system parameters.

Note: Oracle SNMP is not supported on HP OpenVMS platform.

12.2 About the SNMP Management Station

The SNMP management station refers to a node from which managed elements are monitored using the SNMP protocol. Typically, it is a standalone workstation that is on the same network as the managed elements. While this book will consistently use the term SNMP management station, other terms used for it include management console, management system, or managing node.

Because most frameworks use SNMP as a basis for communication, Oracle products that support SNMP can be integrated into virtually every management framework. Third-party products such as CA Unicenter, HP OpenView, Tivoli NetView, Aprisma Spectrum, Sun Solstice, and Castle Rock SNMPc Network Manager provide SNMP Management Station functionality.

12.3 How Enterprise Manager Supports SNMP

Enterprise Manager supports SNMP by integrating with third-party management systems, sharing event information through SNMP traps and extending the monitoring scope of Enterprise Manager by monitoring new devices and targets using SNMP.

There are number of ways that Enterprise Manager uses SNMP as illustrated in [Figure 12-1](#):

1. Sharing event information with third-party management systems by generating SNMP trap notifications from Enterprise Manager to an SNMP management station. For example, you can use SNMP to notify a third-party application that a selected metric has exceeded its threshold.

This method supports SNMP version 1 (SNMPv1).

For more information, see [Section 12.4, "Sending SNMP Trap Notifications"](#) and [Chapter 4, "Using Notifications"](#).

2. Extending the monitoring scope of Enterprise Manager to new entities by receiving SMNP traps or fetching SNMP data from or to the managed entity. By

developing a metadata plug-in to receive SNMP traps, you can enable Enterprise Manager to monitor a product capable of throwing SNMP traps.

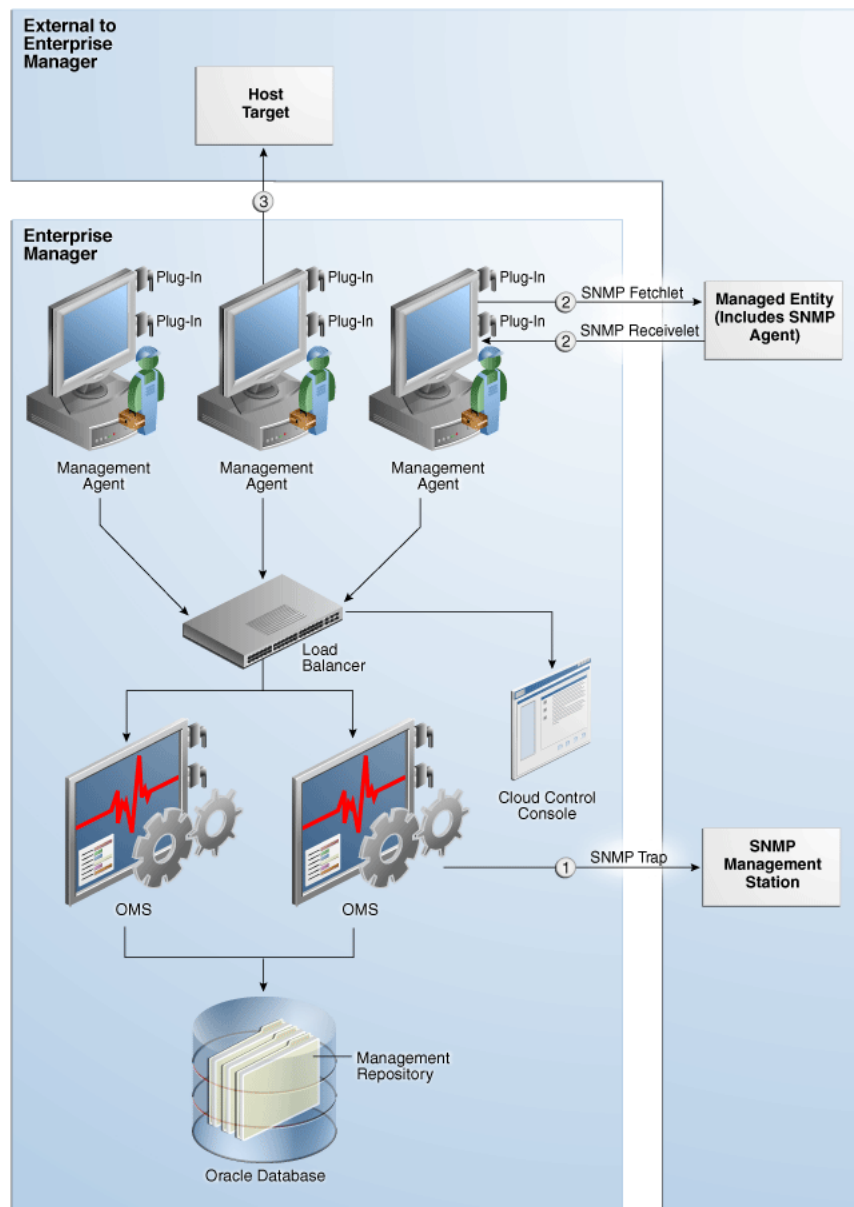
This method supports SNMP version 1 (SNMPv1), SNMP version 2 (SNMPv2c), and SNMP version 3 (SNMPv3).

For more information, see [Section 12.5, "Monitoring External Devices Using SNMP"](#) and the *Oracle Enterprise Manager Cloud Control Extensibility Programmer's Reference*.

3. Creating new metrics to query SNMP agents by using SNMP adapters to allow Management Agents to query native SNMP agents on host targets for Management Information Base (MIB) variable information to be used as metric data.

This method supports SNMP version 1 (SNMPv1), SNMP version 2 (SNMPv2c), and SNMP version 3 (SNMPv3).

For more information, see [Section 12.6, "About Metric Extensions"](#) and [Chapter 8, "Using Metric Extensions"](#).

Figure 12–1 How Enterprise Manager Supports SNMP

12.4 Sending SNMP Trap Notifications

Using the Enterprise Manager notification system, you can share Enterprise Manager event information with other SNMP-enabled third-party applications through SNMP traps. Enterprise Manager supports the SNMPv1 protocol for sending traps. For example, you might want to send event information as traps to a third-party applications from Enterprise Manager when one of the following events takes place:

- a certain metric has exceeded a threshold (metric alert event)
- a target is down (target availability event)
- a job fails (job status change event)

Note: For a full list and description of Enterprise Manager event types, see [Section 3.2.1, "Events"](#).

Using SNMP traps with the notification system is a matter of:

1. Defining a notification method that uses an SNMP trap. For more information, see [Section 4.5, "Sending SNMP Traps to Third Party Systems"](#).
2. Assigning the notification method to a rule. You can edit an existing rule or create a new incident rule. For more information, see [Section 3.3.3, "Setting Up Rule Sets"](#).

12.4.1 About the Management Information Base (MIB)

While SNMP allows Enterprise Manager to send information to third-party SNMP-enabled applications, there might be situations where you want SNMP-enabled applications to obtain information from Enterprise Manager. This is accomplished with the help of MIB variables, and by signing up for SNMP traps. Details of the trap contents can be obtained from the MIB variables.

For more information about the MIB, see [Section 4.6, "Management Information Base \(MIB\)"](#) and [Appendix A, "Interpreting Variables of the Enterprise Manager MIB"](#).

Note: A valid Diagnostic Pack license is required to use the Enterprise Manager MIB variables.

12.5 Monitoring External Devices Using SNMP

It is often critical for an administrator to receive alerts from applications that are not managed by Enterprise Manager. Many of these applications can be configured to trigger SNMP traps when an alert condition takes place. You can receive these traps within Enterprise Manager and start monitoring those applications from Enterprise Manager. Having the capability to receive and analyze SNMP traps raised by such applications allows you extend Enterprise Manager's monitoring and alerting capabilities to these applications and reduce monitoring complexity in your IT environments.

You can configure Enterprise Manager to receive the SNMP traps raised by an application (not managed by Enterprise Manager) and display the traps as alerts in Enterprise Manager.

To receive these traps, you must develop a metadata plug-in to represent the managed entity. Then use an SMMP receivelet or SNMP fetchlet to receive or get monitoring data about that entity.

For more information about developing metadata plug-ins, see the *Oracle Enterprise Manager Cloud Control Extensibility Programmer's Reference*.

12.5.1 About SNMP Receivelets

While monitoring third-party entities in your managed environment, if the status of a third-party network element turns unavailable or if its metric severity conditions (metric thresholds) are met or exceeded, the SNMP Agent of that third-party network element sends a notification to the Management Agent. These notifications are in the form of SNMP traps that get triggered asynchronously upon reaching the performance thresholds, and without any requests from the Management Agent.

Since these traps are based on SNMP, the Management Agent uses SNMP Receivelets to receive and translate these SNMP traps into a form compatible with Oracle Management Service.

For more information about the SNMP receivelet, see the *Oracle Enterprise Manager Cloud Control Extensibility Programmer's Reference* available at:

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

12.5.2 About SNMP Fetchlets

Fetchlets are parameterized data access mechanisms available to map relevant data from a managed element into Enterprise Manager's metric format. In the standards area, Enterprise Manager currently uses SNMP Fetchlets to fetch information from SNMP-enabled entities within your managed environment.

The SNMP fetchlet queries the SNMP Agent for data about the managed entity as defined in the target type's Management Information Base (MIB). For more information about the MIB, see [Section 12.4.1, "About the Management Information Base \(MIB\)"](#) and [Section 4.6, "Management Information Base \(MIB\)"](#).

For more information about the SNMP fetchlet, see the *Oracle Enterprise Manager Cloud Control Extensibility Programmer's Reference* available at:

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

12.6 About Metric Extensions

Metric extensions provide you with the ability to extend Oracle's monitoring capabilities to monitor conditions specific to your IT environment. For example, you can create a new metric for a host target type. Use an SNMP Adapter to allow Enterprise Manager Management Agents to query native SNMP agents on target hosts for Management Information Base (MIB) variable information to be used as the metric data.

For more information about Metric Extensions and the SNMP Adapter, see [Chapter 8, "Using Metric Extensions"](#).

Part II

Security

This section contains the following chapter:

- [Configuring Security](#)

Configuring Security

For information about Enterprise Manager Cloud Control Security, see:
Oracle® Enterprise Manager Cloud Control Security Guide
12c Release 3 (12.1.0.3)

Part III

Generating Reports

This section contains the following chapters:

- [Using Information Publisher](#)

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](#)

14.1 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.

14.2 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**.

14.3 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.

14.3.1 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.

14.3.2 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.

14.3.3 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.

14.4 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.

14.4.1 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.

14.4.2 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.

14.4.3 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.

14.5 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 Enterprise Manager via Mobile Devices

This section contains the following chapter:

- [Remote Access To Enterprise Manager](#)

Remote Access To Enterprise Manager

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 also explains how to connect to the Enterprise Manager desktop version directly from the Safari browser.

The following sections describe setup and use of Cloud Control Mobile:

- [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](#)

To use your iDevice to connect directly to the desktop version, see [Section 15.9](#), "Connecting to Enterprise Manager Desktop Version."

15.1 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

15.2 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

15.3 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.

15.4 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 15.5, "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.

15.5 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.

15.6 Using Cloud Control Mobile in Incident Manager

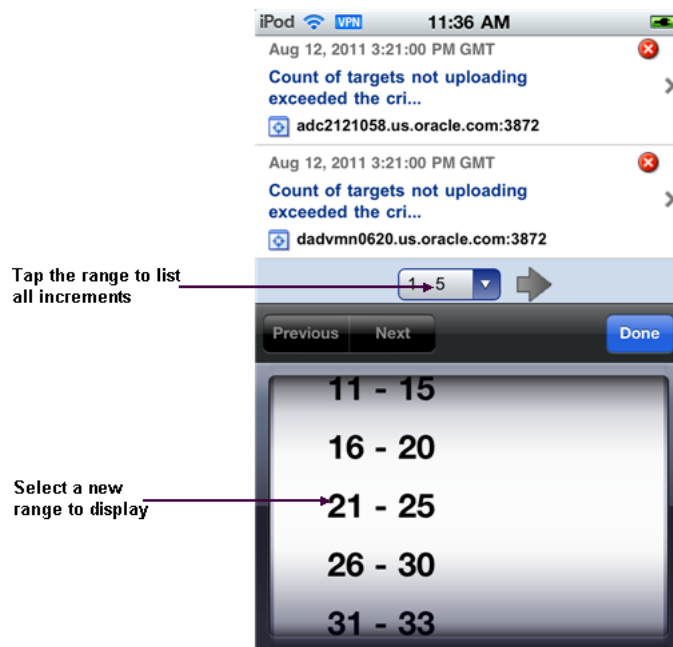
Connecting to Cloud Control remotely, you can do the following in Incident Manager:

- 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 desktop 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 desktop version?

Yes, except for event-related views (standard or custom), which are not available in the app. Any custom views created in the desktop version 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 desktop version, 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 desktop version, 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.

15.7 Working in Cloud Control Mobile

This section covers the following operational tasks in Cloud Control Mobile:

- [Viewing Incidents and Problems](#)
- [Changing Views](#)
- [Performing Actions](#)



15.7.1 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 desktop version. Tap **Incident** there to return to incident details.
- Scroll further and tap **Event List** to see the equivalent of the **Events** tab in the desktop version. 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 desktop version. 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 desktop version. 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.

15.7.2 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.

15.7.3 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.

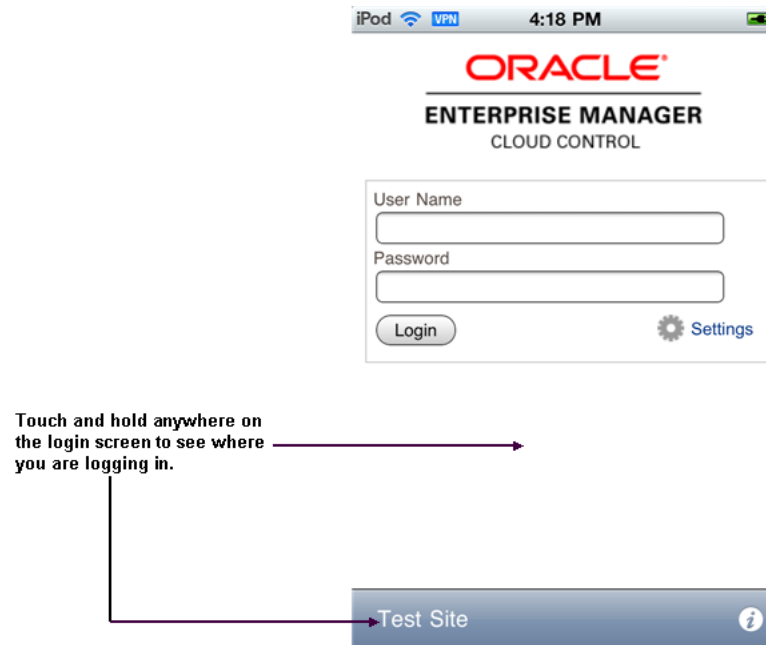
Note that the Acknowledge action may not appear in the action sheet for a variety of reasons; for example, the issue has already been acknowledged or closed, or you do not have the right permissions to acknowledge the issue.

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 desktop version. The only thing missing is the ability to add styles and formatting to the comment.
4. Tap **Save** to complete the action.

15.8 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.

15.9 Connecting to Enterprise Manager Desktop Version

With Enterprise Manager Cloud Control 12c (12.1.0.2), you can log in to the desktop version of Enterprise Manager using Safari on your iDevice, provided your iDevice is on the same network as Enterprise Manager. If you are remote, you may need to establish a VPN connection.

1. Connect to a WiFi or 3G network.
2. Establish a VPN connection if necessary.
3. Open Safari and specify an Enterprise Manager URL.
4. Enter login credentials to access Cloud Control.

With Cloud Control open on your iDevice, you have access to the full feature set. Consider the following as you navigate around the interface:

- Practice gestures to get a sense of how to zoom on a piece of screen real estate.
- Be patient when tapping menu selections; it does not necessarily occur instantaneously.
- Touch and hold a selection to open a context (right-click) menu. This gesture too requires some practice to develop the right sensitivity.
- Not all pages render precisely.
- Pages that have Flex/Flash effects will not render at all.

Note: If you connect to the desktop version of Enterprise Manager through Safari, be sure to set the Safari AutoFill Names and Passwords configuration setting to OFF. Otherwise, the login credentials can be saved to the local store, where they are susceptible to apps scanning for sensitive data.

Part V

Engineered Systems Management

This section contains the following chapters:

- [Discovering and Managing Exadata Targets and Systems](#)

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 is 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](#)

16.1 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 Manager 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. These 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.

16.2 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.

16.3 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.

16.4 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.

16.5 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.

16.6 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

16.7 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](#)
- [Manage IO Resource](#)
- [Accessing Oracle Support Workbench for Exadata Cell](#)
- [Changing the IORM Mode and Updating the IORM Objective](#)

16.7.1 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.

16.7.2 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 IO Resource**.

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.

16.7.3 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.

16.7.4 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 **Manage IO Resource**.

The IORM Performance page appears where you can view the status of the current IORM configuration.

16.7.5 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.

16.7.6 Manage IO Resource

Oracle Exadata Database Machine Cells are added as targets during the database machine discovery workflow and are grouped automatically under the group **Exadata Storage Server Grid**.

To access the IORM Performance page:

1. Select an Exadata Storage Server cell. One way to select the cell:
 - a. From the Targets menu, select **Exadata**.
 - b. Select a DB Machine from the list of Target Names.
 - c. In the Target Navigation pane, expand the Exadata Grid item and click one of the cells.
2. Once you have selected an Exadata Storage Server cell, click the **Exadata Storage Server** menu, select **Administration**, then **Manage IO Resource**.

Once you have accessed the IORM page, you can make the following modifications:

- [Modify the Disk I/O Objective](#)
- [The Inter-Database Plan](#)

For further details on managing I/O resources, refer to the *Managing I/O Resources* chapter in the *Oracle® Exadata Storage Server Software User's Guide*.

16.7.6.1 Modify the Disk I/O Objective

To update the I/O Resource Manager (IORM) settings:

1. Confirm the IORM objective settings before performing the administration task in Enterprise Manager Cloud Control 12c. From the command line, run the following command:

```
# dcli -g cell_group cellcli -e "list iormplan attributes objective"
```

Output should show a value of **off**:

```
cell01: off
cell02: off
cell03: off
.
.
.
cell14: off
```

Note: Starting with Exadata Storage Server Software 11g Release 2 (11.2.3.2) and later, the default output value is **basic** instead of **off**.

2. Navigate to the IORM Performance page as described above.

Note: You can also update a single cell. Expand the Exadata Grid group to view all cells associated with the group. Click the cell you want to update.

The steps to update the IORM settings is the same for a single cell or group of cells.

3. In the Disk I/O Objective drop-down menu, select an objective from the list (**Auto** is the default):
 - **Basic** - optimizes disk IOs based on the IO type. This setting will manage runaway workloads and prevent really long disk queues and latencies.
 - **Auto** - IORM automatically determines the optimization objective as follows. For all consumer groups that are currently issuing I/Os, IORM continuously and dynamically categorizes them as latency-oriented or throughput-oriented. If 75+% of the consumer group's I/Os to disk are small, then it is a latency-oriented consumer group. Otherwise, it is throughput-oriented.
 - **Low Latency** - This setting should be used when Exadata has critical OLTP workloads that require extremely good disk latencies. This setting provides the lowest possible latency by limiting the disk utilization of large I/Os to 50%.

- **Balanced** - This setting should be used when Exadata has critical OLTP and DSS workloads and the user wants to strike a balance between low disk latencies and high throughput. This setting limits the disk utilization of large I/Os to 90%.
- **High Throughput** - This setting should be used when Exadata has critical DSS workloads that require high throughput. This setting provides high throughput performance, at the expense of I/O latency.

Click **Update**. The Exadata Cell Administration Wizard will appear prompting you for the information necessary to complete the Disk I/O Objective configuration:

- a. On the Command page, the Cell Control Command-Line Interface (CellCLI) value should be:

```
# alter iormplan objective = 'auto'
```

Click **Next**.

- b. On the Admin Credentials page, enter the username and password for the selected cells.

Click **Next**.

- c. On the Schedule page, enter a job name (required) and job description (optional). Select an option to start **Immediately** or **Later**. If you select the Later option, enter the time you want the job to run.

Click **Next**.

- d. On the Review page, verify the settings are correct. If there are no changes, click **Submit Command**.

- e. Once the job is successfully submitted, the Job Status page will display.

Click **Return** to return to the I/O Resource Manager (IORM) Settings pane.

4. Click **Get Latest** to refresh the page, which will include your Disk I/O Objective selection.
5. Confirm the IORM objective settings. From the command line, run the following command:

```
# dcli -g cell_group cellcli -e "list iormplan attributes objective"
```

Output should show a value of **auto**:

```
cell01: auto
cell02: auto
cell03: auto
.
.
.
cell14: auto
```

16.7.6.2 The Inter-Database Plan

An inter-database plan specifies how resources are allocated by *percentage* or *share* among multiple databases for each cell. The directives in an inter-database plan specify allocations to databases, rather than consumer groups. The inter-database plan is configured and enabled with the CellCLI utility at each cell.

The inter-database plan is similar to a database resource plan, in that each directive consists of an allocation amount and a level from 1 to 8. For a given plan, the total

allocations at any level must be less than or equal to 100 percent. An inter-database plan differs from a database resource plan in that it cannot contain subplans and only contains I/O resource directives. Only one inter-database plan can be active on a cell at any given time.

You can view the current configured inter-database plan and update an existing Percentage/Share based inter-database plan and a new Percentage/Share based plan can be configured using the Add/Remove options.

You can also view Share, Percentage Radio buttons and a drop down with Basic, Advance options.

Note: If the Exadata plug-in version is 12.1.0.3.0 and earlier or if the Exadata Storage Server version is 11.2.3.1.0 or earlier, the Share, Percentage based inter-database plan radio buttons are not available. You can view only Percentage-based options (that is, the drop-down only displays the Basic, Advance options).

When considering an inter-database plan:

- If Oracle Exadata Storage Server is only hosting one database, then an inter-database plan is not needed.
- If an inter-database plan is not specified, then all databases receive an equal allocation.

For further details on the inter-database plan, refer to the *About Interdatabase Resource Management* section in the *Oracle® Exadata Storage Server Software User's Guide*.

16.7.7 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.

16.7.8 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.

Part VI

Administering Cloud Control

This section contains the following chapters:

- [Personalizing Cloud Control](#)
- [Maintaining Enterprise Manager](#)
- [Maintaining and Troubleshooting the Management Repository](#)
- [Updating Cloud Control](#)
- [Plug-In Manager](#)
- [Patching Enterprise Manager Agents](#)
- [Patching the OMS and Repository](#)
- [Starting and Stopping Enterprise Manager Components](#)
- [Locating and Configuring Enterprise Manager Log Files](#)

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](#)


17.1 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 17-1](#), the menu item displays the SYSMAN user name.

Figure 17–1 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 17–2](#).

Figure 17–2 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 17.2, "Customizing a Region"](#) for details.
 - Click the **View Actions** menu in the upper right corner of the region 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.


17.2 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 17.1, "Personalizing a Cloud Control Page"](#).
2. Click the "ratchet" icon next to the "X" icon within a region, as shown in [Figure 17-3](#). Note that the icon will only be enabled if the region can be customized.

Figure 17-3 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.

17.3 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 17-1](#).

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 17-1](#).
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 Selected** button.
3. Click **OK** when finished.

Plug-In Manager

This chapter provides an overview of Plug-In Manager, and describes how you can use it to view, deploy, administer, and undeploy plug-ins.

In particular, this chapter covers the following:

- [Introduction to Plug-ins](#)
- [Workflow of Plug-In Deployment](#)
- [Introduction to Plug-In Manager](#)
- [Actions that Can Be Performed Using Plug-In Manager](#)
- [Troubleshooting](#)

18.1 Introduction to Plug-ins

This section covers the following:

- [What Is the Enterprise Manager 12c Extensibility Paradigm?](#)
- [What Are Plug-Ins?](#)
- [Are All Plug-Ins Deployed by Default?](#)
- [How Often Are Plug-Ins Released?](#)

18.1.1 What Is the Enterprise Manager 12c Extensibility Paradigm?

Enterprise Manager is system management software that delivers centralized monitoring, administration, and life cycle management functionality for the complete IT infrastructure, including systems running Oracle and non-Oracle technologies.

Enterprise Manager has grown in size and magnitude over the years to offer a spectrum of powerful IT management and monitoring solutions. This growth has led to changes in managing support for new features, enhancements, and bug fixes.

Considering these developments, Oracle has carefully redesigned the architecture of Enterprise Manager in such a way that the framework or the core base on which the product runs is clearly separated from the layer that offers IT solutions by means of features. This new architecture implemented in Enterprise Manager 12c enables Oracle to provide a much stronger framework with capabilities to extend itself seamlessly from time to time for supporting new features and enhancements.

You no longer have to wait for the next release of Enterprise Manager to access the latest monitoring features for released products. The pluggable framework in Enterprise Manager 12c allows target support to be included soon after new versions

of targets ship. You can install a new Enterprise Manager system or upgrade an existing one, as soon as the Enterprise Manager release is made available by Oracle.

Based on the new design, the Enterprise Manager 12c architecture constitutes the following logical parts:

- **EM Platform:** Consists of a set of closely integrated UI and backend services that most monitoring and management functionality in Enterprise Manager depends on. Examples of platform subsystems include the Enterprise Manager target and metric model, the job, event, and provisioning framework. The platform also includes Oracle Management Agent (Management Agent) as well as the core background services such as the data loader, job dispatcher, and notification manager. The platform is delivered as part of an Enterprise Manager release, and can only be upgraded by upgrading to a new version of Enterprise Manager.
- **EM Plug-ins:** Modules that can be plugged to an existing Enterprise Manager Platform to provide target management or other vertical functionality in Enterprise Manager. Plug-ins offer special solutions or new features, for example, connectivity to My Oracle Support, and extend monitoring and management capability to Enterprise Manager, which enable you to monitor a particular target on a host. Plug-ins work in conjunction with OMS and Management Agent to offer monitoring services, and therefore they are deployed to the OMS as well as the Management Agent.

The plug-in releases happen more often than Enterprise Manager Core Platform releases. The plug-ins enable Enterprise Manager 12c to be updated with new features and management support for the latest Oracle product releases, without having to wait for the next platform release to provide such functionality.

18.1.2 What Are Plug-Ins?

Plug-ins are modules that can be plugged into an existing Enterprise Manager Cloud Control deployment to extend target management or other vertical functionality in Enterprise Manager.

At a high level, plug-ins contain archives for monitoring and discovering OMS instances and Management Agents. The archives contain Java and SQL codes, and metadata.

18.1.3 Are All Plug-Ins Deployed by Default?

As a part of Enterprise Manager Cloud Control installation, a set of basic plug-ins is deployed by default. You can deploy other plug-ins to extend the basic functionality of Enterprise Manager Cloud Control.

The plug-ins that are deployed by default, or are shipped out of box are as follows.

- Oracle Database : oracle.sysman.db
- Oracle Fusion Middleware : oracle.sysman.emas
- Oracle MOS (My Oracle Support) : oracle.sysman.mos
- Oracle Exadata : oracle.sysman.xa

18.1.4 How Often Are Plug-Ins Released?

Plug-in releases happen more often than Enterprise Manager Core platform releases. This new pluggable framework enables Enterprise Manager Cloud Control to be

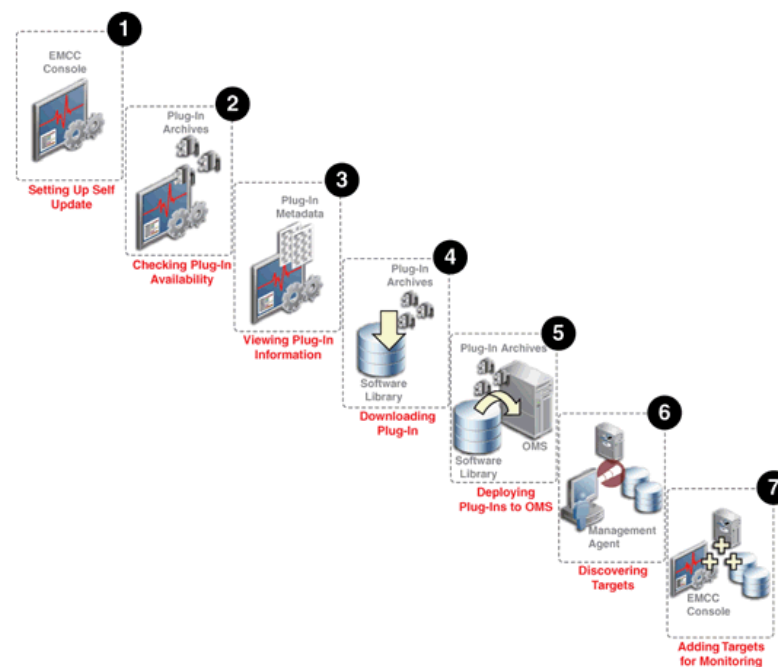
updated with management support for the latest Oracle product releases, without having to wait for the next platform release to provide such functionality.

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 Oracle Database release. You can also work with plug-ins in Offline Mode.

18.2 Workflow of Plug-In Deployment

Figure 18–1 illustrates the workflow of plug-in deployment—how you typically set up the Enterprise Manager infrastructure, deploy plug-ins to OMS, and discovery and monitor targets using the deployed plug-ins.

Figure 18–1 Plug-In Deployment Workflow



Step 1: Setting up Self-Update Console

Self Update console is a common dashboard used for reviewing, downloading, and applying new updates available for Enterprise Manager. The console frees you from having to monitor multiple channels to get informed about new updates that are available from Oracle. The updates automatically downloaded by Self Update include plug-ins. For checking the availability of plug-ins and downloading them to Enterprise Manager, you must set up the Self Update Console. Set up the Self Update Console as described in *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

Step 2: Checking Plug-In Availability

Checking the plug-in availability essentially refers to the act of verifying whether the plug-ins are available on My Oracle Support for download and deployment in Enterprise Manager. This is a prerequisite before downloading plug-ins. To check plug-in availability, follow the steps outlined in [Section 18.4.2](#).

Step 3: Viewing Plug-In Information

Viewing plug-in information refers to the act of viewing basic information related to a particular plug-in, such as the plug-in ID, the plug-in release number, and other basic information. You must view plug-in information to understand what targets and operating systems are certified for plug-ins. You can also check whether or not a particular plug-in has already been deployed. To view plug-in information, follow the steps outlined in [Section 18.4.3](#).

Step 4: Downloading Plug-Ins

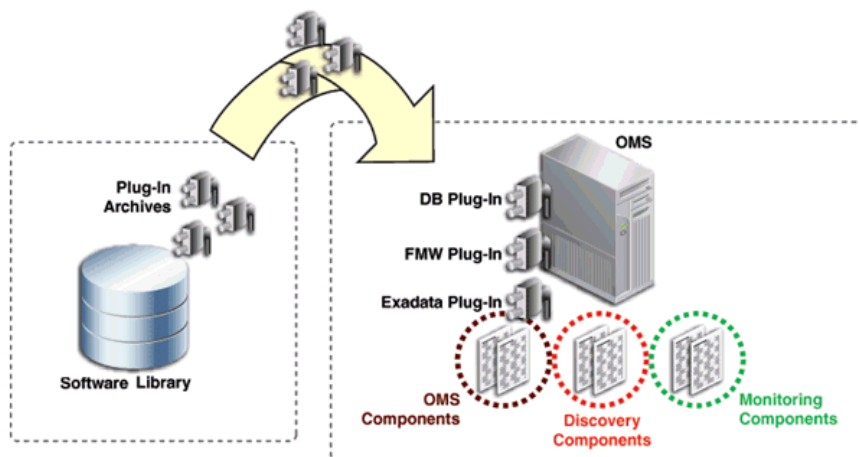
Downloading plug-ins is the act of downloading plug-in archives or components, and its metadata, from My Oracle Support to Oracle Software Library (Software Library), so that they can be deployed suitably for discovering and monitoring certain targets. If you find that a particular target is not being monitored by plug-ins, you must download the required plug-ins. You can download both in online mode and offline mode. To download plug-ins, follow the steps outlined in [Section 18.4.4](#).

Step 5: Deploying Plug-Ins to OMS

Deploying plug-ins to OMS is the next natural course of action once a plug-in is downloaded from My Oracle Support. This is to ensure the OMS capabilities are extended to either manage a new target or to add a new vertical capability. The installation and configuration of plug-ins on the OMS is essentially referred to as *Deployment*. Some plug-ins, when deployed, require the OMS to be re-started.

[Figure 18–2](#) illustrates how plug-ins are deployed to the OMS.

Figure 18–2 Deploying Plug-Ins to OMS



When the plug-in archives are deployed from the Software Library to the OMS, the OMS receives three different components for each plug-in, namely the OMS plug-in components, the discovery plug-in components, and the monitoring plug-in components.

Discovery plug-in components are those components that help in the discovery of unmanaged targets. Monitoring plug-in components are those components that help in the adding of discovered targets to Enterprise Manager Cloud Control Console for monitoring purposes.

To deploy plug-ins on OMS, follow the steps outlined in [Section 18.4.5](#).

Step 6: Discovering Targets

Discovering targets refers to the process of identifying unmanaged hosts and targets in your environment. During discovery of targets, the discovery components of plug-ins are deployed to the Management Agent home. Note that this enables Enterprise Manager Cloud Control to only identify a new target in your environment; it however does not monitor the target.

After converting unmanaged hosts to managed hosts in Enterprise Manager Cloud Control, you must configure automatic discovery of targets on those hosts so that the unmanaged targets running on those hosts can be identified.

For instructions to configure automatic discovery of targets on managed hosts, refer to the Discovering and Monitoring Targets section in the Oracle Enterprise Manager Cloud Control Administrator's Guide, using the following URL:

http://docs.oracle.com/cd/E24628_01/doc.121/e24473/discovery.htm#CBAGJFHC

Note: When you configure automatic discovery of targets on managed hosts, discovery plug-in components are copied to Management Agent.

Once you have configured automatic discovery of targets on managed hosts, you must regularly check for discovered targets so that they can be promoted and monitored in Enterprise Manager Cloud Control.

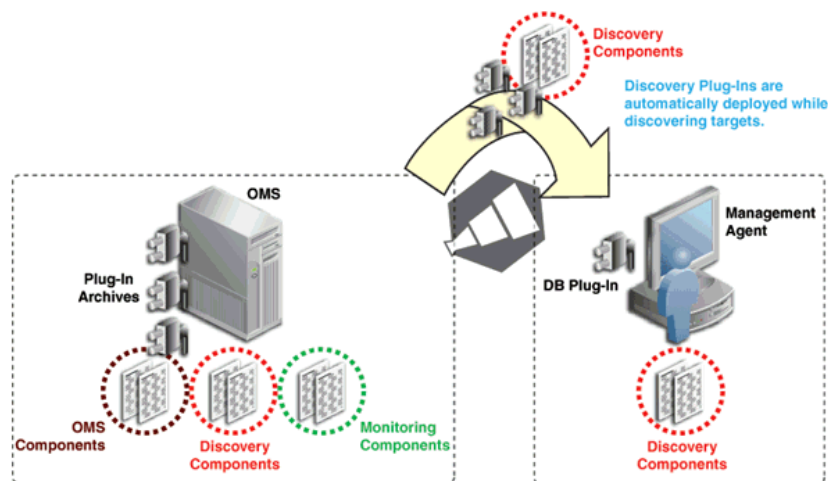
For instructions to check for and promote discovered targets to managed status, refer to the Discovering and Monitoring Targets section in the Oracle Enterprise Manager Cloud Control Administrator's Guide, using the following URL:

http://docs.oracle.com/cd/E24628_01/doc.121/e24473/discovery.htm#CBAFHEHC

Note: The plug-in for a specific target type is automatically deployed to the Management Agent that will monitor targets of that type. For example, if you discover an Oracle Database target, the Discovery Plug-in component of Oracle Database plug-in will be pushed to the Management Agent installed on the database host machine.

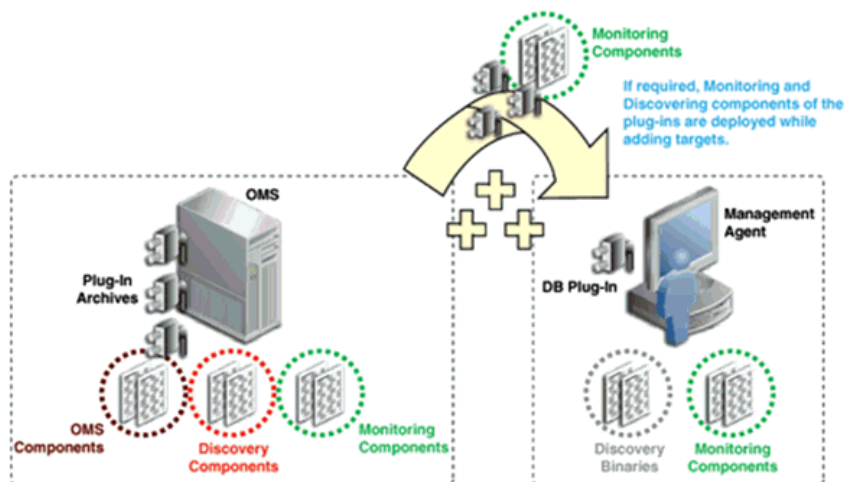
When you check for and promote targets, monitoring plug-in components are copied to Management Agent.

Figure 18–3 illustrates how the discovery plug-in components are deployed to the Management Agent while discovering new targets.

Figure 18–3 Discovering Targets**Step 7: Adding Targets for Monitoring**

Once the targets are discovered, they are added to the infrastructure, so that they can be monitored in Enterprise Manager Cloud Control. While adding targets, the monitoring components of plug-ins are deployed to the Management Agent home.

Figure 18–4 illustrates how the monitoring plug-in components are deployed to the Management Agent while adding targets.

Figure 18–4 Adding Targets

18.3 Introduction to Plug-In Manager

This section introduces you to the Plug-In Manager, and describes some essential concepts you must understand before you start deploying or administering plug-ins. In particular, this section covers the following:

- [What is Plug-In Manager?](#)
- [Accessing and Performing Operations Using Plug-In Manager](#)
- [Actions that Can Be Performed Using Plug-In Manager](#)

18.3.1 What is Plug-In Manager?

Plug-In Manager is a feature of Enterprise Manager Cloud Control, that serves as a single window solution for performing all plug-in deployment-related activities, through GUI as well as CLI. Using Plug-In Manager, you can:

- View plug-ins available for download; plug-ins that have been downloaded; and plug-ins that have been deployed to Cloud Control.
- View certification and critical information about plug-ins such as the name of the plug-in, the vendor who supplied it, the plug-in ID and version, and a short description.
- Deploy plug-ins on OMS.
- Deploy and re-deploy plug-in on Management Agent.
- Create custom plug-in update.
- Undeploy plug-ins from OMS and Management Agent.
- View the status of a plug-in deployment operations.

18.3.2 Accessing and Performing Operations Using Plug-In Manager

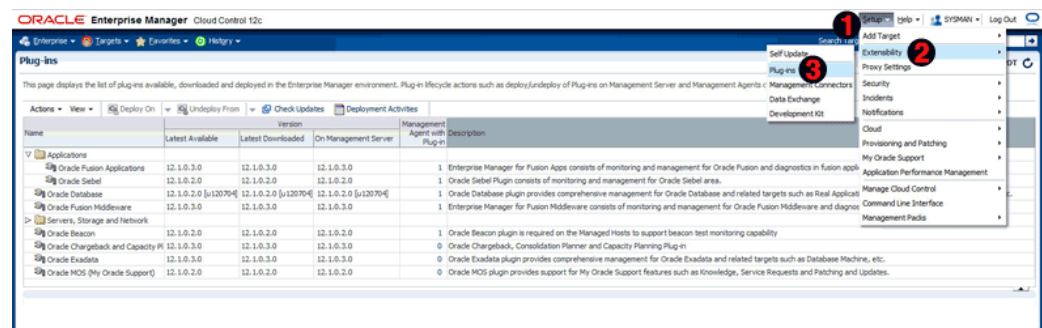
- Accessing Plug-In Manager
- Performing Operations Using Plug-In Manager

18.3.2.1 Accessing Plug-In Manager

To access the Plug-In Manager console, from the **Setup** menu, select **Extensibility**, and then select **Plug-ins**.

Figure 18–5 illustrates how you can access Plug-in Manager.

Figure 18–5 Navigating to Plug-In Manager

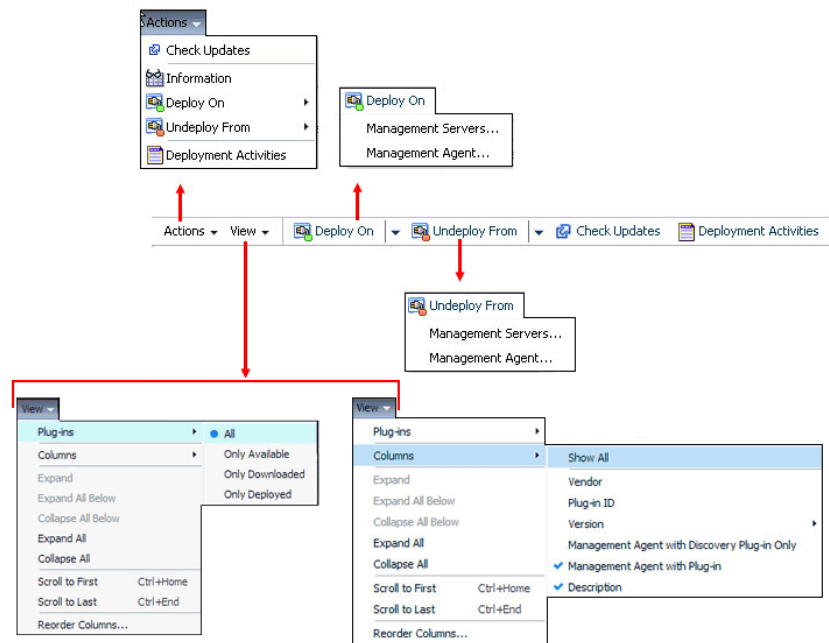


18.3.2.2 Performing Operations Using Plug-In Manager

Using Plug-in Manager, you can deploy, upgrade, redeploy, and undeploy plug-ins.

Figure 18–6 shows the operations you can perform using the Plug-In Manager.

Figure 18–6 Plug-In Manager Operations



18.4 Actions that Can Be Performed Using Plug-In Manager

This section describes the following actions or operations that are supported by the Plug-In Manager.

- Knowing Your Plug-Ins
 - [Customizing Your View](#)
 - [Checking the Availability of Plug-Ins](#)
 - [Viewing Information about Plug-Ins](#)
- Downloading, Deploying, and Upgrading Plug-Ins
 - [Downloading Plug-Ins](#)
 - [Deploying Plug-Ins to Oracle Management Service \(Reduce OMS Restart time and Downtime\)](#)
 - [Upgrading Plug-Ins Deployed to Oracle Management Service](#)
 - [Deploying Plug-Ins on Oracle Management Agent](#)
 - [Upgrading Plug-Ins Deployed to Oracle Management Agent](#)
- Undeploying Plug-Ins
 - [Undeploying Plug-Ins from Oracle Management Agent](#)
 - [Undeploying Plug-Ins from Oracle Management Service](#)
- Advanced Operations with Plug-Ins
 - [Re-deploying Plug-Ins on Oracle Management Agent](#)

- [Deploying Plug-In Patches While Deploying or Upgrading Management Agent \(Create Custom Plug-In Update\)](#)

18.4.1 Customizing Your View

This section tells you how to customize your view, and organize the plug-ins and columns displayed.

18.4.1.1 Customizing Displayed Plug-Ins

Over a period of time, as you download and deploy plug-ins, the number of plug-ins on your list increases. You can sort these plug-ins to view only the ones you require, for example, only the plug-ins available, or only the plug-ins deployed.

In order to customize your view, follow these steps.

1. From the **View** menu, select **Plug-Ins**.
2. From the Plug-Ins menu, select one of the following filters.
 - **All**, using this filter, you can view all plug-ins, including available, downloaded, and deployed plug-ins.
 - **Only Available**, using this filter, you can view the plug-ins that are available for download.
 - **Only Downloaded**, using this filter, you can view the plug-ins that are downloaded.
 - **Only Deployed**, using this filter, you can view the plug-ins that are deployed.

18.4.1.2 Customizing Displayed Columns

By default, only a few columns of information are displayed. Optionally, you can either enable other columns of your interest, or disable ones that are already displayed.

In order to customize the displayed columns, follow these steps.

1. From the **View** menu, select **Columns**.
2. From the Columns menu, select one of the following filters for columns.
 - **Show All**, using this filter, you can view all columns.
 - **Vendor**, using this filter, you can view information about the vendor.
 - **Plug-In Id**, using this filter, you can view the plug-in id.
 - **Version**, this filter has three options you can choose from. They are as follows.
 - **Latest Available**, using this filter, you can view the newest plug-ins that are available.
 - **Latest Downloaded**, using this filter, you can view the plug-ins that have been downloaded recently.
 - **On Management Server**, using this filter, you can view the plug-ins that are deployed to the OMS.
 - **Management Agent with Discovery Plug-Ins Only**, this filter displays the Management Agent which has only Discovery Plug-Ins deployed.
 - **Management Agent with Plug-In**, this filter displays the Management Agent which has any plug-in deployed on it.

- **Description**, this filter displays the description of the plug-ins.

18.4.2 Checking the Availability of Plug-Ins

To check the availability of plug-ins, follow these steps:

1. Set up the Self Update Console as described in *Oracle Enterprise Manager Cloud Control Administrator's Guide*.
2. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
3. On the Plug-ins page, in the Latest Available column of the table, check whether the plug-ins are available.

To refresh the list of available plug-ins, click **Check Updates**. Note that clicking Check Updates will take you to the Self Update page.

18.4.3 Viewing Information about Plug-Ins

This section gives you more information on plug-ins, and functions related to plug-ins. This section covers the following sections:

- [Differentiating Plug-In Releases from Enterprise Manager Platform Releases](#)
- [Identifying Plug-In ID](#)
- [Viewing Targets and Operating Systems Certified for Deployed Plug-Ins](#)
- [Viewing Plug-In Dependencies](#)
- [Verifying Deployed Plug-Ins](#)


18.4.3.1 Differentiating Plug-In Releases from Enterprise Manager Platform Releases

Plug-ins have independent release cycles and release numbers, which may or may not be tied to Enterprise Manager Cloud Control product releases and release numbers.


Plug-in releases typically happen more often than Enterprise Manager platform releases.

[Figure 18–7](#) describes how plug-in releases are numbered.

Figure 18–7 Plug-In Release Number Format



[Figure 18–8](#) describes how Enterprise Manager platform releases are numbered.

Figure 18–8 Enterprise Manager Core Platform Release Number Format

18.4.3.2 Identifying Plug-In ID

To identify the ID of a plug-in, follow these steps:

1. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
2. On the Plug-ins page, in the Plug-in ID column of the table, note the plug-in ID of the plug-in of your interest.

If you do not see this column, from the **View** menu, select **Columns**, then select **Plug-in ID**.

Figure 18–9 illustrates how you can identify the plug-in ID of the Oracle Database plug-in.

Figure 18–9 Identifying Plug-In ID

Name	Plug-in ID	Version		
		Latest Available	Latest Downloaded	On Management Server
Applications				
Databases				
IBM DB2 Database	oracle.em.sldb	12.1.0.2.0	12.1.0.2.0	12.1.0.2.0
Microsoft SQLServer Database	oracle.em.ssmss	12.1.0.2.0		
Oracle Database	oracle.sysman.db	12.1.0.2.0 [u120427]	12.1.0.2.0 [u120427]	12.1.0.2.0 [u120427]
Sybase ASE Database	oracle.em.ssad	12.1.0.2.0	12.1.0.2.0	12.1.0.2.0

18.4.3.3 Viewing Targets and Operating Systems Certified for Deployed Plug-Ins

To view a list of targets and operating systems certified for a deployed plug-in, follow these steps:

1. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
2. On the Plug-ins page, select the plug-in of your interest, and from the **Actions** menu, select **Information**.
3. On the Plug-in Information page, in the **General** tab, review the information provided in the **Certified Targets** and **Certified Operating Systems** tables.

18.4.3.4 Viewing Plug-In Dependencies

To view the dependencies of the preferred plug-in, follow these steps:

1. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
2. On the Plug-ins page, select the plug-in of your interest, and from the **Actions** menu, select **Information**.
3. On the Plug-in Information page, in the **Dependencies** tab, review the information provided in the tables.

18.4.3.5 Verifying Deployed Plug-Ins

To view and administer the deployed plug-ins, from the **Setup** menu, select **Extensibility**, then select **Plug-ins**. Enterprise Manager Cloud Control displays the Plug-ins page, which is essentially the *Plug-In Manager* console.

This section covers the following sections:

- [Identifying the Management Agent and OMS on Which a Plug-In Is Deployed](#)
- [Using EMCLI to Identify Plug-Ins Deployed on Oracle Management Service and Oracle Management Agent](#)

18.4.3.5.1 Identifying the Management Agent and OMS on Which a Plug-In Is Deployed To identify the OMS instances on which the plug-in of your interest is deployed, follow these steps using Enterprise Manager Cloud Control Console:

1. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
2. On the Plug-ins page, select the plug-in of your interest, and from the **Actions** menu, select **Information**.
3. On the Plug-in Information page, in the **Management Servers** tab, review the Oracle Management Services on which the plug-in is deployed.

To identify the Management Agents on which the plug-in of your interest is deployed, follow these steps using Enterprise Manager Cloud Control Console:

1. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
2. On the Plug-ins page, select the plug-in of your interest, and from the **Actions** menu, select **Information**.
3. On the Plug-in Information page, in the **Management Agent** tab, review the Management Agents on which the plug-in is deployed.

18.4.3.5.2 Using EMCLI to Identify Plug-Ins Deployed on Oracle Management Service and Oracle Management Agent

To identify the Plug-ins deployed on OMS, on EM CLI, log in to EMCLI, and enter the following command. The command displays a list of all the plug-ins deployed on the OMS.

```
$emcli login
-username=<EM Console Username>
[-password=<EM Console Password>]
[-force]
$emcli list_plugins_on_server
```

Example 18–1 Sample List of Plug-Ins Deployed on OMS

slc01mib.us.oracle.com:7654_Management_Service

Plug-in Name	Plug-in ID	Version (Revision)
Oracle Database	oracle.sysman.db	12.1.0.4.0
Oracle Fusion Middleware	oracle.sysman.emas	12.1.0.4.0
Oracle MOS (My Oracle Support)	oracle.sysman.mos	12.1.0.5.0

Plug-in Name	Plug-in ID	Version (Revision)
Oracle Exadata	oracle.sysman.sa	12.1.0.4.0

To identify and view all the Plug-ins deployed on Management Agent, on EM CLI, enter the following command:

```
$emcli list_plugins_on_agent
[agent_names="agent1,agent2,agent3..."
[-all] [-include_discovery]
```

Example 18–2 Sample List of Plug-ins Deployed on Management Agent

```
emcli list_plugins_on_agent -agent_names=abc.us.oracle.com:3872
Lists plug-ins on the agent abc.us.oracle.com
emcli list_plugins_on_agent -agent_
names=abcd.us.oracle.com:3872,efgh.us.oracle.com:3872 -include_discovery
Lists plug-ins on both the agents provided along with their discovery components
emcli list_plugins_on_agent -agent_names='adc*,st*93'
Lists plug-ins on all agents with name matching one of the regular expressions
adc* or st*93
emcli list_plugins_on_agent -all
Lists plug-ins on all the management agents.
```

18.4.4 Downloading Plug-Ins

You can download the plug-ins in online or offline mode. Online refers to an environment where you have Internet connectivity to the Enterprise Manager Store. Offline refers to an environment where you do not have Internet connectivity. This section contains the following sections:

- [Downloading Plug-Ins in Online Mode](#)
- [Downloading Plug-Ins in Offline Mode](#)

18.4.4.1 Downloading Plug-Ins in Online Mode

To download the plug-ins in online mode, follow these steps:

1. From the **Setup** menu, select **Extensibility**, then select **Self Update**.
2. On the Self Update page, in the table, click on **Plug-in**.
3. On the Plug-in Updates page, select the plug-in available for download, and click **Download**.

Multiple selection of plug-ins is not supported.

4. In the Schedule Download dialog, select an appropriate option to schedule the download. You can also select **Immediately** which schedules the job for immediate action. Select **Notify Once downloaded** if you want to be informed once the download is complete.
5. Click **Select**.

Enterprise Manager Cloud Control submits a job to download the selected plug-in from the Enterprise Manager Store to the Software Library.

A confirmation dialog appears to confirm that the job has been submitted successfully. In this confirmation dialog, you can click **Job Details** to track the status of the job.

18.4.4.2 Downloading Plug-Ins in Offline Mode

To download the plug-ins in offline mode, follow these steps:

1. Set Enterprise Manager Cloud Control to Offline Mode. To do so, follow these steps.
 - a. From the **Setup** menu, select **Provisioning and Patching**, then select **Offline Patching**.
 - b. In the Online and Offline Settings tab, select **Offline**.
2. From the **Setup** menu, select **Extensibility**, then select **Self Update**.
3. On the Self Update page, click **Check for Updates**.

A message appears with the following URL to an Oracle site from where the updates catalog file can be downloaded.

https://updates.oracle.com/Orion/Download/download_patch/p9348486_112000_Generic.zip

4. From an Internet-enabled computer, download the catalog file using the aforementioned URL.
5. Copy the downloaded catalog file to the OMS host or the Management Agent host where you plan to deploy the plug-ins.
6. Import the catalog file to Enterprise Manager. For instructions, refer to [Importing Catalog Archives](#).
7. On the Self Update page, in the table, click **Plug-in**.
8. On the Plug-in Updates page, select the imported update that is available for download. Click **Download**.

A message appears with a URL to an Oracle site from where the update can be downloaded.
9. From a computer that is connected to the internet, download the update using the aforementioned URL.
10. Copy the downloaded file to the OMS host or the Management Agent host where you plan to deploy the plug-ins.
11. Import the downloaded plug-in archive to Enterprise Manager. For instructions, refer to [Importing Plug-In Archives](#).

18.4.4.3 Importing Catalog Archives

To import a catalog archive, follow these steps:

1. Download the catalog archive as described in [Section 18.4.4.2](#).
2. Execute the following `emcli` command to import the downloaded catalog archive.

```
$emcli import_update_catalog
-file="file"
-omslocal
emcli import_update_catalog
-file="file"
-host="hostname"
```

```
[ -credential_set_name="setname" ] | -credential_name="name"
-credential_owner="owner"
```

Example 18–3 Sample for Importing Catalog Archive

```
$emcli import_update_catalog
  -file="/u01/common/p9984818_121000_Generic.zip"
  -omslocal
```

Imports the master catalog file p9984818_121000_Generic.zip. The file must exist on the OMS host. In a multiple OMS setup, the request can be processed by any OMS, so the file should be accessible from the OMS processing the request. This means that the file must be kept on a shared location that is accessible from all the OMS instances.

```
$emcli import_update_catalog
  -file="/u01/common/p9984818_121000_Generic.zip"
  -host="host1.example.com"
  -credential_set_name="HostCredsNormal"
```

Imports the master catalog file p9984818_121000_Generic.zip that is present on the host host1.example.com. The host must be a managed host target in Enterprise Manager, and the Management Agent on this host must be up and running. The preferred unprivileged credentials for host host1.example.com are used to retrieve the remote file.

18.4.4.4 Importing Plug-In Archives

Import plug-in archives to Oracle Software Library in the following cases:

- When you want to deploy any non-Oracle plug-ins, that is, plug-ins that have been created by a company other than Oracle, and are not available for download on the Self Update console.
- When you want to import other types of entity archives when Self Update is used in offline mode.

To import a plug-in archive, follow these steps:

1. Download the external archive as described in the previous section.
2. Set up the Enterprise Manager Command Line (EM CLI) utility. To do so, from the **Setup** menu, click **Command Line Interface**. Follow the instructions outlined on the Enterprise Manager Command Line Interface Download page.
3. Import the external archive in one of the following ways, depending on where EMCLI is installed.

- If Enterprise Manager server is on the system on which you downloaded the plug-in archive (*.opar file), run the following command:

```
$emcli import_update
  -file="<path to *.opar file>"
  -omslocal
```

The `-omslocal` flag indicates that the plug-in archive path mentioned in the `-file` option is directly accessible to the EM server.

- If Enterprise Manager server 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.

Note: 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.

18.4.5 Deploying Plug-Ins to Oracle Management Service (Reduce OMS Restart time and Downtime)

You can deploy plug-ins to an OMS instance in graphical interface or command line interface. While the graphical interface mode enables you to deploy one plug-in at a time, the command line interface mode enables you to deploy multiple plug-ins at a time, thus saving plug-in deployment time and downtime, if applicable.

Note:

- To view a visual demonstration on how you can deploy a plug-in to the OMS and discover targets using it, from the Oracle Learning Library, select the Oracle Enterprise Manager 12c: Deploy a Plugin video, using the following URL:
https://apex.oracle.com/pls/apex/f?p=44785:24:491956260237501::NO::P24_CONTENT_ID,P24_PREV_PAGE:6000,1
- Plug-ins must be deployed on the OMS prior to being deployed on Management Agents.
- In a multi OMS environment, Plug-in Manager automates plug-in deployment on all the management servers.
- A plug-in upgrade failure could put the Management Repository in an inconsistent state. Oracle recommends that your repository database should be running in archive log mode, and that your backup policies are in place.
- The deployment time varies from one plug-in to another, depending on the volume of data populated in the Management Repository. A page is displayed that allows you to monitor the deployment status, as described in [Section 18.4.5.1](#).
- The deployment of some plug-ins requires the OMS to be stopped, and then restarted. This process occurs automatically as part of the plug-in deployment process.
- While deploying plug-ins to the OMS, OMS plug-in components, discovery plug-in components, and monitoring plug-in components are deployed to the OMS.

To deploy plug-ins to the OMS in graphical mode, follow these steps:

1. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
2. On the Plug-ins page, select the plug-in you want to deploy.
3. From the **Deploy On** menu, select **Management Servers**.
4. In the Deploy Plug-in on Management Servers dialog, enter the Management Repository SYS password, and click **Continue**. Proceed through the steps in the dialog box.
5. Click **Deploy**.

To deploy plug-ins to the OMS in silent mode, follow these steps:

1. Log in to EMCLI as follows:
`$ORACLE_HOME/bin/emcli login -username=sysman`
2. Run the following command if the emcli client is an old version, and does not have all required verbs:
`$ORACLE_HOME/bin/emcli sync`
3. To deploy the plug-ins on the OMS, run the following command:
`$emcli deploy_plugin_on_server`

```
-plugin="plug-in_id[:version]
[-sys_password=sys_password]
[-prereq_check]"
```

Note: For information on plug-in id, refer to [Section 18.4.3.2](#).

For example,

```
$emcli deploy_plugin_on_server
-plugin="oracle.sysman.db:12.1.0.2.0;oracle.sysman.emas:12.1.
0.2.0"
```

Note: The procedure for plug-in deployment remains the same even in a multi-OMS environment. Enterprise Manager automatically detects whether it is a single-OMS or a multi-OMS environment and in case of a multi-OMS environment, Enterprise Manager automatically deploys the selected plug-in on all OMS instances.

If the plug-in deployment on any OMS fails, identify the cause for the failure. If there is a fix or a workaround, fix the problem, and perform the same steps again. The system automatically detects which OMS instances do not have the plug-ins deployed, and deploys them on those servers.

If the problem persists, contact Oracle Support.

18.4.5.1 Tracking the Deployment Status of Plug-Ins on Oracle Management Service

This section describes the procedure of monitoring the deployment status of plug-ins that do not require down time as well as those that do require down time. The following sections are covered in this section:

- [Tracking the Status of Plug-Ins that Do Not Require Down Time](#)
- [Tracking the Deployment Status of Plug-Ins that Require Down Time](#)

18.4.5.1.1 Tracking the Status of Plug-Ins that Do Not Require Down Time To monitor the status of deployment and undeployment operations for plug-ins that do not require down time, follow these steps:

1. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
2. On the Plug-ins page, do one of the following:
 - From the **Actions** menu, select **Deployment Activities**.
 - Select a plug-in, and click the **Recent Deployment Activities** tab at the bottom of the page. Alternatively, you can also run the following command using EMCLI.

```
$emcli get_plugin_deployment_status -plugin_id=<plugin_id>
```

18.4.5.1.2 Tracking the Deployment Status of Plug-Ins that Require Down Time To monitor the status of deployment and undeployment operations of plug-ins that require down time, execute the following command:

```
emctl status oms -details
```


18.4.6 Upgrading Plug-Ins Deployed to Oracle Management Service

This section contains the following sections:

- [Upgrading Across Plug-In Versions Deployed to Oracle Management Service](#)
- [Upgrading Across Plug-In Revisions Within a Plug-In Version Deployed to Oracle Management Service](#)

18.4.6.1 Upgrading Across Plug-In Versions Deployed to Oracle Management Service

You can upgrade across plug-in versions, that is, from one plug-in version to another higher plug-in version or a revision of another higher plug-in version. For example, from Oracle Database plug-in version 12.1.0.1.0 to version 12.1.0.2.0, or from Oracle Database plug-in version 12.1.0.1.0 to version 12.1.0.2.0 [u120427].

To upgrade across plug-in versions deployed to the OMS, follow these steps:

1. Check for the latest available versions and revisions in the Enterprise Manager Store as described in [Section 18.4.2](#).
2. Download them as described in [Section 18.4.4](#).
3. Deploy them to the OMS as described in [Section 18.4.5](#).

18.4.6.2 Upgrading Across Plug-In Revisions Within a Plug-In Version Deployed to Oracle Management Service

This section describes how you can upgrade across plug-in revisions within a plug-in version, that is, from one revision of a plug-in version to another revision of the same plug-in version. For example, from Oracle Database plug-in version 12.1.0.2.0 [u120427] to version 12.1.0.2.0 [u120502].

To upgrade the plug-in revisions deployed to the OMS, follow these steps:

1. Check for the latest available revisions in the Enterprise Manager Store as described in [Section 18.4.2](#).
2. Download them as described in [Section 18.4.4](#).
3. Apply the patches required for the plug-in revision.
 - a. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
 - b. On the Plug-ins page, select the plug-in revision you downloaded, and from the **Actions** menu, select **Information**.
 - c. On the Plug-in Information page, in the **General** tab, in the **Dependencies** section, in the **Bug Fixes Required** table, click **View Bug Fixes**.
 - d. In the Required Bug Fixes dialog, note the bug numbers listed in the Bug Number column of the Management Server Fixes table.
 - e. Download the patches that contain the fixes for the bugs you noted in the previous step.
 - f. Apply the downloaded patches on the OMS.
4. Deploy the latest available revisions to the OMS as described in [Section 18.4.7](#).

18.4.7 Deploying Plug-Ins on Oracle Management Agent

While installing a Management Agent using the Add Host Targets Wizard, all the core discovery plug-ins available on the OMS are automatically deployed to the Management Agent.

For information about discovery plug-ins, refer to [Section 18.4.3](#).

If you want to deploy any additional plug-ins after installing the Management Agent, then follow these steps:

1. Set up the Self Update console.
2. Check whether the plug-ins are available on Enterprise Manager store. For instructions refer to [Section 18.4.2](#).
3. Download the available plug-ins. For instructions, refer to [Section 18.4.4](#).
4. Deploy the downloaded plug-ins to the Management Agent.
 - a. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
 - b. On the Plug-ins page, select the plug-in you want to deploy.
 - c. From the **Deploy On** menu, select **Management Agent**.
 - d. Follow the steps mentioned in the Deploy Plug-ins on Management Agent dialogue box.
 - e. Click **Deploy**.

To deploy plug-ins in EM CLI, use the following command:

```
$emcli deploy_plugin_on_agent  
-agent_names="agent1[;agent2...]"  
-plugin="plug-in_id[:version]"  
[-discovery_only]
```

18.4.7.1 Tracking the Deployment Status of Plug-Ins on Oracle Management Agent

To track the deployment status of plug-ins on Management Agent, refer to [Section 18.4.5.1](#).

18.4.8 Re-deploying Plug-Ins on Oracle Management Agent

Using this option, you can re-deploy plug-ins on Oracle Management Agent. The re-deploy plug-in option reconfigures the same plug-in on the Management Agent, and does not change the configuration details.

```
$emcli redeploy_plugin_on_agent  
{-agent_names="agent1[;agent2...]" | -group_name="group1"}  
-plugin="plug-in_id:version"  
[-redploy_noprompt]
```

Note: While using this option, note that the existing plug-in home will be overwritten, and all applied patches will be lost.

The re-deploy wizard displays the following warning message:

Re-deployment of a plug-in overwrites the existing OracleHome of a plug-in and you will lose any patch(es) that has been applied on plug-in OracleHome.

To continue, click **Yes**.

The redeploy command cannot be used on multiple Management Agents without having Custom Plug-in Update for a plug-in.

If the latest revision of a plug-in needs to be deployed on multiple Management Agents, then run the following command.

```
$deploy_plugin_on_agent allow_revision_update
```

Note: After a metadata plug-in is redeployed, it is recommended that you run the following command.

```
$emcli metric_control -command=flush_metadata_cache
```

The command should be run on all OMS instances.

18.4.9 Upgrading Plug-Ins Deployed to Oracle Management Agent

You can upgrade across plug-in versions, that is, from one plug-in version to another, higher plug-in version or a revision of another, higher plug-in version. For example, from Oracle Database plug-in version 12.1.0.1.0 to version 12.1.0.2.0, or from Oracle Database plug-in version 12.1.0.1.0 to version 12.1.0.2.0 [u120427].

Note: You will upgrade the plug-in versions and revisions only on Management Agents that are already installed in your environment.

When a plug-in is deployed explicitly or a target is promoted on new Management Agents, then the latest plug-in version and revision automatically gets included from the OMS.

To upgrade across plug-in versions deployed to the Management Agent, follow these steps:

1. Check for the latest available versions and revisions in the Enterprise Manager Store as described in [Section 18.4.2](#).
2. Download them as described in [Section 18.4.4](#).
3. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
4. On the Plug-ins page, select the plug-in you want to upgrade.
5. From the **Deploy On** menu, select **Management Agent**.
6. In the Deploy Plug-in on Management Agent dialog, select the version or revision of the plug-in you want to upgrade to., and click **Continue**.
7. Select the preferred Management Agent to upgrade the plug-in on, and click **Continue**. Then click **Next**. And then click **Deploy**.
8. On the Confirmation dialog, click **Close**.

18.4.10 Deploying Plug-In Patches While Deploying or Upgrading Management Agent (Create Custom Plug-In Update)

When a new plug-in is released, it can be downloaded using Self-update. If there are defects with the Management Agent plug-ins, Oracle then releases O-patch style patches. While plug-ins get deployed automatically during target discovery on Management Agents, patches for the plug-ins have to be applied on each plug-in manually.

Custom plug-in update is the user copy of the plug-in, along with patches applied to it. Using the Create Custom Plug-In Update command allows you to create a custom copy of plug-in along with the patches applied in self update. Once the patches are applied, you can create a custom plug-in update of the plug-ins on that Management Agent. The custom plug-in update then becomes a gold image for that plug-in with all the patches applied on that Oracle Home, along with the base plug-in binaries.

After the Custom Plug-in Update is created, any plug-in deployment operation for the plug-in on any Management Agent, using either UI or EMCLI, the new custom copy will be deployed instead of the Oracle supplied version. In this way you don't have to reapply the plug-in patches manually on each pugin home of agent. This custome plug-in image is also used by Agent deployment ot upgrade activity so that the plug-ins getting deployed on these agents are with the patch included.

There are two methods of creating Custom Plug-in Update. The following sections describe the two methods.

- [Creating Custom Plug-In Update Using EMCLI](#)*(recommended)*
- [Creating Custom Plug-In Update Using EDK](#)

18.4.10.1 Creating Custom Plug-In Update Using EMCLI

To create a custom plug-in update, follow these steps:

1. Select a test Management Agent which is up and running on which the preferred plug-in is already deployed. Apply any patches that you want to apply on this plug-in.
2. Perform the required testing.
3. Create a custom plug-in update using the following command:

```
$emcli create_custom_plugin_update  
-agent_name="agent_name"  
-plugin_id="plugin_id"
```

Note: To overwrite and update your current custom plug-in update, use the `overwrite` option.

```
$emcli create_custom_plugin_update
-agent_name="agent_name"
-plugin_id="plugin_id"
[-overwrite]
```

This command creates and imports a custom plug-in update from an existing Management Agent where the selected plug-in is deployed. The custom plug-in update will be used for all subsequent plug-in deployments on any Management Agent, in place of Oracle supplied versions.

Custom plug-in update is created as per plug-in type. If a custom plug-in update is created, and after three days, a patch is applied, in order to include the patch, the custom plug-in update will have to be created again.

To view a list of all Custom Plug-in Updates created, run the following command.


```
$emcli list_custom_plugin_updates
```

To view a list of patches included in a particular Custom Plug-in Update, run the following command.

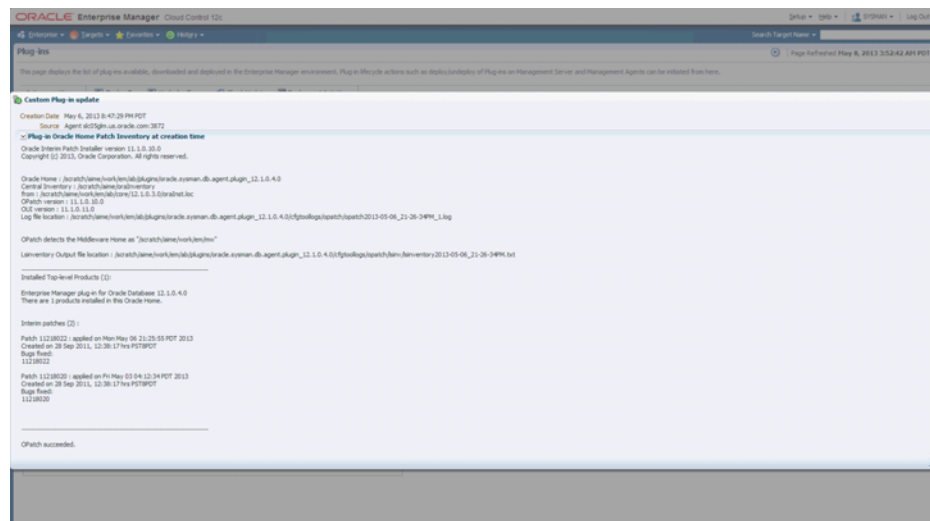
```
$emcli list_patches_in_custom_plugin_update -plugin=<plugin_id>:<version> [-discovery]
```

On the Plug-in Manager console, when you select a plug-in, if a Custom Plug-in Update exists, an icon is displayed beside the version identifier, indicating that that particular plug-in version is customized in the environment, with a list of patches. [Figure 18–10](#) displays the Custom Plug-in Update icon. Once the custom plug-in update exists, it will be used by the Management Agent for deployment and upgrade automatically.

Figure 18–10 Custom Plug-in Update Icon



When you click the Custom Plug-in Update icon, the page that displays the information on Custom Plug-in Update is displayed. [Figure 18–11](#) displays the Custom Plug-in Update information page.

Figure 18–11 Custom Plug-in Update Information Page

18.4.10.2 Creating Custom Plug-In Update Using EDK

To create custom plug-in update using EDK, follow these steps.

1. Download EDK, using the UI or EMCLI, on the Management Agent Host.

To download EDK using UI, from the **Setup** menu, select **Extensibility**, and then select **Development Kit**.

To download EDK using EMCLI, run the following command.

```
$emcli get_ext_dev_kit
```

2. Run the following command.

```
$empdk create_custom_plugin_update -out_dir <output_dir>
-agent_state_dir <agent_state_dir>
```

```
-agent_oracle_home <agent_oracle_home>
-plugin_id <plugin_id>
```

For help with empdk commands, run the following command.

```
$/empdk -help
```

The plug-in update is saved to on a local directory as a .zip file. The .zip file has to be copied to an OMS instance. Once the .zip file is created, from the OMS Home, run the following command to import the custom plug-in update.

```
$emcli import_plugin_update -archive=<archive path>
```

On the Plug-in Manager console, when you select a plug-in, if a Custom Plug-in Update exists, an icon is displayed beside the version identifier, indicating that that particular plug-in version is customized in the environment, with a list of patches. Figure 1-1 displays the Custom Plug-in Update icon.

18.4.11 Undeploying Plug-Ins from Oracle Management Agent

To undeploy plug-ins from the Management Agent, follow the steps:

1. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
2. On the Plug-ins page, select the plug-in you want to undeploy, and from the **Actions** menu, select **Undeploy From**, then select **Management Agent**.
3. In the Undeploy Plug-in From Management Agent dialog, click **Add** and add the Management Agents from which you want to undeploy the plug-in. Click **Continue**. Then click **Undeploy**.
4. On the Confirmation dialog, click **Close**.

To monitor the undeployment operation, click **Show Status**.

Note: Undeploying a plug-in from Management Agent removes all the targets that were monitored by the plug-in.

Undeployment of a plug-in from the Management Agent restarts the Management Agent. The Management Agent does not monitor any target during downtime.

To undeploy a plug-in using EM CLI, use the following command:

```
$emcli undeploy_plugin_from_agent
-plugin="pluginId"
{-agent_names="agent1[;agent2...]" | -all_discovery_only_agents}
```

To undeploy all versions of `oracle.sysman.db2` plug-ins from all Management Agents where only Discovery Plug-ins are deployed, use the following command:

```
$emcli undeploy_plugin_from_agent -plugin=oracle.sysman.db2
-all_discovery_only_agents
```

18.4.12 Undeploying Plug-Ins from Oracle Management Service

To undeploy plug-ins from the OMS, follow the steps:

1. First, undeploy all plug-ins from all Management Agents. To do so, follow the steps mentioned in [Section 18.4.11](#).

2. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
3. On the Plug-ins page, select the plug-in you want to undeploy, and from the **Actions** menu, select **Undeploy From**, then select **Management Servers**.
4. In the Undeploy Plug-in From Management Server dialog, enter the Management Repository SYS password, and click **Continue**. Then click **Undeploy**.
5. On the Confirmation dialog, click **Close**.

To monitor the undeployment operation, click **Show Status**.

To undeploy a plug-in in EM CLI, use the following command:

```
$emcli undeploy_plugin_from_server
-plugin="plug-inId"
[-sys_password="sys_password"]
```

Note: When a metadata plug-in is undeployed/redeployed, it is recommended that you run the following command.

```
$emcli metric_control -command=flush_metadata_cache
```

The command should be run in each OMS environment instance.

If you want to undeploy only the plug-ins from the OMS, and not the entire Enterprise Manager system, then use the Plug-ins page within the Enterprise Manager Cloud Control Console. For instructions, see the *Oracle Enterprise Manager Cloud Control Administrator's Guide*. **Do NOT use runInstaller to undeploy only the plug-ins.**

18.5 Troubleshooting

This section contains information on troubleshooting plug-in related issues. The following sections are covered in this section:

- [Understanding Plug-In Homes](#)
- [Troubleshooting OMS Plug-In Deployment and Upgrade Issues](#)
 - [Troubleshooting OMS Plug-In Deployment Issues](#)
 - [Rollback and Resume OMS Plug-In Upgrade](#)
- [Troubleshooting Management Agent Plug-In Deployment and Upgrade Issues](#)
 - [Troubleshooting Management Agent Plug-In Deployment Issues](#)
 - [Troubleshooting Management Agent Plug-In Upgrade Issues](#)

18.5.1 Understanding Plug-In Homes

Plug-in homes are essentially Oracle homes that are dedicated for plug-ins. The plug-in home for plug-ins deployed to the OMS is different from the plug-in home for plug-ins deployed to the Management Agent. Each plug-in home will be its own ORACLE_HOME with a dependency on the OMS. Since plug-in homes are registered in the oraInventory, they should not be manually deleted or manipulated.

[Figure 18-12](#) shows the plug-in homes for plug-ins deployed to the OMS and the Management Agent, which is installed with the OMS, in the middleware home.

Figure 18-12 Plug-In Homes in Enterprise Manager Installation

```

<middleware_home>{
  |__wlsserver_10.3{
  |__jdk16{
  |__oms{
  |__plugins{
  |__agent{
    |__plugins{
    |__core{
      |__12.1.0.1.0{
        |__agent_inst{
        |__sbin{
        |__agentimage.properties{
      |__gc_inst{
      |__Oracle_WT{
      |__oracle_common{
      |__utils{
      |__logs{
      |__modules{
      |__user_projects{
      |__ocm.rsp{
      |__registry.dat{
      |__domain-registry.xml{
      |__registry.xml{

```

Figure 18-13 indicates the plug-in home for plug-ins deployed to a standalone Management Agent.

Figure 18-13 Plug-In Home for Standalone Management Agent

```

<agent_base_directory>
|__core
|  |__12.1.0.1.0
|  |__plugins{
|  |__agent_inst
|  |__sbin
|  |__plugins.txt
|  |__plugins.txt.status
|  |__agentimage.properties

```

18.5.2 Troubleshooting OMS Plug-In Deployment and Upgrade Issues

If the deployment of a new plug-in fails, the system automatically recovers. When the automatic recovery is complete, all OMS instances are started. If the upgrade of an existing plug-in fails, manual system recovery is required.

This section provides troubleshooting tips related to the following topics:

- [Troubleshooting OMS Plug-In Deployment Issues](#)
- [Rollback and Resume OMS Plug-In Upgrade](#)

18.5.2.1 Troubleshooting OMS Plug-In Deployment Issues

If plug-in deployment to the OMS fails, first check the details of the deployment, using the following commands.

- If the OMS is down, use the following command.

```
$emctl status oms -details
```

- If the OMS is running, use the following command.

```
$emcli get_plugin_deployment_status
```

Note: When the status of the OMS is displayed, review the log files that are displayed in the output.

It is recommended that you take a backup of the Repository in case of a failure in the Recovery.

Review the `pluginca` log file available in the following location. Use them to debug the issue, and if you raise a service request to Oracle Support, then make sure you append these to the service request.

```
$<OMS_HOME>/cfgtoollogs/pluginca/*
```

Note: When you install an additional OMS by cloning an existing, running OMS instance, the plug-in deployed to the source OMS are automatically carried over to the cloned OMS as well. Therefore, you do not have to redeploy the plug-ins on the cloned OMS.

In case of multi OMS environment, the `OMS_HOME` in the log file path indicates the root folder of the OMS where the failure occurs.

For information about installing an additional OMS, refer to the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.

18.5.2.2 Rollback and Resume OMS Plug-In Upgrade

If plug-in upgrade fails, then do one of the following:

- Raise a Service Request to find out if there is a possibility of recovering from the current Management Repository.
- Rollback to the latest backup of the Management Repository.
- If you have not taken a backup of the Management Repository, diagnose and resolve the issue that is causing the plug-in upgrade to fail.

Then, run the following command to resume the plug-in upgrade:

```
$<OMS_HOME>/bin/emctl resume_plugin_upgrade
```

This command automatically detects and resumes the last failed upgrade session. Once the plug-in upgrade succeeds, the OMS restarts automatically. The same deployment ID is updated with the current status of the operation. In case of a multi-OMS environment, the plug-in upgrade happens on the first OMS, and then on all other additional OMS instances.

- If flashback is enabled, the section number will be printed. The warning message ID archive is not enabled in the Repository DB.

18.5.3 Troubleshooting Management Agent Plug-In Deployment and Upgrade Issues

This section provides troubleshooting tips related to the following topics:

- [Troubleshooting Management Agent Plug-In Deployment Issues](#)
- [Troubleshooting Management Agent Plug-In Upgrade Issues](#)

18.5.3.1 Troubleshooting Management Agent Plug-In Deployment Issues

If plug-in deployment to the Management Agent fails, then review the log file available in the following locations.

`agent_inst/sysman/log/*`

`agent_inst/sysman/registry.xml`

`agent_inst/install/logs/*`

18.5.3.2 Troubleshooting Management Agent Plug-In Upgrade Issues

If plug-in upgrade fails, then review the log file available in the following locations.

- To review the log files using the UI, follow these steps.
 1. From the **Setup** menu, select **Extensibility**, and then select **Plug-ins**.
 2. Select the preferred plug-in, and review the information displayed in the **Recent Deployment Activities** tab.
 3. Click the link in the **Action** column for the preferred Management Agent. From the **Deployment Steps** tab, select the job name. Selecting the job name opens the job details wizard.
- The detailed logs for Management Agent upgrade and deployment are available at the following location.

`agent_inst/install/logs/agentplugindeploy_N.log`

In the aforementioned location, N refers to the internal ID. Check the latest log files in the location.

- While filing an SR, upload the following log files.

`agent_inst/install/logs/*`

`agent_inst/sysman/log/*`

`agent_inst/sysman/registry.xml`

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 ensure 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 Manager, 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](#)
- [Health Overview](#)
- [Repository](#)
- [Controlling and Configuring Management Agents](#)
- [Management Services](#)
- [Viewing Enterprise Manager Topology and Charts](#)

19.1 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. 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.
- 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).
- Upgrade Management Agents

19.2 Health Overview

The Health Overview provides a comprehensive overview of OMS and Repository operation and performance, and therefore allows you to view the overall health of your Enterprise Manager environment.

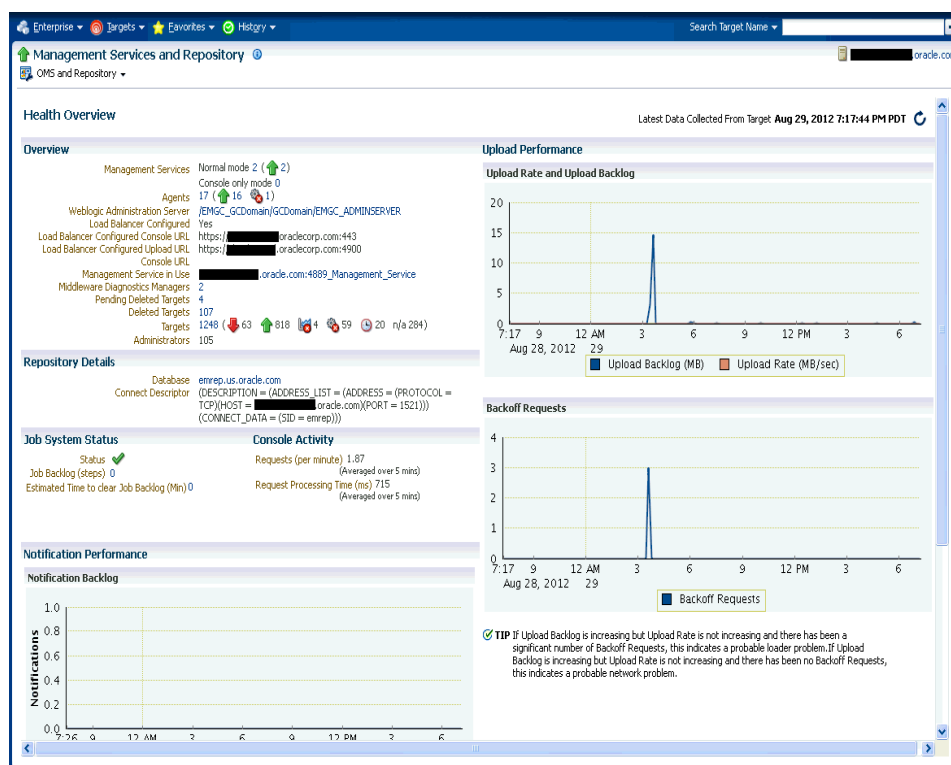
Accessing the Health Overview

From the **Setup** menu, select **Manage Cloud Control** and then **Health Overview**.

All major areas of Enterprise Manager are represented.

- **Overview:** Provides key information for active Management Services such as the Management Agents, the WebLogic Administration Server, total number of monitored targets, number of administrators, and server load balancer (SLB) upload and console URLs, provided SLB is configured. If configured, the SLB upload and console URLs are also displayed.
- **Repository Details:** Provides physical information about the Management Repository and the host on which the database is located. You can drill down into the database home page for more information and carry out administrative operations.
- **Job System Status:** Displays key operational parameters of the Enterprise Manager Job service. For detailed information, you can click on the status icon to drill down into the Enterprise Manager Job Service home page.
- **Console Activity:** Displays the overall load on the Enterprise Manager console through the average number of requests per minute and the average time required to process those requests.
- **Alerts:** Provides details on the metric errors recorded and when an alert was triggered. In-context links to Incident Manager are also provided.
- **Performance Charts:** Upload Backlog and Upload Rate, Backoff Requests, Notification backlog. You can drill down into any chart to view detailed metric information.

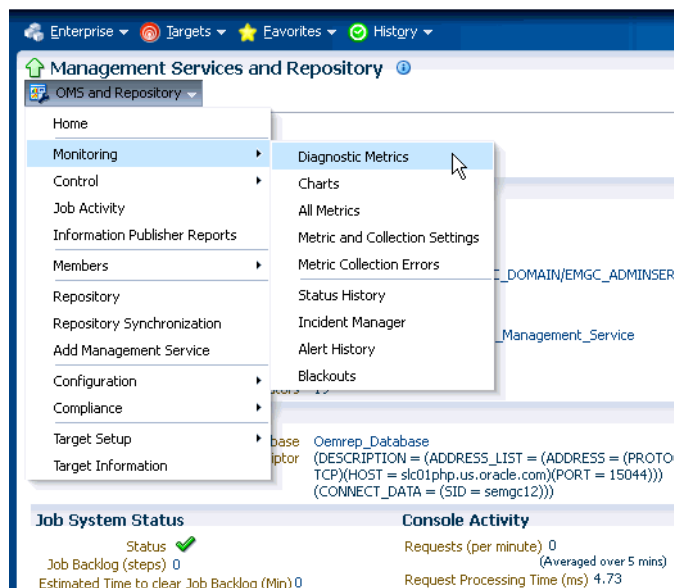
Figure 19–1 Health Overview Page



From this page, you can carry out all monitoring and management operations using the **OMS and Repository** menu.

Note: The Diagnostic Metrics page is intended for use by Oracle Support when diagnosing issues with the OMS. The page can be accessed by selecting **Monitoring** and then **Diagnostic Metrics** from the **OMS and Repository** menu.

Figure 19–2 Management Services and Repository Operations




19.3 Repository

The Repository 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 key Repository Details and Backlog in Repository Collection are provided. The Scheduler Status region provides the status of the scheduler and the number of Job Queue Processes.

Accessing the Repository Page

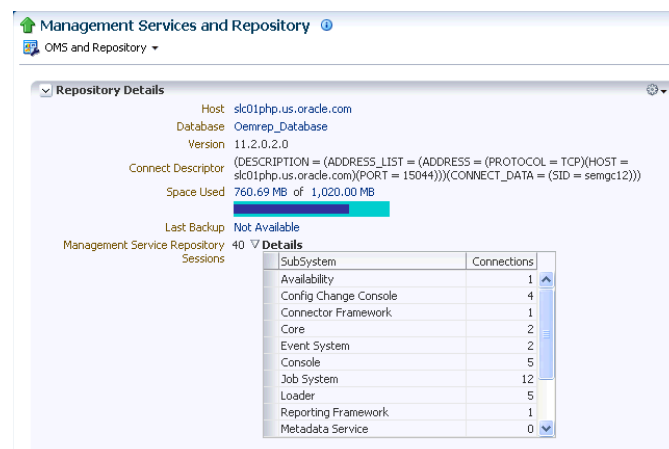
From the **Setup** menu, select **Manage Cloud Control** and then **Repository**.

Figure 19–3 Repository Page



From the **Repository Details** region, you can expand **Management Service Repository Sessions** details to view the exact number of repository connections per individual Enterprise Manager subcomponent such as the event system, console, job system, or connector framework.

Figure 19–4 Repository Sessions Per Subcomponent



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.

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).

Advanced Queuing (AQ) provides database-integrated message queuing functionality that allows application messages to be stored persistently and propagated between queues on different machines and databases. The Management Services AQ Status region list the details of DEQUEUE and ENQUEUE status for messages specific to the Management Services AQ.

19.4 Controlling and Configuring Management Agents

Beginning with Enterprise Manager Cloud Control 12c, controlling Management Agents can be performed directly from the Enterprise Manager console. This provides a central point where all Management Agents within your monitored environment can be compared, configured and controlled.

19.4.1 Manage Cloud Control Agents Page

The Agents page lists all Management Agents within your monitored environment. This page also includes misconfigured, blocked and both upgradable and non-upgradable Agents.

Accessing the Agent Page

From the **Setup** menu, select **Manage Cloud Control** and then **Agents**.

Figure 19–5 Agent List Page

Name	Secure Upload	Status	Incidents	Last Successful Load	Monitored Targets	Broken Targets
oracle.com:1838	Yes	Up	0	Jun 19, 2013 11:06:57 PM PDT	23	0

Misconfigured and Blocked Agents

A *misconfigured* Agent is an Agent that is not able to perform a heartbeat or upload data to the Oracle Management Service (OMS) due to invalid configuration or invalid data. Agent misconfiguration alerts are triggered by the following metrics:

- Consecutive metadata upload failure count

- Consecutive ping failure count
- Consecutive severity upload failure count
- OMS Agent time skew

If the Agent heartbeat or upload requests are failing consistently, and the problem cannot be resolved in a timely manner, you can manually *block* the Agent to prevent excessive load on the OMS. When you block an Agent, the OMS rejects all heartbeat or upload requests from the blocked Agent. However, even though blocked Agents continue to collect monitoring data, it will not be able to upload any alerts or metric data to the OMS. Once the Agent configuration problem is resolved, you must manually *unblock* the Agent to resume normal operation.

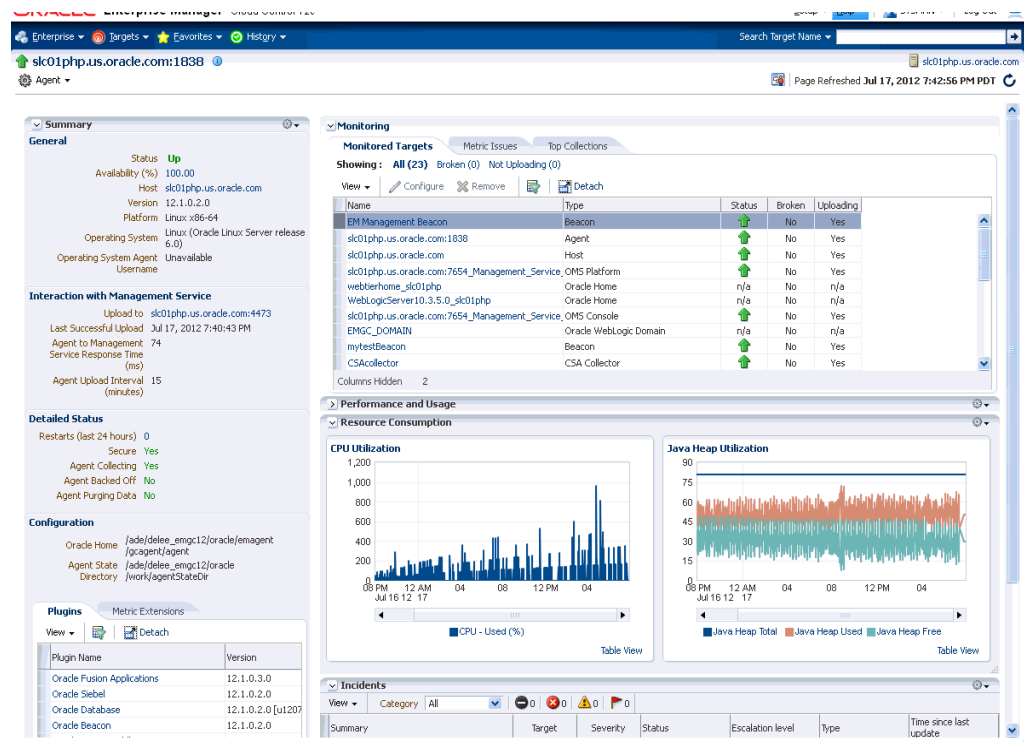
Note: Before unblocking the Agent, ensure that all issues related to Agent misconfiguration have been resolved.

From this page, you can also initiate the Agent upgrade process. For more information about upgrading Agents see "[Upgrading Multiple Management Agents](#)" on page 19-12.

19.4.2 Agent Home Page

The Agent home page provides details for a single Agent. This page also lets you drill down for more detailed information. You can access an Agent home page by clicking on a specific Agent from in the Agent list page or by selecting it from the All Targets page.

Figure 19–6 Agent Home Page



- The **Summary** region provides primary details of the Agent such as its status and availability. The Interaction with Management Service region provides details on

the communication between the OMS and the Agent and metric extensions and management plug-ins deployed in the Agent.

- The **Status** region provides further details on the Agent status such as the number of restarts, the action that the Agent is performing currently.
- The **Performance, Usage and Resource Consumption** charts provide further details on the Agent in graphical format.
- The **Incidents** region lists the incidents recorded for the Agent.
- The **Monitoring** region provides details on the targets that are being monitored by the Agent. You can filter targets in this region by All, Broken, and Not Uploading. Separate tabs within the Monitoring section display Metric Issues and Top Collections.

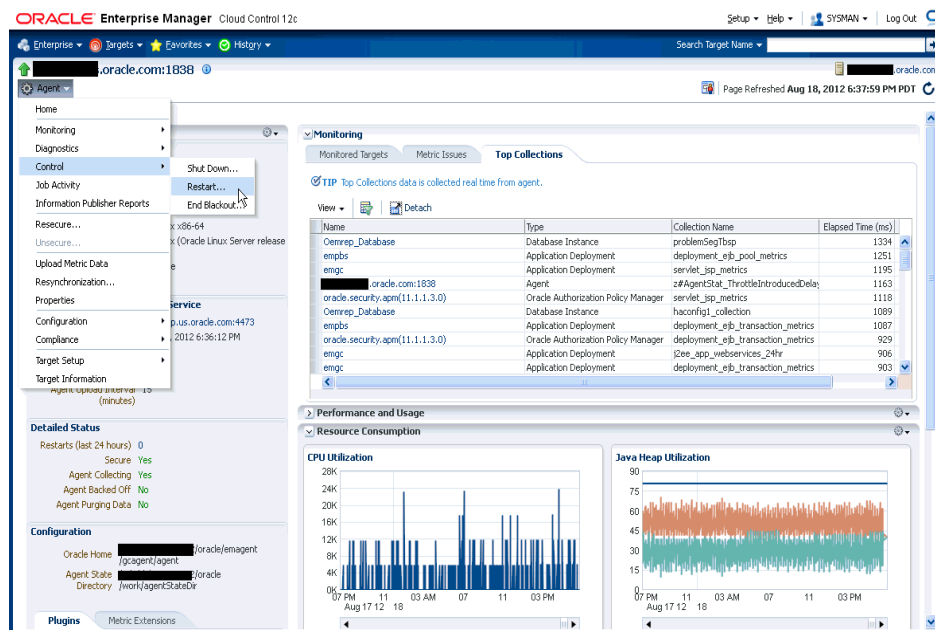
19.4.3 Controlling a Single Agent

Control operations for a single Agent can be performed on the Agent home page for that Agent.

1. Navigate to the desired 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 19–7 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 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.

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.

Agent Control When Using a Server Load Balancer

When choosing the Agent Secure/Resecure options in a multi-OMS environment with a server load balancer (SLB), the Agent will be secured/resecured against the SLB automatically without administrator intervention.

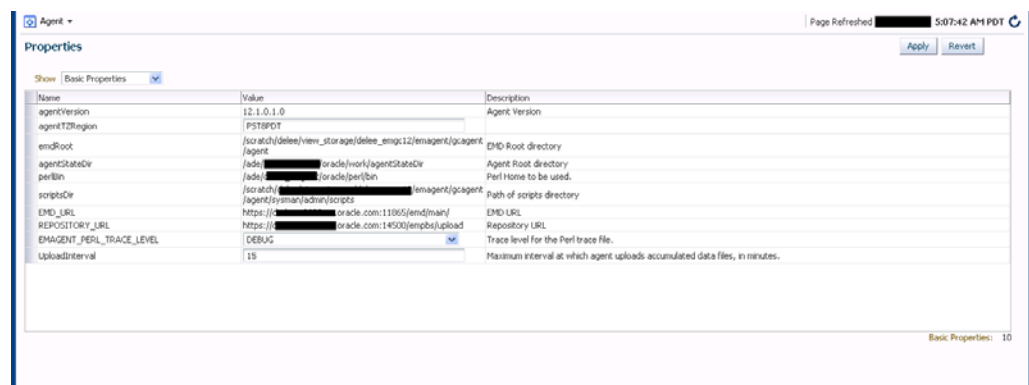
19.4.4 Configuring Single Management Agents

Configuration operations for a single Agent can be performed from the Agent home page. To access the Agent properties page:

1. Navigate to the desired 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 19–8 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.

19.4.5 Controlling Multiple Management Agents

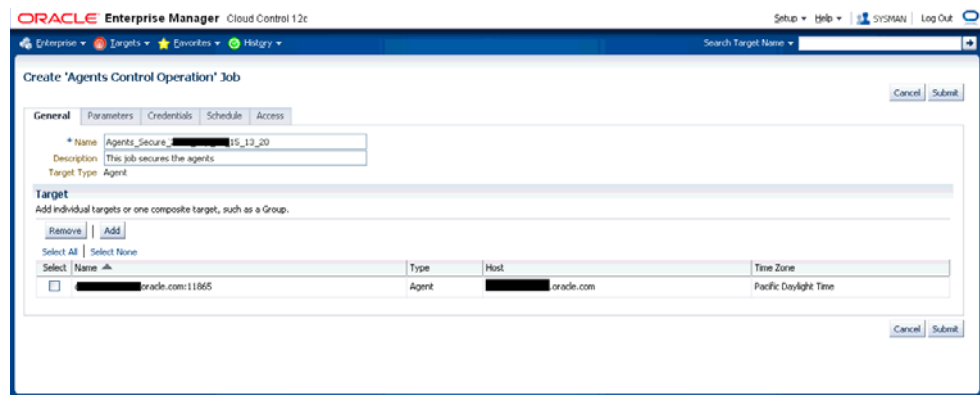
In order to perform control operations on multiple Management Agents, Enterprise Manager makes use of the Job system to automate repetitive tasks. Therefore, you must have Job privileges for controlling multiple Management Agents through a single action. To access

1. From the **Setup** menu, select **Manage Cloud Control and then Agents**. The Agent page displays.

- 2. Select multiple Management Agents from the list.
- 3. Click one of the control operation buttons (**Start Up/Shutdown/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 19–9 Multiple Agent Control Operation: Job Creation

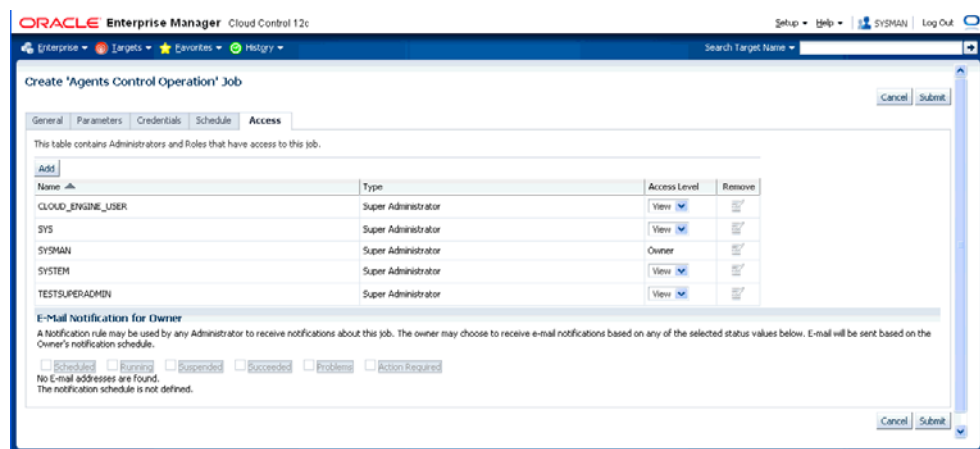


In the Jobs page, you can view the chosen Management Agents in Target section in the General tab. You can add more Management Agents by clicking the **Add** button. You then provide the parameters for the operation in the **Parameters** tab, if needed. The credentials must be specified in the **Credentials** 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 19–10 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.

19.4.6 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 **Manage Cloud Control** and then **Agents**. The Agent 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.

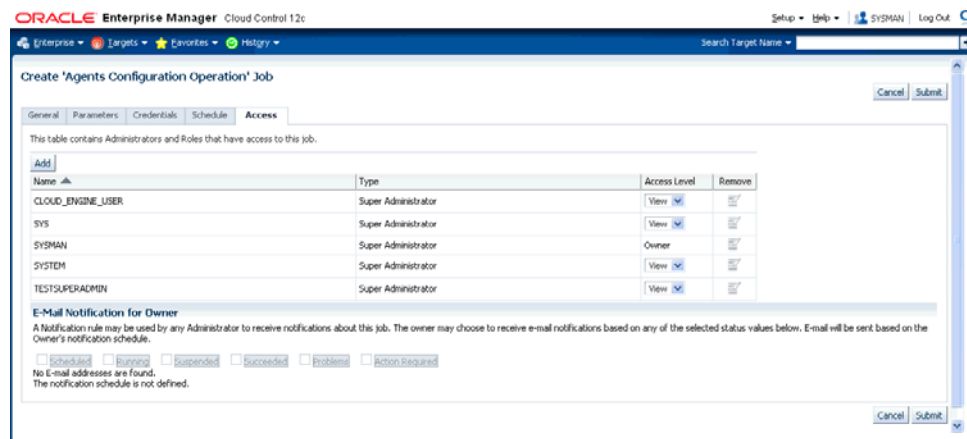
Figure 19-11 Agent Properties Page

The screenshot shows the Oracle Enterprise Manager Cloud Control interface. The main title is "Create 'Agents Configuration Operation' Job". Below the title are tabs for "General", "Parameters", "Credentials", "Schedule", and "Access". The "Parameters" tab is selected, displaying a table titled "Agent Configuration Properties".

Name	Value	Description
agentTZRegion		
CLASSPATH		Additional classpath used for launching agent.
agentJavaDefines		Additional java flags used for launching agent.
proxyHost		hostname of HTTP proxy used to connect to targets. Not used for upload to EM repository.
proxyPort		port number of HTTP proxy used to connect to targets. Not used for upload to EM repository.
dontProxyFor		comma-separated list of domains that should not use HTTP proxy when connecting to targets. Not used for upload to EM repository.
REPOSITORY_PROXYHOST		hostname of HTTP proxy used to connect to EM repository.
REPOSITORY_PROXYPORT		port number of HTTP proxy used to connect to EM repository.
REPOSITORY_PROXYUSER		username for an authenticated HTTP proxy used to connect to EM repository.
REPOSITORY_PROXYPWD		password for an authenticated HTTP proxy used to connect to EM repository.

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 19–12 Multi-Agent Configuration: Job Access

19.4.7 Upgrading Multiple Management Agents

When you upgrade to Enterprise Manager Cloud Control 12c Release 2, you essentially upgrade your Oracle Management Service (OMS) from 12c Release 1 (12.1.0.1) to 12c Release 2 (12.1.0.2), but not your target Oracle Management Agents (Management Agents). To mass-upgrade your Management Agents, access the Upgrade Agents page. To access this page, follow these steps:

1. From the **Setup** menu, select **Manage Cloud Control**, then select **Agents**.
2. Click **Upgradable**, then select the Management Agents you want to upgrade.
3. Click **Upgrade**.

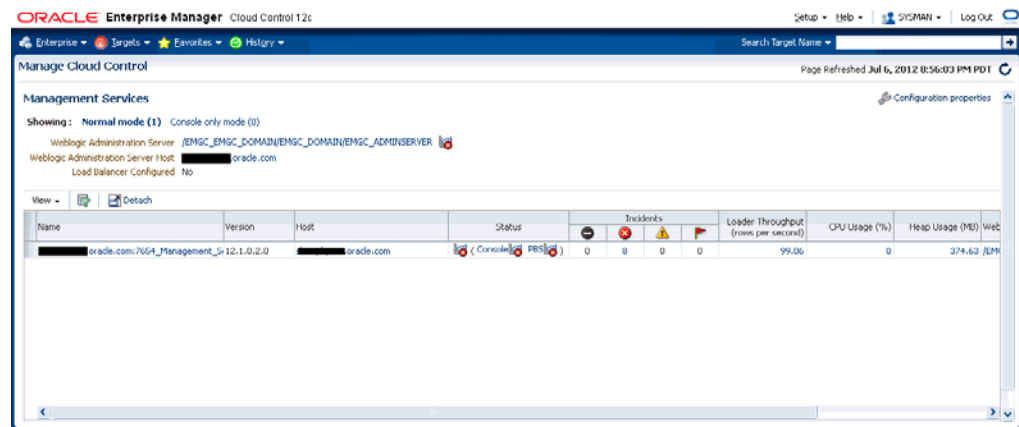
Alternatively, to access the Upgrade Agents page, from the **Setup** menu, select **Manage Cloud Control**, then select **Upgrade Agent**. For more information on upgrading Management Agents, refer Upgrading Oracle Management Agents.

19.5 Management Services

The Management Services page displays the list of Management Services, their status, incidents, the loader throughput, CPU usage, and the JVM memory usage metrics. In addition, the Management Services displayed can be filtered by Normal Mode, Console Only, PBS only and Standby Management Services.

Accessing the Management Services Page

From the **Setup** menu, select **Manage Cloud Control** and then **Management Services**.

Figure 19-13 Management Services Page

When an OMS is decommissioned, status will be shown as down and a delete button will be displayed, allowing you to remove the decommissioned OMS from Enterprise Manager.

By clicking on the individual Management Service, you can drill down to that OMS target home page for explicit information and status.

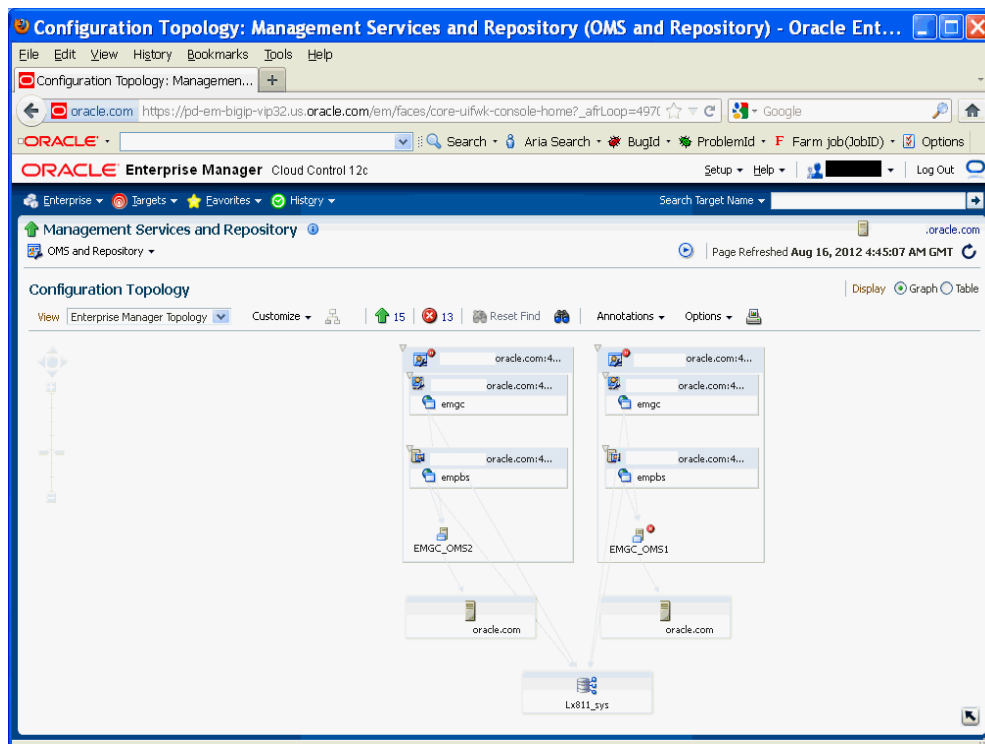
19.6 Viewing Enterprise Manager Topology and Charts

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, select **Manage Cloud Control** and then **Health Overview**.
2. Click on the **OMS and Repository** menu to display available operations that can be performed from this page.
3. Select **Members** and then **Topology**.

Figure 19–14 Enterprise Manager Topology

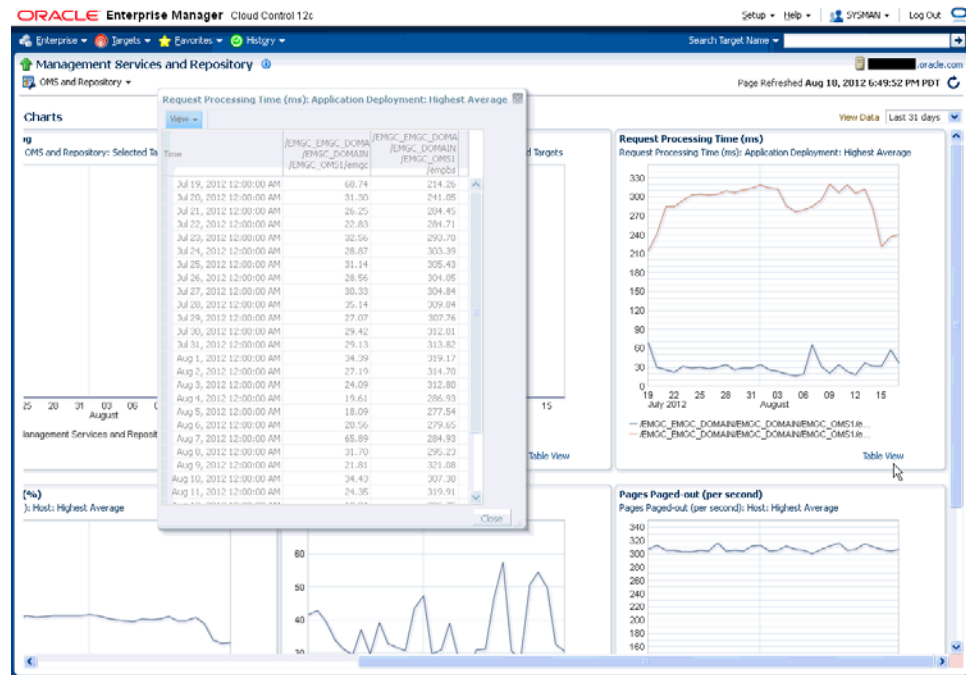
Enterprise Manager Charts

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 **Manage Cloud Control** and then **Health Overview**.
2. Click on the **OMS and Repository** menu to display available operations that can be performed from this page.
3. Select **Monitoring** and then **Charts**.

Figure 19-15 Enterprise Manager Charts



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](#)
- [Dropping and Recreating the Management Repository](#)
- [Troubleshooting Management Repository Creation Errors](#)
- [Cross Platform Enterprise Manager Repository Migration](#)

20.1 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 archive backups and 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

- Oracle recommends that you not use SQL Plan Management (SQL plan baselines and capture) with the Enterprise Manager Cloud Control repository. If you do need to use it for a specific problem, shut it off immediately after using. Issues with the Enterprise Manager Cloud Control repository may occur when using SQL Plan Management, such as very poor SQL performance using unverified plans, and deadlocks between SQL Plan Management capture and the Enterprise Manager security VPD.

20.2 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.

20.2.1 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 20-1](#).

Table 20–1 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 20–2](#).

Table 20–2 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

20.2.2 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.

20.2.3 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 20–3](#) 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 20–3 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

To modify the retention period for any of the above tables, execute the following command:

```
SQL> execute gc_interval_partition_mgr.set_retention('SYSMAN',
<table name>, <number of partitions to retain>);
```

Replace the <table name> by name of table as listed above. The API will allow you to change the number of partitions retained only.

For example, to modify the default retention time for the table EM_METRIC_VALUES from 7 partitions to 14 partitions, follow these steps:

1. Use SQL*Plus to connect to the repository database as the SYSMAN user.
2. Check the current value of the retention periods:

```
SQL> select table_name, partitions_retained
from em_int_partitioned_tables
where table_name in ('EM_METRIC_VALUES','EM_METRIC_VALUES_HOURLY','EM_METRIC_
VALUES_DAILY');
```

TABLE_NAME	PARTITIONS_RETAINED
EM_METRIC_VALUES	7
EM_METRIC_VALUES_HOURLY	32
EM_METRIC_VALUES_DAILY	12

3. To modify the default retention time for the table EM_METRIC_VALUES from 7 partitions to 14, execute the following command:

```
SQL> execute gc_interval_partition_mgr.set_
retention('SYSMAN', 'EM_METRIC_VALUES', 14);
```

4. Verify that the retention period has been modified:

```
SQL> select table_name, partitions_retained
from em_int_partitioned_tables
where table_name in ('EM_METRIC_VALUES','EM_METRIC_VALUES_HOURLY','EM_METRIC_
VALUES_DAILY');
```

TABLE_NAME	PARTITIONS_RETAINED
EM_METRIC_VALUES	14
EM_METRIC_VALUES_HOURLY	32
EM_METRIC_VALUES_DAILY	12

20.2.4 Modifying Data Retention Policies When Targets Are Deleted

By default, when you delete a target from the Cloud 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;
```

20.2.5 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**, then 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.

20.2.6 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.

```
SQL> 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.

```
SQL> 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.

```
SQL> 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:

```
SQL> SELECT value
FROM v$parameter
WHERE name='sessions';
```

```
SQL> 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:

```
SQL> 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.

20.3 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.

20.3.1 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
```

Note: Do not use the database version of the Repmanager script. It does not delete all components which will result in a failed re-installation.

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 dropall* indicates that you want to drop the Management Repository, MDS, OPSS, APM, and Schemas. If you use *drop*, the command drops only 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_DATA=(SERVICE_NAME=service_name)))"
-sys_password efkl34lmn -action dropall
```

See Also: "Establishing a Connection and Testing the Network" in the *Oracle Enterprise Manager Licensing Information* for more information about connecting to a database using connect descriptors.

20.3.2 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 a Connect Descriptor to Identify the Management Repository Database](#)

20.3.2.1 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 Enterprise Manager Licensing Information* 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_DATA=(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

20.4 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.

20.4.1 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 20-11.

20.4.2 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*.

20.4.3 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 20.3.1, "Dropping the Management Repository"](#)

If the RepManager script drops the repository successfully, try creating the Management Repository again.

If the RepManager `-action drop/dropall` fails for any reason, perform the following steps:

1. Apply the Bundle Patch to the 12c OMS home. Refer to My Oracle Support Note 1393173.1: Enterprise Manager Cloud Control Installation Instructions for Bundle Patch 1 and 12.1.0.2 Plugins for instructions.
2. Stop the OMS and verify that all the WLS / OMS processes have been stopped in the OMS home:

```
cd <OMS_HOME>/bin
emctl stop oms -all
```

Note: You should use the `-all` option so that the Admin Server is stopped as well

Verify that there are no WLS / OMS processes still running:

```
$ ps -ef | grep EMGC
$ ps -ef | grep java
```

3. Drop the repository objects using "Repmanager dropall" command:

```
cd <OMS_HOME>/sysman/admin/emdrep/bin
RepManager <database hostname> <database listener port> <database sid> -action
dropall -dbUser sys -dbPassword <sys user password> -dbRole sysdba -mwHome
<Middleware Home> -mwOraHome <Middleware Home> -oracleHome <OMS Home>
```

For example:

```
RepManager repomachine.domain 1521 orcl -action dropall -dbUser sys -dbPassword
oracle123 -dbRole sysdba -mwHome /home/oracle/Middleware
-mwOraHome /home/oracle/Middleware -oracleHome /home/oracle/Middleware/oms
```

4. Log in to the Repository Database as sys or any DBA user and verify that all the repository objects have been dropped:

```
SQL> select username,account_status from dba_users where username in ('SYSMAN',  
'SYSMAN_MDS','MGMT_VIEW','SYSMAN_BIP','SYSMAN_APM','BIP','SYSMAN_OPSS','SYSMAN_  
RO') ;
```

```
SQL> select owner,synonym_name from dba_synonyms where table_owner in  
( 'SYSMAN', 'SYSMAN_MDS','MGMT_VIEW','SYSMAN_BIP','SYSMAN_APM','BIP','SYSMAN_  
OPSS','SYSMAN_RO') ;
```

```
SQL> select tablespace_name from dba_tablespaces where tablespace_name like  
'MGMT%';
```

```
SQL> select comp_name from SCHEMA_VERSION_REGISTRY;
```

None of the above queries should return any rows. If any of the above queries return any rows, then raise an SR with Oracle Support.

Note: The above solution is applicable if the OMS is in working condition. If the OMS home is not available or not intact, raise an SR with Oracle Support.

20.5 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 the 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 raises the 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]

If the source and target is not on release 12c 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*.

20.5.1 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.

- The target database should meet all the pre-requisites for the Enterprise Manager Repository software requirements mentioned in the *Oracle Enterprise Manager Installation Guide*.
- If the source and target database are on release 10gR2 and higher rdbms versions, and provided they are meeting other prerequisites, cross platform transportable database migration 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.

20.5.2 Methodologies

The following sections discuss two methodologies for a repository migration:

- [Cross Platform Transportable Database](#)
- [Migration Using Physical Standby](#)

20.5.2.1 Cross Platform Transportable Database

Oracle's transportable database feature allows users to quickly move a user tablespace across Oracle databases. It is the most efficient way to move bulk data between databases. With the cross platform transportable database, you can transport tablespaces across platforms.

Cross platform transportable database allows a database to be migrated from one platform to another (use with Data Pump or Import/Export). The following set of steps for migration using Transportable Database can be used for migrations between same-endian platforms:

1. Verify whether migration is possible on the destination platform from v\$db_transportable_platform.

```
SQL> select platform_name from v$db_transportable_platform;
```

You may see a list of platforms similar to the list below:

```
Microsoft Windows IA (32-bit)
Linux IA (32-bit)
HP Tru64 UNIX
Linux IA (64-bit)
HP Open VMS
Microsoft Windows IA (64-bit)
Linux x86 64-bit
Microsoft Windows x86 64-bit
Solaris Operating System (x86)
HP IA Open VMS
```

Solaris Operating System (x86-64)

2. Verify that the external tables and files exist in the database.

```
SQL> set serveroutput on
SQL> declare x boolean;
       2 begin x := dbms_tdb.check_external;
       3 end;
       4 /
```

The following output results:

The following external tables exist in the database:

SH.SALES_TRANSACTIONS_EXT

The following directories exist in the database:

SYS.SUBDIR, SYS.SS_OE_XMLDIR, SYS.MEDIA_DIR, SYS.LOG_FILE_DIR,

SYS.DATA_FILE_DIR, SYS.XMLDIR, SYS.DATA_PUMP_DIR, SYS.ORACLE_OCM_CONFIG_DIR

The following BFILEs exist in the database:

PM.PRINT_MEDIA

Enter the following command:

```
SQL> select directory_path from dba_directories;
```

The following output results:

```
DIRECTORY_PATH
-----
-
/home/oracle/app/oracle/product/11.2.0/dbhome_1/demo/schema/order_
entry//2002/Sep
/home/oracle/app/oracle/product/11.2.0/dbhome_1/demo/schema/order_entry/
/home/oracle/app/oracle/product/11.2.0/dbhome_1/demo/schema/log/
/home/oracle/app/oracle/product/11.2.0/dbhome_1/demo/schema/sales_history/
/ade/b/1191423112/oracle/rdbms/xml
/home/oracle/app/oracle/product/11.2.0/dbhome_1/demo/schema/product_media/
/home/oracle/app/oracle/admin/orcl/dpdump/
/home/oracle/app/oracle/product/11.2.0/dbhome_1/ccr/state

8 rows selected.
```

Enter the following command:

```
SQL> select directory_path||'/'||location External_file_path
from dba_directories a, dba_external_locations b where
a.directory_name=b.directory_name;
```

The following output results:

```
EXTERNAL_FILE_PATH
-----
/home/oracle/app/oracle/product/11.2.0/dbhome_1/demo/schema/sales_
history//salelv3.dat
```

Enter the following command:

```
SQL> @tgt_get_bfile_dirs.sql
```

The following output results:

The following directories contain external files for BFILE columns

Copy the files within these directories to the same path on the target system

```
/home/oracle/app/oracle/product/11.2.0/dbhome_1/demo/schema/product_media/
```

There are 1 directories, 4 total BFILES

```
SQL> @tgt_get_bfiles.sql
External files for BFILE column AD_GRAPHIC in table PM.PRINT_MEDIA
/home/oracle/app/oracle/product/11.2.0/dbhome_1/demo/schema/product_
media//monitor.jpg
/home/oracle/app/oracle/product/11.2.0/dbhome_1/demo/schema/product_
media//mousepad.jpg
/home/oracle/app/oracle/product/11.2.0/dbhome_1/demo/schema/product_
media//keyboard.jpg
/home/oracle/app/oracle/product/11.2.0/dbhome_1/demo/schema/product_
media//modem.jpg
```

3. Stop the OMS.

```
emctl stop oms -all
```

Enter the following SQL commands:

```
SQL> shutdown immediate
Database closed.
Database dismounted.
ORACLE instance shut down.
```

```
SQL> startup mount;
ORACLE instance started.
```

```
Total System Global Area 1473089536 bytes
Fixed Size                  1336596 bytes
Variable Size               1124076268 bytes
Database Buffers            335544320 bytes
Redo Buffers                 12132352 bytes
Database mounted.
```

4. Open the database in read-only mode.

```
SQL> alter database open read only;
```

Enter the following SQL commands:

```
SQL> set serveroutput on
SQL> declare
  2  retcode boolean;
  3  begin
  4  retcode := dbms_tdb.check_db('Linux IA (64-bit)',dbms_tdb.skip_none);
  5  end;
  6  /

SQL> declare
  2  retcode boolean;
  3  begin
  4  retcode := dbms_tdb.check_db('Linux x86 64-bit',dbms_tdb.skip_none);
  5  end;
  6  /
```

PL/SQL procedure successfully completed.

5. Generate the RMAN conversion script.

```
[oracle]$ ./rman
Recovery Manager: Release 11.2.0.1.0 - Production on Fri dd-mm-yy 12:10:29 2012
Copyright (c) 1982, 2009, Oracle and/or its affiliates. All rights reserved.
RMAN> connect target /
```

```
connected to target database: ORCL (DBID=1308105793)
RMAN> convert database on target platform
2> convert script '/tmp/convert_mydb.rman'
3> transport script '/tmp/transport_mydb.sql'
4> new database 'mydb'
5> format '/tmp/mydb%U'
6> db_file_name_convert '/home/oracle/app/oracle/oradata/mydb/', '/tmp';

Starting conversion at source at dd-mm-yy
using target database control file instead of recovery catalog
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=135 device type=DISK

External table SH.SALES_TRANSACTIONS_EXT found in the database

Directory SYS.SUBDIR found in the database
Directory SYS.SS_OE_XMLDIR found in the database
Directory SYS.MEDIA_DIR found in the database
Directory SYS.LOG_FILE_DIR found in the database
Directory SYS.DATA_FILE_DIR found in the database
Directory SYS.XMLDIR found in the database
Directory SYS.DATA_PUMP_DIR found in the database
Directory SYS.ORACLE_OCM_CONFIG_DIR found in the database

BFILE PM.PRINT_MEDIA found in the database

User SYS with SYSDBA and SYSOPER privilege found in password file
channel ORA_DISK_1: starting to check datafiles
input datafile file number=00001
name=/home/oracle/app/oracle/oradata/orcl/system01.dbf
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00
channel ORA_DISK_1: starting to check datafiles
input datafile file number=00002
name=/home/oracle/app/oracle/oradata/orcl/sysaux01.dbf
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00
channel ORA_DISK_1: starting to check datafiles
input datafile file number=00007
name=/home/oracle/app/oracle/oradata/orcl/mgmt.dbf
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00
channel ORA_DISK_1: starting to check datafiles
input datafile file number=00003
name=/home/oracle/app/oracle/oradata/orcl/undotbs01.dbf
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00
channel ORA_DISK_1: starting to check datafiles
input datafile file number=00008
name=/home/oracle/app/oracle/oradata/orcl/mgmt_ad4j.dbf
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00
channel ORA_DISK_1: starting to check datafiles
input datafile file number=00005
name=/home/oracle/app/oracle/oradata/orcl/example01.dbf
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00
channel ORA_DISK_1: starting to check datafiles
input datafile file number=00006
name=/home/oracle/app/oracle/oradata/orcl/mgmt_depot.dbf
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00
channel ORA_DISK_1: starting to check datafiles
input datafile file number=00004
name=/home/oracle/app/oracle/oradata/orcl/users01.dbf
channel ORA_DISK_1: datafile checking complete, elapsed time: 00:00:00
Edit init.ora file /tmp/init_mydb00nbs6dl_1_0.ora. This PFILE will be used to
```

```

create the database on the target platform
Run RMAN script /tmp/convert_mydb.rman on target platform to convert datafiles
Run SQL script /tmp/transport_mydb.sql on the target platform to create
database
To recompile all PL/SQL modules, run utlirp.sql and utlrp.sql on the target
platform
To change the internal database identifier, use DBNEWID Utility
Finished conversion at source at dd-mm-yy

```

6. Copy all the required files to a temporary location and mount on the target machine.

```

convert_mydb.rman
init_mydb.ora
transport_mydb.sql (and other data files listed in rman script [convert_
mydb.rman] and redolog files)

```

7. Execute the scripts generated in Step 5 on the target machine.

The RMAN script contains convert datafile commands. The SQL script contains control file creation, invalidating objects, and recompiling objects. On the target machine, execute the following:

```

RMAN> connect target /
RMAN> @/home/oracle/migrate/convert_mydb.rman

```

8. Ensure the database is up and running and the database is registered with the listener.

```

RMAN> STARTUP NOMOUNT PFILE = '/home/oracle/migrate/init_
mydb00nbs6dl_1_0.ora'; database is already started

```

9. Start the OMS to ensure the admin server is up.

```

[oracle]$ ./emctl status oms
[oracle]$ ./emctl start oms

```

10. Stop the OMS.

```

[oracle]$ ./emctl stop oms

```

11. Update repository database connection details.

```

[oracle]$ ./emctl config oms -store_repos_details -repos_
conndesc "(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=evildead.idc.oracle.com) (PO
RT=1521))) (CONNECT_DATA=(SID=mydb)))" -repos_user SYSMAN
-repos_pwd Oracle123

```

If there are multiple Oracle Management Services in this environment, run this store_repos_details command on all of them.

12. Restart the OMS.

```

[oracle]$ ./emctl start oms
[oracle]$ ./emctl status oms

```

```

Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.
WebTier is Up
Oracle Management Server is Up

```

13. Relocate Management Services and the Repository target.

20.5.2.2 Migration Using Physical Standby

The following steps describe the process you can use to migrate a repository using Physical Standby. This method can be used when the source and target platforms are supported. See My Oracle Support Note:413484.1 for details of which platform combinations are supported.

1. Install the database ORACLE_HOME on the target machine. The binaries should be the same version as the source.

If the target machine is Windows, install and configure CYGWIN on the Windows box for Management Agent deployment.

2. Deploy the Management Agent to the target server.
3. Create a Physical Standby as described in the Data Guard documentation.
4. Configure the Data Guard broker as described in the Data Guard Broker documentation.
5. Shutdown the OMS.

```
emctl stop oms -all
```

6. Check the OMS connect descriptor.

```
./emctl config oms -list_repos_details
```

7. Switchover the database using dgmgrl.

Use the following commands:

```
DGMGRL> switchover to target_db
verify
show configuration
show database target_db
show database source_db
```

8. Start the OMS admin server.

```
emctl start oms -admin_only
```

9. Update connect descriptor to point to the Standby Database.

```
eemctl config oms -store_repos_details -repos_conndesc
"(DESC= )" -repos_user sysman

[oracle]$ ./emctl config oms -store_repos_details -repos_conndesc
"(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP) (HOST=sample2.us.company.com) (PORT=1521)) (
CONNECT_DATA=(SERVER=DEDICATED) (SERVICE_NAME=test_win)))" -repos_user sysman
Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
Copyright (c) 1996, 2012 Oracle Corporation. All rights reserved.
Enter Repository User's Password :
Successfully updated datasources and stored repository details in
Credential Store.
```

If there are multiple Oracle Management Services in this environment, run this store_repos_details command on all of them.

10. Stop all the Oracle Management Services.

```
emctl stop oms -all
```

11. Start the OMS.

```
emctl start oms
```

12. Relocate Oracle Management Services and the Repository.

```

emctl config emrep -agent <agent_name> -conn_desc

[oracle]$ ./emctl config emrep -agent sample2.us.company.com:3872 -conn_desc
"(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=sample2.us.company)(PORT=1521))(CONNECT_DATA=(SERVER=DEDICATED)(SERVICE_NAME=test_win)))"
Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
Copyright (c) 1996, 2012 Oracle Corporation. All rights reserved.
Please enter repository password:
Enter password :
Login successful
Moved all targets from sample.us.company.com:3872 to
sample2.us.company.com:3872
Command completed successfully!
Enter password :
Login successful
Moved all targets from sample.us.company.com:3872 to
sample2.us.company.com:3872
Command completed successfully!

```

20.5.3 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.

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.

This chapter contains the following sections:

- [Using Self Update](#)
- [Setting Up Self Update](#)
- [Applying an Update](#)
- [Accessing Informational Updates](#)
- [Acquiring or Updating Management Agent Software](#)

21.1 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.

21.1.1 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 updatable entity.

Entity Types That Can Be Updated

Examples of updatable 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 Prerequisite Checks
- Compliance Content
- Diagnostic Checks

21.2 Setting Up Self Update

Before you can use the Self Update feature, you must satisfy these prerequisites:

- My Oracle Support credentials have been set up using the SYSMAN user. 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 \(Optional\)](#)

21.2.1 Setting Up Enterprise Manager Self Update Mode

In order to set up or modify the Enterprise Manager Self Update feature, you must have Enterprise Manager Super Administrator privileges.

1. Log in 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 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 and schedule updates from the Enterprise Manager Update Store.

21.2.2 Assigning Self Update Privileges to Users

Enterprise Manager administrators must have the requisite privileges to use the Self Update feature. The Enterprise Manager Super Administrator must assign the following Self Update roles/privileges to these administrators:

- *View any Enterprise Manager Update*—User can view the Self Update console and can monitor the status of download and apply jobs.

- *Self Update Administrator*—User can schedule download and apply jobs. User can also suppress/unsuppress updates. This privilege implicitly contains the *View any Enterprise Manager Update* privilege.
- *EM_INFRASTRUCTURE_ADMIN*—User can perform all self update operations. This role implicitly contains the *Self Update Administrator* privilege.

By default, the Super Administrator will be granted EM_INFRASTRUCTURE_ADMIN privilege.

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.

21.2.3 Setting Up the Software Library

The Software Library 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. In the context of applying updates, it is the "local store" that entities are downloaded to before deployment.

If the Software Library is not already set up in your environment, see Chapter 8, "Configuring Software Library," for instructions on the various ways you can configure the Software Library.

21.2.4 Setting Up the EM CLI Utility (Optional)

If you plan to apply software updates in offline mode, 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
```

For example:

```
https://emcc_host:emcc_port/em/console/emcli/download
```

21.3 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](#)

21.3.1 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 21.3.2, "Applying an Update in Offline Mode"](#) for details.

1. From the **Setup** menu, select **Extensibility**, then select **Self Update**.
2. Click **Check Updates** to submit a job to check for new updates from Oracle. Click **OK** to close the confirmation message.
3. When the job completes, select the desired entity type, then select **Open** from the **Actions** 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.

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.

21.3.2 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 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 that 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 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_catalog` command to import the file into the Oracle Management Service instance or the Management Agent you want to update.
7. Review the update from Self Update Home and click **Download** in the **Actions** menu. A message displays with a URL and instructions.
8. Click **Apply** in the **Actions** menu to apply the update.

21.4 Accessing Informational Updates

The Self Update feature also serves as a news feed, providing new product announcements, news stories, industry updates, and any number of other items of interest to the Oracle community. These informational updates occur on an ad hoc basis and typically include useful links where you can obtain additional information and download items.

1. From the **Setup** menu, select **Extensibility**, then select **Self Update**.
2. On the Self Update page, click the **Informational Updates** link at the top-right corner, as shown in [Figure 21-1](#). The link includes the number of new updates. A number appears only if there are new (unread) updates.

Figure 21-1 Informational Updates Link on the Self Update Home Page



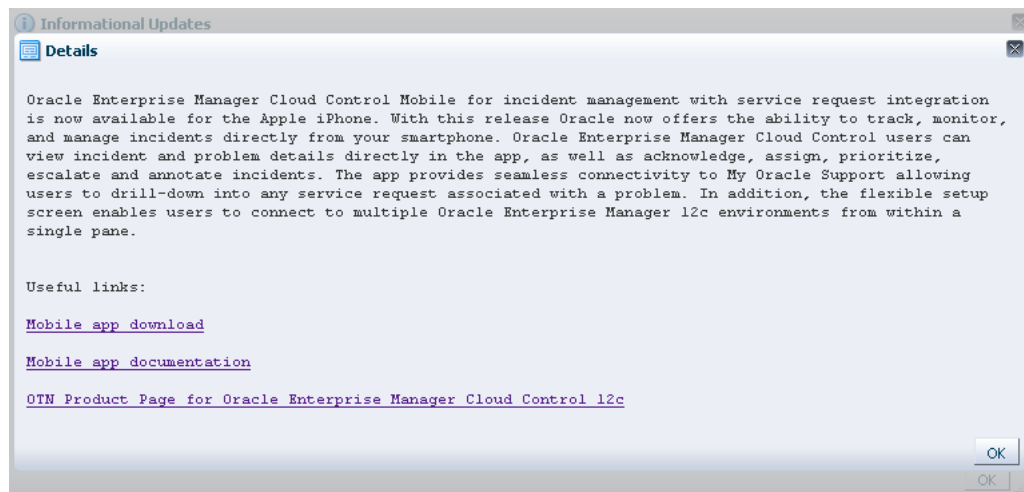
The Informational Updates dialog opens.

3. Select an update notification in the table and click **Details**.

A popup appears describing the new product and listing applicable links.

Figure 21–2 shows the informational update announcing availability of Enterprise Manager Cloud Control Mobile, a new iPhone app that enables you to connect to Enterprise Manager remotely. Notice in this case that the announcement includes a link to iTunes where you can go to download the app.

Figure 21–2 Informational Update for Mobile App



4. Click **OK** to close the details display and return to the table of announcements.

By default, the table displays only unread announcements. You can choose to display all or only read announcements. You can also toggle selected items between read and unread states. Note that if you mark an item as read, you are doing so for all users. A warning to this effect appears.

21.5 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 Add Host Targets wizard.

For instructions on obtaining Management Agent software in both online and offline modes, see the section "Meeting Management Agent Software Prerequisites" in the *Oracle® Enterprise Manager Cloud Control Basic Installation Guide*.

Patching Enterprise Manager Agents

This chapter shows how you can patch the Enterprise Manager Agents. There are two approaches to apply the Agent patches: automated approach, and the manual approach. Oracle recommends using the patch plans (automated approach) for patching as 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:

- [Overview](#)
- [Automated Agent Patching \(Recommended Option\)](#)
- [Manual Agent Patching](#)

Note: Oracle recommends using the Automated Patching mechanism using the Patch Plans available within Enterprise Manager Console to carry out all your patching operations. For a large scale deployments, you can use EM CLI. For information about patching using EM CLI, see *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

If you are unable to patch using patch plans, or using the command line interface, you can use the Manual Patching approach, though it is not recommended.

22.1 Overview

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.

Patch plans are used to patch the Agent targets. Patch Plans can be accessed using GUI and also EMCLI. For information about accessing the patch plans, see [Section 22.2.1](#).

In Enterprise Manager 12c, 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.

22.2 Automated Agent Patching (Recommended Option)

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).

Note: Oracle recommends that you set preferred credentials for all the Management Agent targets before beginning the patching process.

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.

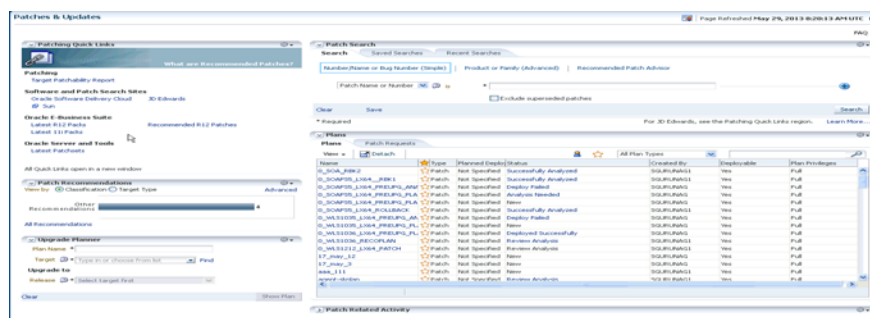
This section includes the following:

- [Accessing Patch Plans](#)
- [Viewing Patch Recommendations](#)
- [Searching Patches](#)
- [Applying Management Agent Patches](#)
- [Verifying the Applied Agent Patches](#)
- [Validating Agent Patch Errors](#)

22.2.1 Accessing Patch Plans

To access the Patch plans using GUI, follow these navigation steps: 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.

22.2.2 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.
-

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.

22.2.3 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.

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.

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.

22.2.4 Applying Management Agent Patches

To apply a Management Agent patch using patch plans, follow these steps:

Note: Applying the core agent patch, or the plug-in agent patch follows the same patching process (as listed in this section.)

For a large scale deployments, you can use EM CLI. For information about patching using EM CLI, see *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

Ensure that the patches selected have the same platform, as the targets being patched. For example, Linux x86 patches can only be applied on 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 [Section 22.2.2](#)

- 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 [Section 22.2.3](#)

3. Select one or more patches, and from the **Add to a Plan** menu, select one of the following options:

- If you select **Add to New**, then the Create New Plan dialog box is displayed. Enter a unique **Plan Name**.

From the **Select a target** menu, you can perform on of the following operations:

Select **Agent**, and click search icon. From the list of agents displayed in the table, you can either select a single target or multiple targets as follows:

- Provide the target name in the search field to select a particular target.
- To add multiple target to the plan, keep the Control key pressed, and select the desired targets.

Select a **Group**, click search icon, and then select the desired group. 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.

Click **Create Plan**

Name	Type	Oracle Hc
slc04zbn.us.oracle.com:387 Agent		/scratch/
slc04zbn.us.oracle.com:387 Agent		/scratch/

The patch selected along with the associated target gets added to the plan.

- If you select **Add to Existing**, then Add Patch To Plan dialog box is displayed. The actual process of updating the plan is a two step process as follows:

Step1: Select Plan

Click an existing plan to select it, then click **Select Targets**.

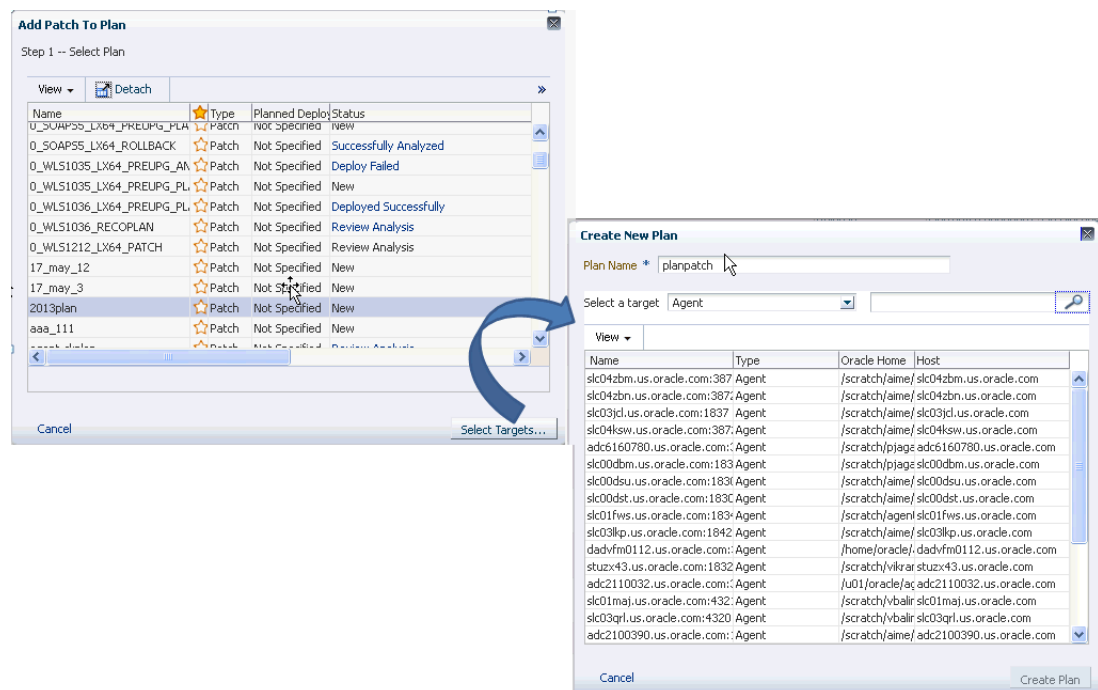
Step2: Select Targets to Plan

To search and select targets, follow one of these approaches:

- If you know the target name, 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.
- From the **Select a target** menu, select **Agent**, and click the search icon. From the list of agents displayed in the table, select one or more targets and click **Add Patch To Plan**. You can also perform a group search by providing the group name in the search field to add all the targets in that particular group to the plan.


After selecting the targets, click **Create Plan**. The selected patches and associated targets are added to the plan after validating for conflicts.

Figure 22–1 Add Patch To Plan



4. 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 or Cluster-Agents, then you may see a warning message **Issues with Adding Patch**.



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.

However, if there is a mismatch between the platform of the patch selected, and that of the target, you may see one of the following warnings:

- a. **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.
- b. **Platform Mismatch:** This warning message appears when the platform of the patch does not match the target platform. All the other homogeneous targets get added to the plan, however, the mismatched target is ignored, and the plan is created without this target. Click **View Plan** to see the details.

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

5. 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.

6. 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.

7. Click **Next**.

8. 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**.

9. Click **Next**.

10. On the Deployment Options page, in the Where to stage section select **Yes** to retain the default location (%emd_emstagedir%) available on the target machine. If you select **No**, then you must provide a stage location to pre-stage the patches. Ensure that you provide a mounted file system or an NFS shared location for staging the patches. Table 22–1 describes an example of the directory structure you must follow for staging the Agent patches and OPatch patches.

Table 22–1 Pre-Stage Location Example

Patch Type	Stage Location	Where to unzip the patch
Agent Patches	/scratch/<user_name>/	Create a directory structure as follows: /scratch/<user_name>/<patch_number>/ Unzip the patch inside the above location
OPatch	/scratch/<user_name>/	If you want to upgrade the OPatch component before patching, then ensure that OPatch is also available at the same folder level. For example: /scratch/<user_name>/6880880 Unzip the OPatch downloaded from ARU inside the above location

In the Oracle Home Credentials section, select **Use** to use the existing Oracle Home Preferred Credentials that you have set. Oracle recommends using preferred credentials to apply patches to Management Agent home, as the preferred credentials can be reused in any patch plans, and saves you the trouble of setting it multiple times. On the other hand if you use override option, you need to specify the named credentials for each patch plan for the same agent target, which can be cumbersome.

For more information about setting Preferred credentials using Cloud Control, see *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*. For setting Host Preferred Credentials Using SSH Key Credentials, see the *Enterprise Manager Cloud Control Security Guide*.

Note: If you have not set the Normal and Privileged Oracle Home Credentials for your target, then the plan becomes non-deployable, which means that you will not be able to proceed to analyze or deploy your plan.

Also, note that the Named Credentials of type SSH Key Credentials cannot be set as the preferred credentials for the *Normal Host Credentials* or *Privileged Host Credentials* credential set for Oracle Home targets.

To Override the existing credentials, you must specify if you would like to Override the credentials for each of the Oracle home targets in your plan, or all the Oracle home targets, to do so select Each Home or All homes from the **Specify For** menu. Click Edit to edit the **Normal Oracle Home Credentials** and **Privileged Oracle Home Credentials**, and click **Validate Credentials** to verify the accuracy of the credentials provided.

In the Customization section, you can choose one of the following options to resolve a patch conflict:

- **Stop at Conflicts** - This option stops the plan from being deployed when a conflict is detected.
- **Force Apply** - This option will rollback the conflicting patches in the Oracle Home and forcefully apply the incoming patches.
- **Skip Conflicts** - This option will apply only the non-conflicting patches and skips the conflicting ones.

In the Rollback section, ensure that the **Plan deployment rolls back the patches in this plan** option is not selected if you would like to apply the patches. If you select

the rollback option, it will mean that the plan is not for applying the patches, but instead for rolling back patches. Therefore, only if you want to rollback the selected patches from the selected targets, this rollback option should be selected.

If you do not have the latest version of OPatch patch on the targets to be patched, then you can enable the **OPatch Upgrade** option in the OPatch Upgrade section.

11. Click Next.

- 12.** 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.

Note: Alternately, you can directly go to the Review and Deploy page to deploy the Agent Patches without analyzing the plan. If you do so, a Deploy Job is submitted which internally analyzes the plan, and deploys it on successful analysis. If there are any issues in the Analysis phase, then the *split plan feature is triggered*, in which a plan is split into two plans, one plan with the failed targets and another plan with the successful targets. The successful plan is available for deployment, while the failed plan can be reanalyzed and deployed separately.

To track the progress of the job, click **Show Detailed Results**. 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, you must contact Oracle Support.

See Also: For more information about the common errors during the Validation phase, see [Section 22.2.6](#)

13. Click Next.

- 14.** On the Review & Deploy page, review the details you have provided for the patch plan, then click **Deploy**.

Important: A new split plan feature has been introduced with Enterprise Manager 12.1.0.3 patch set release 2. According to this feature, upon failure, a plan is split into two plans, one plan with the failed targets, and another plan with the successful targets. The successful plan is available for deployment, while the failed plan can be reanalyzed and deployed separately. The plan details are provided in the validations page, or review and deploy page.

The plan may fail due to several validation errors. For more information about these errors, see [Section 22.2.6](#).

For an example that captures split plan feature, see [Figure 22-2](#).

A job is submitted, to track the progress of the job, click **Show Detailed Results**.

Note: For a demonstration about how to apply patches on Enterprise Manager 12c Agents using Cloud Control Console, see *My Oracle Support* note 1359221.1.

22.2.5 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.

In the Agent home page under Summary section, in the Configuration sub-section, click **Oracle Home and Patch Details** to view all the jobs that have run on the Oracle Home Target of that Agent.

4. In the Patch Advisories region, select **Patches Applied** tab to verify all the patches that have been successfully applied on the target.

22.2.6 Validating Agent Patch Errors

Here are some of the errors that you may typically see during the Validation phase of patch plans:

- [Oracle Home Credentials Not Set](#)
- [Target Is Down](#)
- [Conflict Check Analysis](#)
- [User Not a SuperUser](#)
- [Patch Not Staged or Found](#)

Note: For all the errors mentioned in this section. Oracle recommends contacting the customer support to resolve the error.

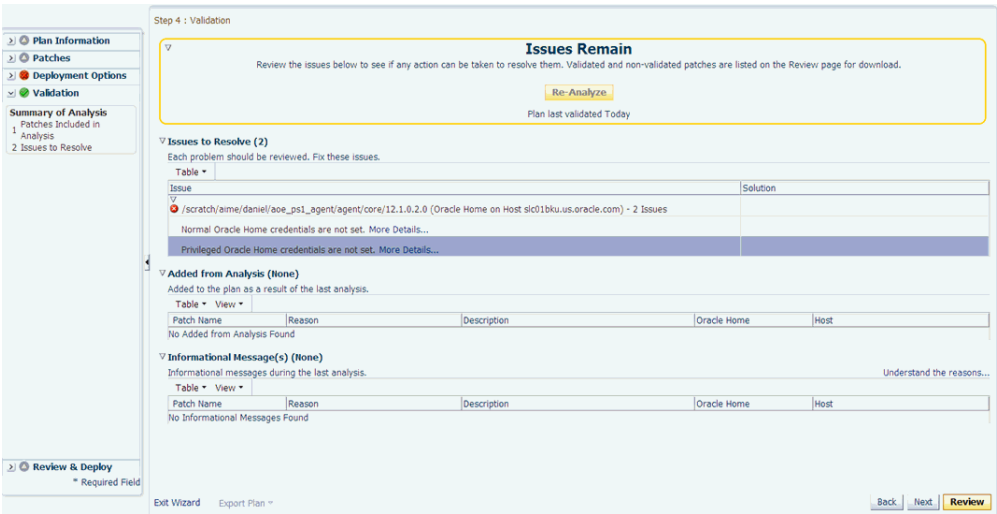
- **Oracle Home Credentials Not Set**

This error occurs when the Oracle Home credentials like Privileged Oracle Home credentials or the Normal Oracle Home credentials are not set.

Note: This error occurs only when you submit the patch plan using EMCLI. Using Enterprise Manager Cloud Control interface, you will not be allowed to proceed to validate and deploy the patch plan unless you set the credentials.

Note:

For example if you apply Linux x86 patch on a Linux x64 target, then the plan fails.



Workaround: You can use the Deployment Options page available in your patch plan to set the Privileged Oracle Home credentials or the Normal Oracle Home credentials.

■ **Target Is Down**

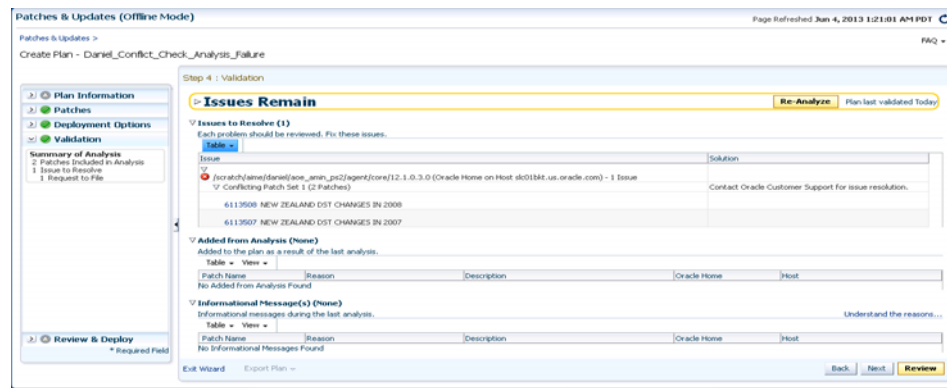
This error occurs when the target added for patching is not up and running.



Workaround: Start the target, and Re-Analyze the plan.

■ **Conflict Check Analysis**

This error occurs when there is a conflict in the patches added.



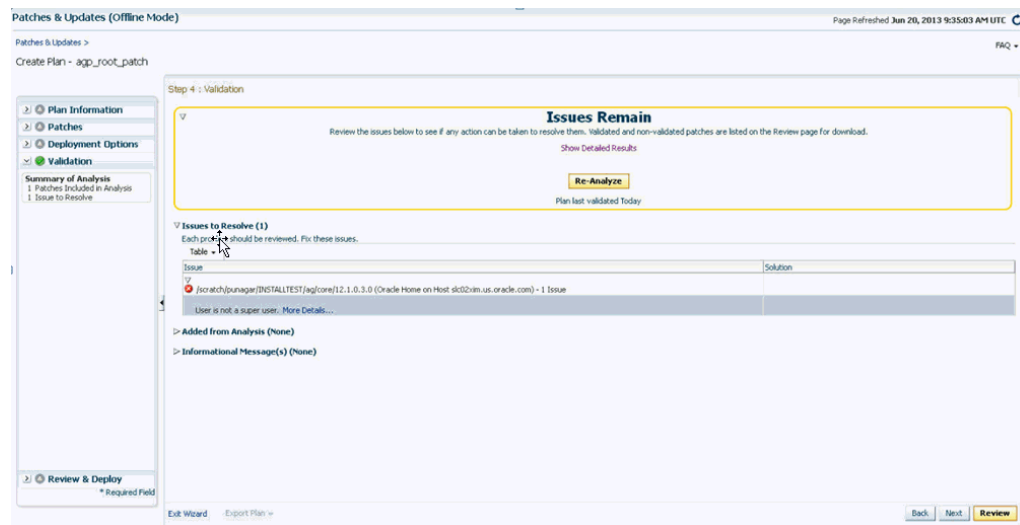
Workaround:

Approach 1: Contact Support to get a merged patch.

Approach 2: Choose the Advanced OPatch Options to force apply the patch. But choosing this option and applying the patch will make you loose earlier patch changes.

■ User Not a SuperUser

This error occurs when the user running the patch plan does not have root access.



Workaround:

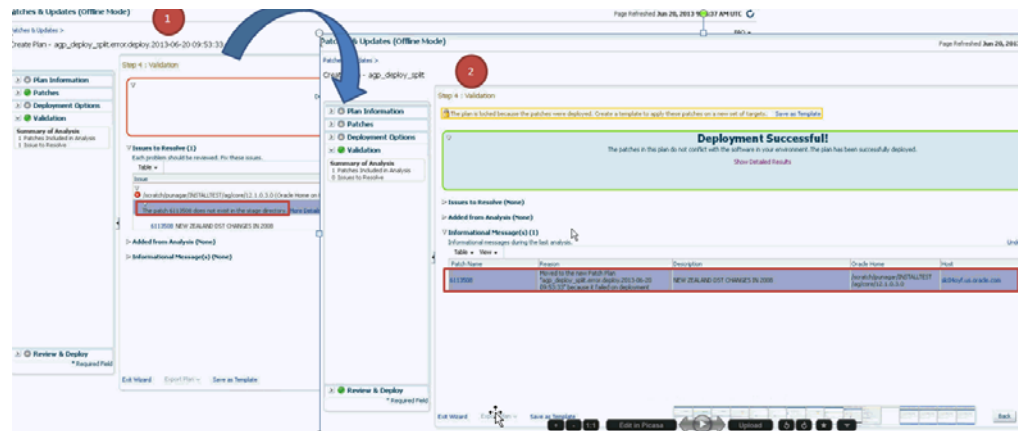
Step 1: Create a new credential with root access.

Step 2: Check if the privilege delegation has been done for that host. For more information about privilege delegation, see *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

Step 3: Re-analyze the plan

■ Patch Not Staged or Found

This error occurs when the patch is not present in the stage location.

Figure 22–2 Split Plan Feature (with patch not found error)**Split Plan Feature:**

The following graphic Uses the Split Plan Feature Introduced in Enterprise Manager 12.1.0.3. Basically, page 1 displays the error with which the plan fails. But, since the plan consisted of two patches, one of which was successful, the split plan feature enabled. Because of this, the successful targets are deployed. Page 2 essentially captures how you need to create a new patch plan to stage and deploy the patch appropriately for the plan to be successful.

Workaround

Ensure that the patch is available in the stage location, and then re-analyze the plan.

22.3 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.

Patching the OMS and Repository

Important: The patching methodology discussed in this chapter can only be used with Enterprise Manager Cloud Control Release 12.1.0.3 and later.

OPatchauto (introduced with version 11.1 of the OPatch utility) automates the patching process by generating custom patching instructions based your particular environment and then automatically applies the patch.

This chapter covers the following topics:

- [OPatch Automation](#)
- [Required OPatchauto Parameters](#)
- [Prerequisites for Running OPatchauto](#)
- [Using OPatchauto](#)

23.1 OPatch Automation

With OPatchauto, you can automatically patch a typical OMS configuration (core, plug-in homes) with minimal intervention.

OPatchauto performs many of the pre-patch checks such as:

- Configuration-based prerequisite checks
- Patch-based binary prerequisite checks

OPatchauto performs end-to-end configuration patching. Configuration patching is the process of patching a target based on its configuration. By incorporating the configuration information into the patch process, OPatchauto is able to simplify patching tasks by automating most of the steps.

23.1.1 Supported OMS Configurations and OPatchauto Patchability

- Single OMS – OMS application that runs from a single OMS instance of the system. OPatchauto performs patching and deployment operations
- Multiple OMS – OMS applications that run on two or more machines. The OMSs are connected by the Oracle WebLogic domain and separate managed servers. There is a one-to-one mapping between the managed servers and the separate OMS bits residing on a single machine. OPatchauto only provides context-sensitive steps (text and HTML) that the administrator will need to carry out manually.

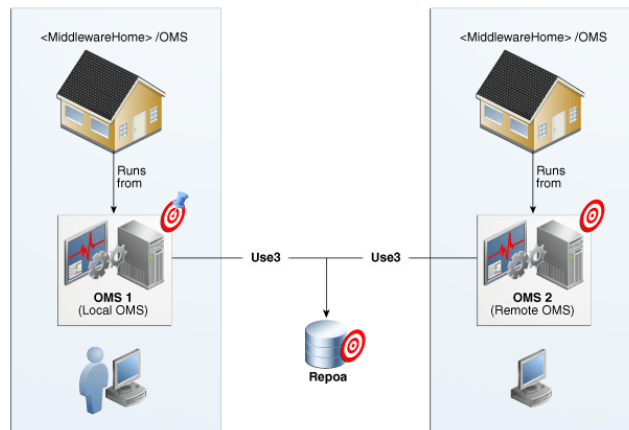
- Standby OMS – A HA (high availability) OMS configuration of Enterprise Manager which permits servicing requests when the primary OMS is down. When configured in this way, the standby OMS is managed by a different domain (managed by Oracle WebLogic Administration Server). OPatchauto only provides context-sensitive steps (text and HTML) that the administrator will need to carry out manually.
- Single Instance Database or Real Application Cluster - shared or Real Application Cluster (RAC)

Example: Multi-OMS System

The following figure illustrates a multi-OMS deployment. The following terms are used:

- *Administrator*: Person installing patches to the OMS core and plug-in homes.
- *Local OMS*: OMS instance on which the administrator runs OPatchauto.
- *Remote OMS*: OMS instances on other machines (within the same OMS domain as the local OMS) where the administrator has not started any patching operations.

Figure 23–1 Simple Multi-OMS System



For a single OMS system (primary), OPatchauto will execute the patching steps. For multi-OMS systems and standby OMS systems, OPatchauto will generate customized patching instructions/commands for the environment in text and HTML formats. Administrators can execute these instructions to patch the various OMSs.

23.1.2 OUI Inventory Configurations

Apart from the target (or) instance-based configurations, OPatchauto utilizes installation configuration relationships established in the Oracle Universal Installer (OUI) inventory as core and plug-in Oracle Homes. A typical OMS 12c home from the OUI inventory is organized as follows:

```
<Middleware Home>
|___platform home
|___plugin homes
|   |___oracle.sysman.db.oms.plugin_12.1.0.0.0
|   |___oracle.sysman.emas.oms.plugin_12.1.0.0.0
|   |___oracle.sysman.mos.oms.plugin_12.1.0.0.0
|   .
|   .
```

23.1.3 Supported Patch Format

Beginning with Enterprise Manager 12.1.0.3, Enterprise Manager patches have been converted to a *System patch* format in order to support patch automation.

What is a System Patch?

A System patch contains several sub-patches whose locations are determined by a file called *bundle.xml* in the top level directory of the patch. The sub-patches are intended for different sub-systems of a system that correspond with the OMS core and plug-in home organization.

A typical System patch format is organized as follows:

```
<System patch location - directory>
|_____ Readme.txt (or) Readme.html
        bundle.xml
        automation
        |_____ apply_automation.xml
                rollback_automation.xml
        Sub-patch1
                |_____ etc/config/inventory.xml
                |_____ etc/config/actions.xml
                |_____ files/Subpatch1 'payload'
        Sub-patch2
                |_____ etc/config/inventory.xml
                |_____ etc/config/actions.xml
                |_____ files/Subpatch1 'payload'
```

Notes:

- For Enterprise Manager release 12.1.0.2 or below, OPatchauto is not supported for the released one-off patches. For these older releases, you must use OPatch and follow the patch README instructions.
- OPatchauto and System patches are only supported by Enterprise Manager release 12.1.0.3 and above.

23.1.4 Supported Patching Methodologies

OPatchauto supports rolling mode only for System patches without any automation (binary-only patching through OPatchauto). For all other artifacts (MRS, SQL), OPatchAuto only supports complete system downtime patching operations.

Refer to the patch README for the explicit information on supported patching methodologies.

23.2 Required OPatchauto Parameters

OPatchauto for the Enterprise Manager OMS will prompt for the following input parameters when performing patching operations. These parameters were determined at the time of Enterprise Manager installation.

- Oracle WebLogic Admin Sever URL & port number
- Oracle WebLogic Administration Server username

- Oracle WebLogic Administration Server password

Because OPatchauto requires this input for each patching operation, OPatchauto provides the ability to encrypt the username and password via WebLogic encryption APIs and pass this information using a property file when running OPatchauto *apply* and *rollback* operations. The next section discusses how to create a property file.

23.2.1 Creating a Property File

The automated patching functionality achieved using opatchauto expects Oracle WebLogic Administration Server URL and credentials as an input for patching and configuration detection operations. Primarily, the Oracle WebLogic Administration server is the host that manages the Managed Server where the OMS instance is deployed. If you do not want to set the credentials every time you are prompted while patching the OMS, you can update the property file. OPatch allows you to repeatedly provide the inputs using property file option.

Note: The property file for a Primary OMS and Standby OMS are different, as they are in different domains.

To create a property file for OPatch:

Run the following script to create the Oracle WebLogic encrypted configuration and key files.

On UNIX:

```
$ OPatch/wlskeys/createkeys.sh -oh <full path of platform home> -location  
<location to put the encrypted files>
```

On Windows:

```
OPatch\wlskeys\createkeys.cmd -oh <full path of platform home> -location <location  
to put the encrypted files>
```

When prompted, enter the credentials of the Oracle WebLogic Administration Server that manages the Managed Server on which OMS instance is deployed. Two files are generated with the file names: config and key.

Create the property file with the following entries:

```
AdminServerURL=t3s://<host address from where admin server is running>:<port of  
the admin server>  
AdminConfigFile=<'config' file location>  
AdminKeyFile=<'key' file location>
```

23.3 Prerequisites for Running OPatchauto

Before running an OPatchauto patching session, you must ensure the following configuration and inventory-based prerequisites are satisfied: Configuration-based conditions that have to be honored for OMS automation is given below.

- The Enterprise Manager Software library must be configured.
- The Oracle WebLogic Administration Server that controls the OMS instance (currently to be patched) through a managed server must be up and running.
- Ensure that the Oracle Database, which houses the OMS Management Repository, and its listener are up and running.

- Ensure that you have the latest version of OPatch in the OMS platform home of each host. The latest OPatch is available from My Oracle Support, and release version 11.1.0.0.0. Patch 6880880

If you do not have the latest OPatch version, follow the instructions outlined in the My Oracle Support note 224346.1 available at:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=224346.1>

- Check your patch README to determine whether there are any specific prerequisites to be executed based on patch and patching methodologies.

Checking System Prerequisites

To make sure all prerequisite checks pass and no errors occur during the OPatchauto patching session, Oracle recommends running the following commands on each OMS instance (in your OMS system).

```
$<Platform Home>//OPatch/patchauto apply -analyze
```

Must be run from the System patch location (for *apply* operations)

Note: Check the Patch README and the instructions given for chosen patching methodologies.

OR

```
$<Platform Home>/OPatch/patchauto rollback -analyze -id <list of sub-patches to be rolled back for System patch>
```

Example 23–1 opatchauto apply -analyze Output

```
$opatchauto apply /omsSystemPatchV2/core_plugin_patches/1111123 -property_file /scratch/property_file -analyze
```

OPatch Automation Tool

Copyright (c) 2013, Oracle Corporation. All rights reserved.

```
OPatchauto version : 11.1.0.10.1
```

```
OUI version : 11.1.0.11.0
```

```
Running from : /multioms/middleware/oms
```

```
Log file location : /multioms/middleware/oms/cfgtoollogs/patch/patch2013-06-09_08-20-25AM_1.log
```

opatchauto log file:

```
/multioms/middleware/oms/cfgtoollogs/patchauto/1111123/patch_oms_2013-06-09_08-20-30AM_analyze.log
```

Configuration Validation: Success

```
Running apply prerequisite checks for patch(es) "111111" and Oracle Home "/multioms/middleware/plugins/oracle.sysman.emas.oms.plugin_12.1.0.4.0"...
```

Please monitor OPatch log file:

```
/multioms/middleware/plugins/oracle.sysman.emas.oms.plugin_
```

```
12.1.0.4.0/cfgtoollogs/patch/111111_Jun_09_2013_08_20_25/ApplyPrereq2013-06-09_08-20-53AM_8.log
```

Patches "111111" are successfully analyzed for Oracle Home

```
"/multioms/middleware/plugins/oracle.sysman.emas.oms.plugin_12.1.0.4.0"
```

```
Running apply prerequisite checks for patch(es) "111114" and Oracle Home
```

```
"/multioms/middleware/oms"...
```

```
Please monitor OPatch log file:
/multioms/middleware/oms/cfgtoollogs/opatch/111114_Jun_09_2013_08_20_
25/ApplyPrereq2013-06-09_08-20-55AM_8.log
Patches "111114" are successfully analyzed for Oracle Home
"/multioms/middleware/oms"

Copying all logs to: /multioms/middleware/oms/cfgtoollogs/opatch/2013-06-09_
08-20-25AM_SystemPatch_111123_1

Log file location: /multioms/middleware/oms/cfgtoollogs/patchauto/111123/opatch_
oms_2013-06-09_08-20-30AM_analyze.log

opatchauto succeeded.
```

Example 23-2 opatchauto rollback -analyze *output*

```
36084 OPatch]$ ./opatchauto rollback -id 111111,111114 -property_file
/scratch/property_file -analyze
OPatch Automation Tool
Copyright (c) 2013, Oracle Corporation. All rights reserved.

OPatchauto version : 11.1.0.10.1
OUI version        : 11.1.0.11.0
Running from       : /singleoms/middleware/oms
Log file location  :
/singleoms/middleware/oms/cfgtoollogs/patch/opatch/opatch2013-06-09_08-26-02AM_1.log

opatchauto log file:
/singleoms/middleware/oms/cfgtoollogs/patchauto/SystemPatch/opatch_oms_
2013-06-09_08-26-07AM_analyze.log

Configuration Validation: Success

Running rollback prerequisite checks for patch(es) "111111" and Oracle Home
"/singleoms/middleware/plugins/oracle.sysman.emas.oms.plugin_12.1.0.4.0"...
Please monitor OPatch log file:
/singleoms/middleware/plugins/oracle.sysman.emas.oms.plugin_
12.1.0.4.0/cfgtoollogs/patch/111111_Jun_09_2013_08_26_
02/RollbackPrereq2013-06-09_08-26-16AM_8.log
Patches "111111" are successfully analyzed for Oracle Home
"/singleoms/middleware/plugins/oracle.sysman.emas.oms.plugin_12.1.0.4.0"

Running rollback prerequisite checks for patch(es) "111114" and Oracle Home
"/singleoms/middleware/oms"...
Please monitor OPatch log file:
/singleoms/middleware/oms/cfgtoollogs/patch/111114_Jun_09_2013_08_26_
02/RollbackPrereq2013-06-09_08-26-19AM_8.log
Patches "111114" are successfully analyzed for Oracle Home
"/singleoms/middleware/oms"

Copying all logs to: /singleoms/middleware/oms/cfgtoollogs/patch/2013-06-09_
08-26-02AM_SystemPatch_111123_1

Log file location:
/singleoms/middleware/oms/cfgtoollogs/patchauto/SystemPatch/opatch_oms_
2013-06-09_08-26-07AM_analyze.log

opatchauto succeeded.
```

Note

- Once the analysis finishes, you can refer to the OPatchauto log to see what steps would be executed by OPatchauto in non `-analyze` mode. The log file contains references to the HTML and text output file HTML containing detailed steps.
- If you are analyzing a standby OMS system, you must include the `-standby` option.

23.4 Using OPatchauto

OPatchauto must be run from the platform home of the OMS being patched. The `ORACLE_HOME` environment variable must be set as the platform home or provided using the OPatchauto `"-oh"` option. For example:

```
<Platform Home>/OPatch/patchauto apply <patch>
```

Minimum Required OPatchauto Version: 11.1.0.10.1

Ensuring You Have the Latest Version of OPatch

OPatchauto uses the OPatch utility to apply the patch. For this reason, you must ensure that you have the latest version of OPatch 11.1 on all OMS instance platform homes. If you not sure which version of OPatch resides on your system, run the following command:

```
<Platform Home>/OPatch/patchauto version
```

To download the latest version of OPatch, follow the instructions outlined in the My Oracle Support note 224346.1 available at the following location:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=224346.1>

Patching Quickstart

Using OPatchauto typically involves the following phases:

1. Determining Whether Your System Meets OPatchauto System Requirements

```
Run patchauto apply -analyze
```

The `apply -analyze` command simulates an OPatchauto apply session by running all prerequisite checks, when possible, without making changes to the system (either bits or configurations). This command does not apply the patch.

See "[Prerequisites for Running OPatchauto](#)" on page 23-4 for additional information.

2. Determining What System Patches Currently Exist on Your System

```
Run patchauto lspatches
```

See "[lspatches](#)" on page 23-20 for more information.

3. Obtaining Patches from My Oracle Support (MOS)

OPatchauto requires that the required platform or plug-in System patches be obtained from My Oracle Support and downloaded to the OMS instance on which OPatchauto is to be run.

See "[My Oracle Support: Searching for Patches](#)" on page 23-8 for more information.

4. Applying a Patch

```
Run $patchauto apply <patch>
```

The `apply` command applies all patches within a specified System patch to the platform home from which OPatchauto is command is run.

See ["Running opatchauto apply"](#) on page 23-12 for more information.

5. Deinstalling Individual Sub-patches of a System Patch

Run `$opatchauto rollback -id <list of comma separated sub-patches of System patch>`

Note: For a complete list of sub-patches of the System patch, refer to the patch README.

If, after applying the patch, the system is not stable, the most likely cause is the patch itself. Contact Oracle Support. They will recommend that you remove the patch using the `opatchauto rollback` command.

See ["Running opatchauto rollback"](#) on page 23-13 for more information.

23.4.1 My Oracle Support: Searching for Patches

The first step in the patching process is to determine what patches you need from My Oracle Support (MOS). MOS is the single source of truth for patching. You can access MOS at the following location:

<https://support.oracle.com>

Once you have logged in, you have access to interactive support tools and information that simplify searching for and obtaining the requisite patches for your Oracle environment. The following figure shows the MOS home page.

Figure 23–2 My Oracle Support Main Page

The screenshot shows the My Oracle Support (MOS) main page. At the top, there's a navigation bar with tabs: Dashboard, Knowledge, Service Requests, Patches & Updates, Community, and Certifications. Below this, the page is organized into several sections:

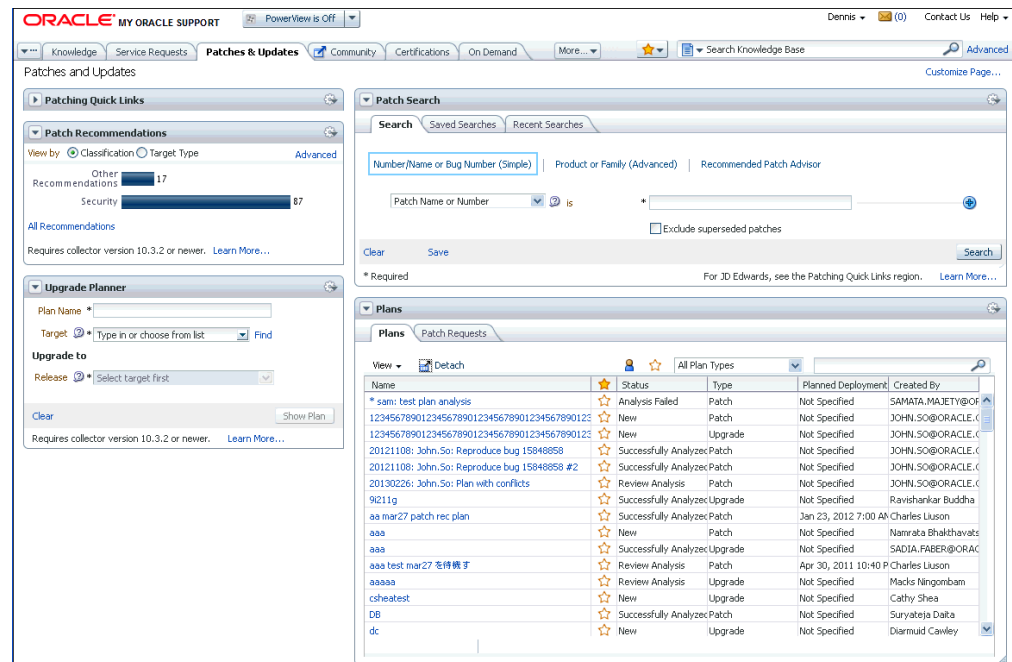
- News:** Includes links like "Get Proactive! Oracle Support Best Practices at COLLABORATE 13", "Do You Know the 3 Oracle Hardware Essentials?", "Welcome SelectMinds Customers - Start Here", and "My Oracle Support - Release Notes and Known Issues".
- Getting Started:** Features a "User Resource Center" with links to "Webcasts and Recordings", "Discover the Get Proactive Portfolio", and "Quick Video Training".
- Knowledge Base:** Contains a search bar and tabs for "Search & Browse", "Recently Viewed", "Recent Searches", and "Favorite Articles". Below the search bar, it says "Use the search tab first. Recent Searches will be listed here."
- Service Requests:** Displays a table of open requests. The table has columns for "Problem Summary", "SR Number", "Product", "Severity", and "Contact".
- Knowledge Articles:** Shows a list of recent updates, including "WS: 07: HCM Updates/Issues to Review After A9.3 Upgrade", "ASM Alert: After A Rebalance Operation, An ASM Disk Goes Offline In A RAC ASM Configuration", and "Performance Fix For Engineering Change Order (ECO) Flow From Agile to Oracle E-Business Suite".

Note: You can find complete documentation about MOS at the following location:

http://docs.oracle.com/cd/E25290_01/index.htm

My Oracle Support contains many features and capabilities that are grouped under tabs across the top of the application. Of primary interest is the *Patches and Updates* tab shown in the following figure.

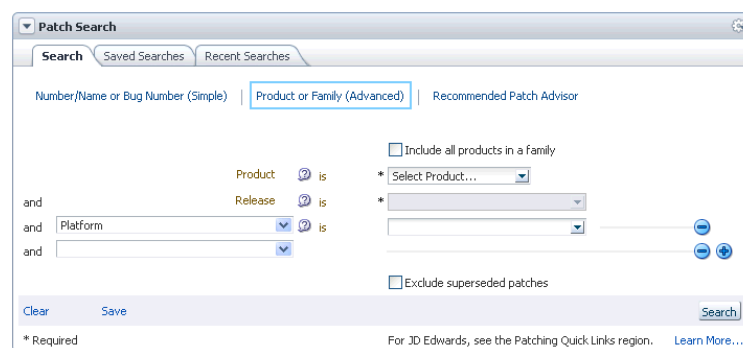
Figure 23–3 MOS Patches and Updates



Click on the tab to access the *Patches and Updates* page. From here, you can search for the patches based on the OMS patch area (core, plug-in, or combination). The following examples illustrate OMS patch searches for the various patch areas.

From the Search tab, click **Product or Family (Advanced)**.

Figure 23–4 Searching by Product or Family



Example: Searching for OMS Core Patches

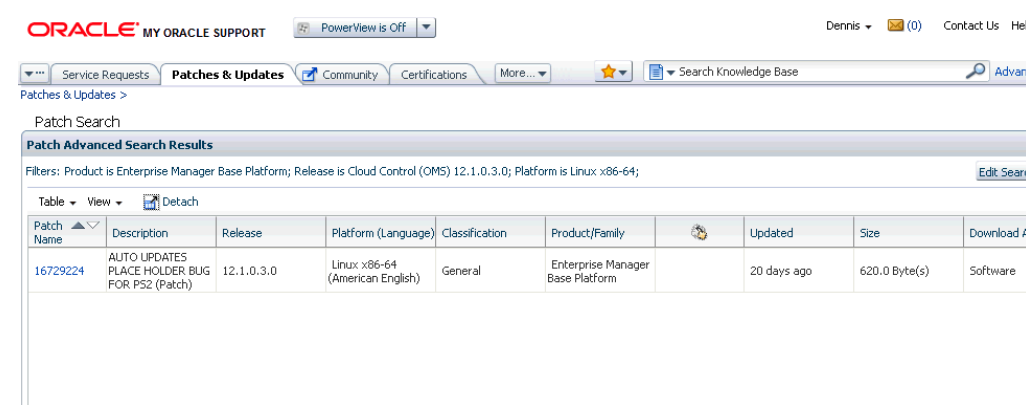
To search for Enterprise Manager OMS core patches, enter the following search parameters:

- **Product:** Enterprise Manager Base Platform
- **Release:** Cloud Control (OMS) 12.1.0.3.0

- **Platform:** Linux x86-64

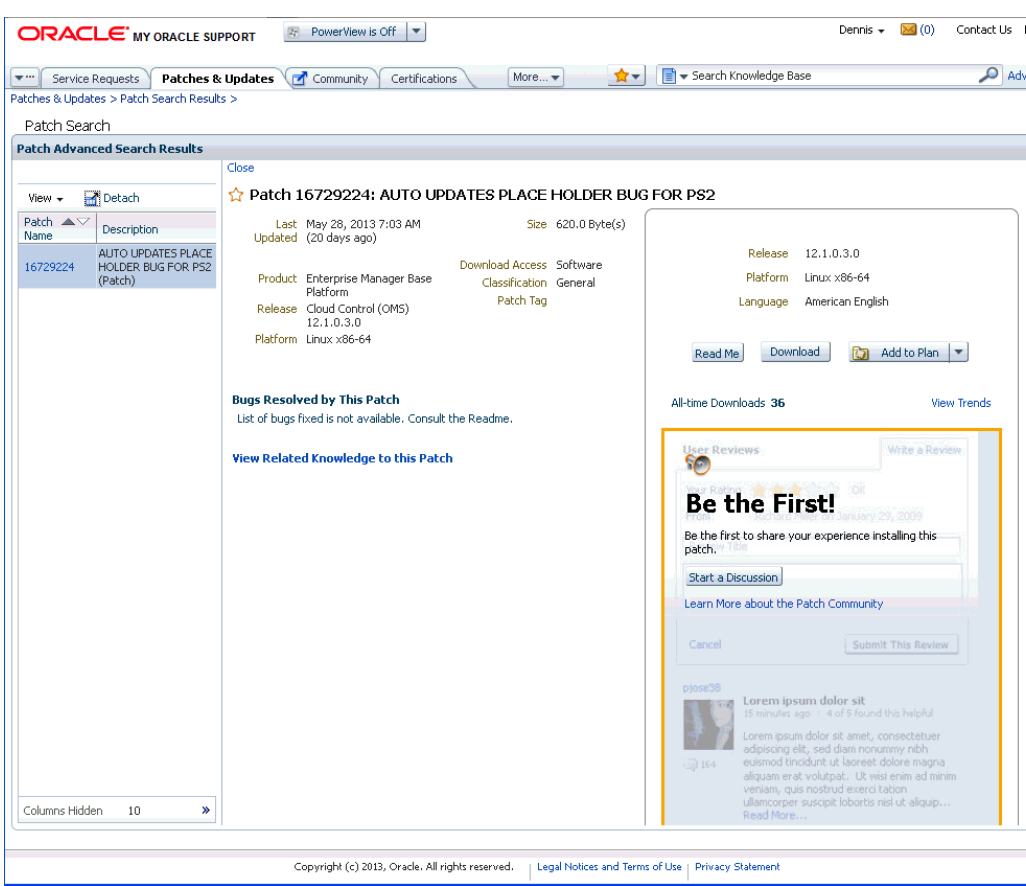
Clicking **Search** displays the following results.

Figure 23–5 OMS Core Patch Search Results

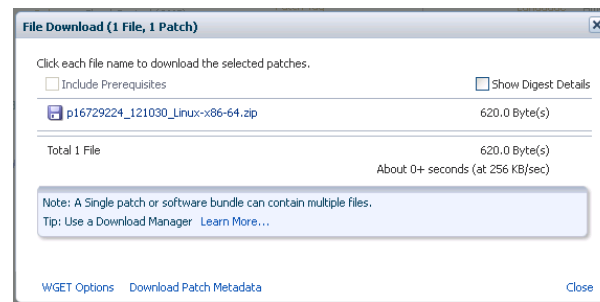


By clicking on patch 16729224 you are taken to the patch page where you can view bugs resolved by this patch, related Knowledge Articles, or view a generic patch README.

Figure 23–6 OMS Core Patch Page



Click **Download** to save the patch .ZIP file to your local system.

Figure 23–7 Patch Download Dialog**Example: Searching for OMS Plug-in Patches**

To search for Enterprise Manager OMS plug-in patches (Database plug-in is used for this example), enter the following search parameters:

- **Product:** Enterprise Manager for Oracle Database
- **Release:** DB Plug-in (OMS) 12.1.0.3.0
- **Platform:** Linux x86-64

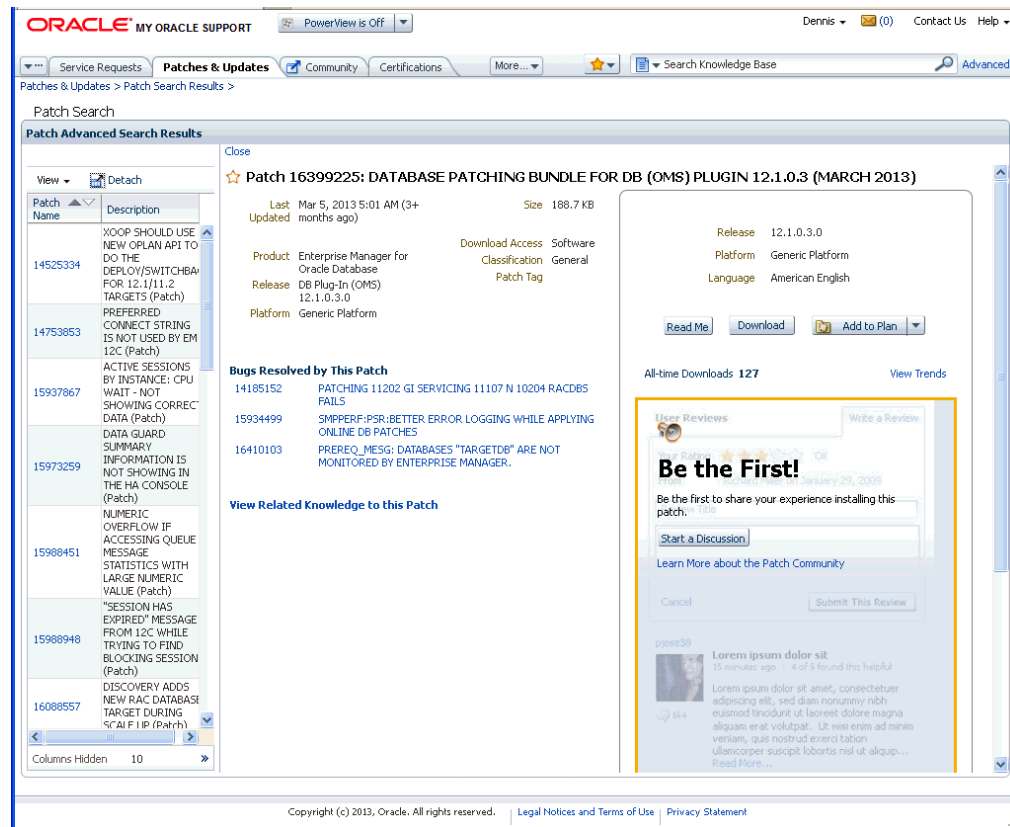
In the search results, click on patch 16399225.

Figure 23–8 OMS Plug-in Patch Search Results

Patch Name	Description	Release	Platform (Language)	Classification	Product/Family	Updated	Size	Download Access
16088557	DISCOVER ADD NEW RAC DATABASE TARGET DURING SCALE UP (Patch)	12.1.0.3.0	Generic Platform (American English)	General	Enterprise Manager for Oracle Database	3+ months ago	139.4 KB	Software
16213295	LOCKS IN EVENTBASEDPROCES JOB TYPE CAUSE JOB SYSTEM HANG (Patch)	12.1.0.3.0	Generic Platform (American English)	General	Enterprise Manager for Oracle Database	1+ month ago	15.5 KB	Software
16273848	WHEN RENAMING THE DATABASE SYSTEM TARGET THE CLOUD CONSOLE HANGS 12C (Patch)	12.1.0.3.0	Generic Platform (American English)	General	Enterprise Manager for Oracle Database	1+ month ago	28.0 KB	Software
16396459	TRACKING BUG FOR DB PLUGIN OMS SIDE FIXES FOR BUG 16038579 (Patch)	12.1.0.3.0	Generic Platform (American English)	General	Enterprise Manager for Oracle Database	3+ months ago	21.8 KB	Software
16399225	DATABASE PATCHING BUNDLE FOR DB (OMS) PLUGIN 12.1.0.3 (MARCH 2013) (Patch)	12.1.0.3.0	Generic Platform (American English)	General	Enterprise Manager for Oracle Database	3+ months ago	188.7 KB	Software
16423417	ASM_PATCHES METRIC GIVES ERROR FOR PRE-12.1ASM (NEEDS TO VERSION QUERY FOR 12.1) (Patch)	12.1.0.3.0	Generic Platform (American English)	General	Enterprise Manager for Oracle Database	2+ months ago	169.2 KB	Software
16438087	TRACKING BUG FOR DATA MASKING PATCH ON 12.1.0.3.0 DB PLUGIN (Patch)	12.1.0.3.0	Generic Platform (American English)	General	Enterprise Manager for Oracle Database	1+ month ago	382.2 KB	Software

On the patch page, in addition to downloading the patch and viewing the README, you can also view/access bugs resolved by this patch.

Figure 23–9 OMS Plug-in Patch Page



23.4.2 Running `opatchauto apply`

Once you have downloaded the patch, see the patch README for explicit patch details and instructions on applying the patch. You can find the README at the following location

```
<System patch location>/README.txt (or) README.html
```

As you step through the patching operations in the README, running `opatchauto apply` (depending on the configuration that is patched, primary or standby) will generate a custom, environment-specific version of the README for patching operations for the primary site multi-OMS or standby site OMS systems. For a primary site single OMS system, running `opatchauto apply` will perform patching and deployment operations.

On your local OMS instance, run the following command from the top level System patch directory:

```
$<Platform home>/OPatch/opatchauto apply <patch>
```

Note: Unlike `opatchauto analyze`, you should not run `opatchauto apply` on every OMS instance. OPatchauto will either execute all patching and deployment operations, or will generate environment-specific steps that include complete configuration aspects of the System.

For a multi-OMS primary site, OPatchauto will generate environment-specific README that includes patching steps in both text and HTML formats.

23.4.3 Running `opatchauto rollback`

See the patch README for explicit patch details and instructions on deinstalling the patch. You can find the README at the following location

<System patch location>/README.txt (or) README.html

As you step through the patch deinstallation operations in the README, running `opatchauto rollback` (depending on the configuration that is patched, primary or standby) will generate a custom, environment-specific version of the README for patching operations for the primary site multi-OMS or standby site OMS systems. For a primary site single OMS system, running `opatchauto rollback` will perform the deinstallation operations.

On your local OMS instance, run the following command from the top level System patch directory:

```
$<Platform home>/OPatch/opatchauto rollback -id <list of comma separated
sub-patches of System patch>
```

Note:

- Unlike `opatchauto analyze`, you should not execute the `opatchauto rollback` command on every OMS instance. OPatchauto will either execute all patching and deployment operations, or will generate environment-specific steps that include complete configuration aspects of the System.
- The list of sub-patches within the System patch can be retrieved from patch README.

The list of sub-patches listed in System patch README may differ from the patches that are actually installed. During System patch installation, some sub-patches may be skipped (not installed).

For a multi-OMS primary site, OPatchauto will generate environment-specific README that includes patching steps in both text and HTML formats.

23.4.4 Running `opatchauto lspatches`

After the patch is applied or rolled back, you can run the `opatchauto lspatches` command to generate a comprehensive Oracle home – patches map of the OMS instance homes and installed patches.

Example 23–3 Sample `opatchauto lspatches` Output

```
$ ./opatchauto lspatches
```

```
OPatch Automation Tool
Copyright (c) 2013, Oracle Corporation. All rights reserved.

Oracle Home:/singleoms/middleware/oms
1111114
13983293

Oracle Home:/singleoms/middleware/plugins/oracle.sysman.emas.oms.plugin_12.1.0.4.0
1111111

The following groups of patch(es) are applied as System Patch bundle(s):
1111114,1111111

For more details on individual homes, execute '$ORACLE_HOME/OPatch/patch
lsinventory -details -oh <desired home>'
```

23.4.5 Running opatchauto version

To determine the version numbers of the various OPatchauto utilities (Opatch, OPlan, OsysModel) that reside on your system, you can run `opatchauto version`.

Example 23–4 `opatchauto version` *Output*

```
$ ./opatchauto version
OPatch Version: 11.1.0.10.1
OPlan Version: 12.1.0.1.3
OsysModel build: Mon Jun 03 19:41:02 PDT 2013
opatchauto succeeded.
```

23.4.6 Patching a Standby OMS System

If you have configured a standby OMS for High Availability, refer to [Chapter 28, "Enterprise Manager Disaster Recovery"](#) and [Appendix D, "Standby OMSs using Standby WebLogic Domain"](#) for information on how to patch Standby OMS systems.

OPatchauto Command Syntax

This section provides a comprehensive listing and description of all OPatchauto commands used to patch an OMS.

Important: OPatchauto commands must be run from the OMS Platform home.

OPatchauto Commands

The OPatchauto commands are run from the OMS Platform home out of the standard OPatch directory. The Platform home must be set as \$ORACLE_HOME. In the following generic example, an OPatchauto command is run from a Platform home

```
<Platform home>/OPatch/opatchauto apply <PATH_TO_PATCH_DIRECTORY>
```

where <PATH_TO_PATCH_DIRECTORY> is the full path to the System patch top level directory.

OPatchauto consists of four primary commands.

- apply
- rollback
- version
- lspatches

OPatchauto help

You can view online help for any command (except *version*) by specifying the -help option.

For example:

- opatchauto -help
- opatchauto apply -help
- opatchauto rollback -help
- opatchauto lspatches -help

Apply

Apply a System patch to OMS instance homes. You must specify the patch location or the current directory will be used as the patch location.

Important: OPatchauto must be run from the platform home. ORACLE_HOME environment variable must be set as the platform home or provided using the -oh option.

You must run the *Apply* command directly from the System patch location.

Syntax

```
<Platform home>/OPatch/patchauto apply
    [-jre <Path to JRE>]
    [-nonrolling]
    [-invPtrLoc <Path to oraInst.loc>]
    [-property_file <Path to property file>]
    [-analyze]
    [-silent]
    [-oh <Platform home path>]
    [-standby]
```

Command Options

Table 23–1 *Apply*

Option	Description
jre	Instructs OPatchauto to use a Java environment other than the JRE in the platform home.
nonrolling	If patch is rolling patch as per the README, opatchauto can execute in nonrolling fashion, if this option is provided.
invPtrLoc	Overrides the default inventory pointer file to be used by OPatchauto with a administrator-specified oraInst.loc file.
property_file	<p>The administrator defined property file for OPatchauto to use. The path to the property file should be absolute.</p> <p>The keys for OPatchauto are:</p> <p>'AdminConfigFile' - Encrypted file for Administration Server administrator of the OMS instance domain.</p> <p>'AdminKeyFile' - Encrypted file for Administration Server password of the OMS instance domain.</p> <p>'AdminServerURL' - Administration Server URL of OMS instance domain.</p> <p>(Example: t3s://<host address>:<port number>)</p> <p>The Key, value pair is of the format 'x=y' where 'x' is an opatchauto/patch understood key and each pair is separated by newline in the property file.</p> <p>In order to create encrypted files for Oracle WebLogic Administration Server username & password, run the following script:</p> <pre>\$ORACLE_HOME/OPatch/wlskeys/createkeys.sh</pre> <p>(.cmd for Windows)</p> <p>to obtain the files and load them through a custom file using the property_file option.</p>

Table 23–1 (Cont.) Apply

Option	Description
analyze	This option performs prerequisite checks that includes both configuration and binary prerequisite checks. It simulates an apply operation (does not apply the patch). The command 'opatchauto apply' with this option must be run on each OMS instance to make sure all prerequisites pass for patching operations.
silent	This option refers to silent mode of invocation.
oh	The parameter passed through this option will override ORACLE_HOME environment variable. This parameter (if this option is used) must be the core platform home.
standby	This option should be used when patching a standby OMS infrastructure.

Rollback

Rollback sub-patches of a System patch from OMS instance home. Administrator specifies the sub-patch IDs of the System patch. You can obtain the sub-patch IDs by running the `opatchauto lspatches` command. See ["Running opatchauto lspatches"](#) on page 23-13.

Important: OPatchauto must be run from the platform home. ORACLE_HOME environment variable must be set as platform home or provided via the `-oh` option.

Syntax

```
opatchauto rollback -id <sub patches ID of a System patch>
                    [-id <list of comma separated sub-patches of System patch>]
                    [-invPtrLoc <Path to oraInst.loc>]
                    [-jre <LOC>] [-silent] [-nonrolling]
                    [-property_file <path to property file>]
                    [-analyze] [-oh <Platform home path>]
                    [-standby ]
```

Options

Table 23–2 Rollback

Option	Description
id	List of sub-patches of a System patch. For a complete list of sub-patches, see the System patch README.
jre	This jre option instructs OPatchauto to use the JRE (java) from the specified location instead of the default location under Oracle Home.
nonrolling	The nonrolling option instructs OPatchauto to run the patching session run in 'nonrolling' mode. Before the patching session can start, the following prerequisites must be met: <ul style="list-style-type: none">■ The stack on the local node must be running■ All remote nodes must be down.
invPtrLoc	The invPtrLoc option is used to locate the Central Inventory Pointer File (<code>oraInst.loc</code>) . Input for this option is the path to the <code>oraInst.loc</code> file.

Table 23–2 (Cont.) Rollback

Option	Description
property_file	<p>The administrator defined property file for opatchauto to use. The path to the property file should be absolute.</p> <p>The keys for 'opatchauto' are:</p> <p>'AdminConfigFile' - Encrypted file for Administration Server administrator of OMS instance domain.</p> <p>'AdminKeyFile' - Encrypted file for Administration Server password of OMS instance domain.</p> <p>'AdminServerURL' - Administration Server URL of OMS instance domain.</p> <p>(Example: t3s://<host address>:<port number>)</p> <p>The Key, value pair is of the format 'x=y' where 'x' is opatchauto/patch understood key and each pair is separated by newline in the property file.</p> <p>In order to create encrypted files for Oracle WebLogic Administration Server username & password, Use</p> <p>\$ORACLE_HOME/OPatch/wlskeys/createkeys.sh (.cmd for Windows) to get the files and load it thorough a custom file by 'property_file' option.</p>
analyze	<p>This option helps us to do dry-run prerequisite checks that includes both configuration and binary prerequisite checks. The command 'opatchauto rollback' with this option must be run on each OMS instance to make sure all prerequisites pass for patching operations.</p>
silent	<p>This option refers to silent mode of invocation.</p>
oh	<p>The parameter passed through this option will override ORACLE_HOME environment variable. This parameter (if this option is used) must be the core platform home.</p>
standby	<p>This option should be used when patching a standby OMS infrastructure.</p>

lspatches

Display patches and their Oracle home relationships with reference to OMS instances homes.

Syntax

```
opatchauto lspatches [ -invPtrLoc <Path to oraInst.loc> ]  
                    [-jre <LOC> ]  
                    [-oh <Platform Oracle home> ]
```

Options

Table 23–3 *lspatches*

Option	Description
jre	This jre option instructs OPatchauto to use the JRE (java) from the specified location instead of the default location under Oracle Home.
invPtrLoc	The invPtrLoc option is used to locate the Central Inventory Pointer File (oraInst.loc) . Input for this option is the path to the oraInst.loc file.
oh	The parameter passed through this option will override ORACLE_HOME environment variable. This parameter (if this option is used) must be the core platform home.

version

The `version` command shows the current version number of the OPatch utility, dependent OPlan version, and the osysmodel version.

Important: OPatchauto must be run from the platform home.

Syntax

```
<GI_HOME>/OPatch/patchauto version [-invPtrLoc <Path to oraInst.loc>]
                                     [-jre <LOC>]
                                     [-oh <ORACLE_HOME>]
                                     [-help] [-h]
```

Options

The following table describes the options available for the `version` command.

Table 23–4 `version` Command Options

Option	Description
-invPtrLoc	The <code>invPtrLoc</code> option is used to locate the Central Inventory Pointer File (<code>oraInst.loc</code>) . Input for this option is the path to the <code>oraInst.loc</code> file.
-jre	This <code>jre</code> option instructs OPatchauto to use the JRE (java) from the specified location instead of the default location under Oracle Home.
-oh	The <code>oh</code> option specifies the Oracle Home to work on. This takes precedence over the environment variable <code>ORACLE_HOME</code> .

Standby OMS Patching

There are two methods used to implement a standby OMS:

1. Standby OMS - Replicating OMS, Software library and Repository Components

If a standby site is set up through file system replication (Enterprise Manager 12.1.0.3 and later), there is no need to patch the standby site to keep it in sync with the primary site. When you apply patches to the primary OMS, the changes are automatically replicated at the standby site either manually or via automatic storage replication. For more details on this approach, see ["Enterprise Manager Disaster Recovery"](#) on page 28-1.

Using the data and storage replication, the standby site does not need to be patched.

2. Standby OMS Configured in Parallel with the Primary OMS System

When patches are applied on the primary site OMS, they must also be applied on the standby site Management Services. Note that patches typically update the Oracle Homes (via the `OPatch apply` command) and may require scripts to be run against the Management Repository. On the standby site, it is sufficient to update the Oracle Homes (using `OPatchauto`) and skip the running of scripts on the Management Repository because database changes are automatically propagated to the standby site using Data Guard. For more information, see [Appendix D, "Standby OMSs using Standby WebLogic Domain."](#)

To patch the standby site, run `opatchauto apply` and `opatchauto rollback` with the standby option.

Troubleshooting

This chapter describes common OPatchauto problems that may occur during patching operations or the analyze phase.

This chapter covers the following:

- [OPatchauto Troubleshooting Architecture](#)
- [OPatchauto Log Management Architecture](#)
- [Logs for Oracle Support](#)
- [OPatchauto: Cases Analysis, Error Codes, and Remedies/Suggestions](#)
- [Special Error Cases for OPatchauto OMS Automation](#)

OPatchauto Troubleshooting Architecture

In order for OPatchauto to fully automate the patching process, it accesses various tools/utilities to carry out different patching tasks in their respective phases. The primary tools/utilities outside of OPatch and OPatchauto are:

- `emctl stop oms` - Life cycle
- `emctl start oms` - Life cycle
- `emctl applypatch`, `emctl rollbackpatch` – Apply, rollback SQL changes in the OMS repository SYSMAN schema respectively
- `emctl register`, `emctl deregister` – Register, de-register metadata services with the right XMLs for MRS artifacts as per patch metadata instructions respectively

These tools/utilities are accessed during the patching process. Note that failure during invocation of these utilities can also happen and the errors & remedies for those commands are not handled in this document. They need to be followed up with Oracle Support for details. However, OPatchauto will trap errors from these commands output, push it to appropriate logs and announce it to the administrator and finally to support..

Apart from the above external tools/utilities, OPatchauto uses the following internal utilities to do binary patching operations. They have separated log files generated by OPatchauto. The internal utilities are patch binary prerequisite checks and patch binary apply, rollback operations.

OPatchauto Log Management Architecture

This section refers to the information through logs published by OPatchauto as part of its patching operations. This knowledge is needed for the administrator to obtain the appropriate logs from right area to troubleshoot and inform Oracle Support for further analysis. The following annotated example shows `opatchauto apply` output that displays the various log files that are created when running OPatchauto.

Sample OPatchauto Apply Output

```
$. /opatchauto apply /omsSystemPatchV2/core_plugin_patches/1111123 -property_file
/scratch/property_file -analyze
```

```
OPatch Automation Tool
```

```
Copyright (c) 2013, Oracle Corporation. All rights reserved.
```

OPatchauto version : 11.1.0.10.1

OUI version : 11.1.0.11.0

Running from : /multioms/middleware/oms

Log file location : /multioms/middleware/oms/cfgtoollogs/patch/patch2013-06-09_08-20-25AM_1.log

Consolidated OPatch log file – This log file contains all details of core OPatch session call references.

opatchauto log file:
/multioms/middleware/oms/cfgtoollogs/patchauto/1111123/patch_oms_2013-06-09_08-20-30AM_analyze.log

OPatchauto log file – This log file contains configuration information, traces of calls to individual commands, SDK calls and also mentions what step(s) passed, failed.

In case of failure of steps, it contains the failed step as command and remaining list of steps to be executed. This is useful to proceed so that administrator can resolve the failure, execute the remaining steps and then reach a good state in patching.

Configuration Validation: Success

Running apply prerequisite checks for patch(es) "111111" and Oracle Home
"/multioms/middleware/plugins/oracle.sysman.emas.oms.plugin_12.1.0.4.0"...

Please monitor OPatch log file:
/multioms/middleware/plugins/oracle.sysman.emas.oms.plugin_12.1.0.4.0/cfgtoollogs/patch/111111_Jun_09_2013_08_20_25/ApplyPrereq2013-06-09_08-20-53AM_8.log

Binary prerequisite log file for a sub-patch '111111'

Patches "111111" are successfully analyzed for Oracle Home
"/multioms/middleware/plugins/oracle.sysman.emas.oms.plugin_12.1.0.4.0"

Running apply prerequisite checks for patch(es) "111114" and Oracle Home
"/multioms/middleware/oms"...

Please monitor OPatch log file:
/multioms/middleware/oms/cfgtoollogs/patch/111114_Jun_09_2013_08_20_25/ApplyPrereq2013-06-09_08-20-55AM_8.log

Binary prerequisite log file for a sub-patch '111114'

Patches "111114" are successfully analyzed for Oracle Home
"/multioms/middleware/oms"

Copying all logs to: /multioms/middleware/oms/cfgtoollogs/patch/2013-06-09_08-20-25AM_SystemPatch_111123_1

All sub-logs, command logs go to this directory. This is useful to know what output each individual commands, execution threw on the screen. They are kept in separate logs (one for each command executed as part of automation). The naming of the individual logs are self-explanatory.

Log file location: /multioms/middleware/oms/cfgtoollogs/patchauto/111123/patch_oms_2013-06-09_08-20-30AM_analyze.log

OPatchauto log file will be printed at the end.

opatchauto succeeded.

Log output to a consolidated directory

As shown in the example above, there is a reference to pushing of all logs to consolidated log directory. The following line in the trace example shows this consolidation log directory.

...

```
Copying all logs to: /singleoms/middleware/oms/cfgtoollogs/patch/2013-06-09_08-41-33AM_SystemPatch_1111123_1
```

...

This consolidated log directory would contain the following files (here with reference to analyze example for apply).

```
$ ls -l /singleoms/middleware/oms/cfgtoollogs/patch/2013-06-09_08-41-33AM_SystemPatch_1111123_1

total 80

-rw-r--r-- 1 admin dba 7759 Jun 9 08:42 1111111_apply2013-06-09_08-42-23AM_11.log

-rw-r--r-- 1 admin dba 5343 Jun 9 08:42 1111111_ApplyPrereq2013-06-09_08-41-41AM_8.log

-rw-r--r-- 1 admin dba 8489 Jun 9 08:42 1111114_apply2013-06-09_08-42-20AM_10.log

-rw-r--r-- 1 admin dba 5440 Jun 9 08:42 1111114_ApplyPrereq2013-06-09_08-41-43AM_8.log

-rw-r--r-- 1 admin dba 120 Jun 9 08:41 AdminServerStatusPrerequisites_2013-06-09_08-41-40AM.log

-rw-r--r-- 1 admin dba 136 Jun 9 08:42 emctl_register_logmgmt_2013-06-09_08-42-28AM.log

-rw-r--r-- 1 admin dba 10740 Jun 9 08:42 opatch2013-06-09_08-41-33AM_1.log

-rw-r--r-- 1 admin dba 8252 Jun 9 08:42 rcu_applypatch_original_patch_2013-06-09_08-42-25AM.log

-rw-r--r-- 1 admin dba 66 Jun 9 08:41 RepositoryStatusPrerequisites_2013-06-09_08-41-40AM.log

-rw-r--r-- 1 admin dba 71 Jun 9 08:41 Swlib_Prerequisite_2013-06-09_08-41-40AM.log

-rw-r--r-- 1 admin dba 3221 Jun 9 08:42 temp_apply_automation.xml

-rw-r--r-- 1 admin dba 3242 Jun 9 08:42 temp_rollback_automation.xml

$
```

All the individual log files of each invocation, commands are finally copied to a consolidated place as highlighted above. Each command naming convention is self-explanatory and it indicates the actual operations being performed in automation. The 'opatchauto' log file will refer the individual log files so that administrator can easily connect to individual files to refer to any failure.

Logs for Oracle Support

If the administrator wants to contact Oracle Support, the administrator must provide the following references to Support.

- Administrator interface trace(s).
- Consolidated log directory as zip
- OPatch log file
- OPatchauto log file
- Output of 'opatchauto lspatches' command on all OMS instance homes.

OPatchauto: Cases Analysis, Error Codes, and Remedies/Suggestions

Refer to the following table for common OPatchauto error codes.

Table 23–5 *OPatchauto Error Codes*

Error Code	Description	Remedy/Suggestion
231	Wrong Oracle WebLogic Administration Server URL and/or invalid credentials	Correct the interview inputs and run OPatchauto again.
234	Malformed Oracle WebLogic Administration Server URL	If the Oracle WebLogic Administration Server URL is already defaulted (value given), type <enter>. If it is not given, construct the Oracle WebLogic Administration Server URL as t3s://<WebLogic Administration Server host address>:<WebLogic Administration Server port>.of the domain that controls the managed server on which the OMS is deployed.
235	Unable to connect to OMS repository	Check the OMS repository connectivity for SYSMAN administrator and run OPatchauto again.
236	OUI central inventory read issue	Check if the OUI inventory is locked by some other processes. Check if OUI inventory is readable.
238	Patch binary prerequisite checks failure	Check OPatchauto, OPatch, patch binary prerequisite log files for more details on the errors. If not resolved, contact Oracle Support.

Table 23–5 (Cont.) OPatchauto Error Codes

Error Code	Description	Remedy/Suggestion
240 - 251	Binary updates (or) deployment failure	<ul style="list-style-type: none"> This is a case for single OMS system. Patching steps are decided by OPatchauto but it failed to execute steps. OPatchauto will print the failed executed step and the remaining steps to be executed for completion of patching operations. Administrator need to contact Oracle support with logs, resolve why it failed and then must execute manually the failed step and steps referred by OPatchauto (in OPatchauto log file) to complete operations. In case of multi OMS (or) stand by OMS patching operations, failure of individual commands that got executed through text/html output must be brought to support notice for further diagnosis. After the failure condition is resolved, administrator need to execute the failed steps and further steps mentioned in HTML (or) text output to complete the patching operations.
233	<p>Software library not configured</p> <p>OMS repository connectivity not achieved. (post successful check of the same during credential inputs</p> <p>Oracle WebLogic Administration Server not reachable (post successful check of the same during credential inputs)</p>	Check the OPatchauto log file for the failure.

OPatchauto: External Utilities Error Codes

The following table lists exit codes for external utilities that OPatchauto uses for life cycle and deployment. If the deployment (or) life cycle fails through OPatchauto, the administrator can search individual log files for the error messages shown in the *Error Message/Recommendation* column.

Table 23–6 OPatchauto External Utilities Error Codes

Exit Code	Error Message / Recommendation
34	Displays the usage of the command.
35	Unable to read password! Exiting...
36	Unable to get a connection to the repository! Exiting...
37	The Plug-in is not deployed on this Management Server. The plug-in has to be deployed first to register metadata for that plug-in.
38	Input file does not exist
39	This operation is not supported by service.
40	Metadata operation is skipped.
41	Error occurred during Metadata registration.
42	Error occurred during Metadata de-registration.

Special Error Cases for OPatchauto OMS Automation

This section provides issue resolution information for special cases when using OPatchauto. This information will allow the administrator to handle these issues easily with less need for support team intervention.

Windows analyze failure due to lock of files by Oracle WebLogic Administration Server

In Windows operating systems, it has been noticed that some of the Enterprise Manager related files (used for patching) are locked by running of Oracle WebLogic Administration Server. As OPatchauto required Oracle WebLogic Administration Server to be RUNNING for the configuration detection, we need to perform the following steps to make sure that this conflict with respect to environment and patching is removed.

1. As per the patch README & guide, Administrator would try to run `opatchauto analyze` to check if configuration and binary prerequisites succeeds for each of the OMS instances homes. If in Windows, the files are locked due to Oracle WebLogic administration server, we should get binary prerequisite checks failure.

```
Apply binary prerequisite checks failed for patch(es) "111112".
```

```
Copying all logs to: C:\MW_130~1\oms\cfgtoollogs\opatch\2013-06-19_02-55-10AM_SystemPatch_111112_1
```

```
OPatch failed: Some of the prerequisite checks to apply or rollback the patch(es) have failed.
```

Error occurred during the Run patch(es) binary prerequisite checks phase.

```
Detail: Some of the prerequisite checks to apply or rollback the patch(es) have failed.
```

```
Log file location: C:\MW_130~1\oms\cfgtoollogs\opatchauto\111112\opatch_oms_201
```

```
3-06-19_02-55-14AM_analyze.log
```

```
Recommended actions: Please check opatchauto, opatch, patch(es) binary prerequisite log files for more details on the errors. If not resolved, please contact Oracle Support.
```


opatchauto failed with error code = 238

2. Once binary patch prerequisite failure occurs, Administrator is recommended to open the individual patch binary prerequisite checks log and check the failure. If there are active files locked, the individual prerequisite *checkActiveFilesAndExecutables* would have failed in the consolidated OPatch prerequisites report (inside the log) and it would also show the files that are currently active.

```
, CheckActiveServicesForWindows=[ Prerequisite Status: NOT ATTEMPTED (serious failure), Prerequisite output:
```

The details are:

```
Executable "C:\MW_130~1\oms\bin\oradim.exe" was not found in the Oracle Home, so not checking for active services.]
```

```
, CheckActiveFilesAndExecutables=[ Prerequisite Status: FAILED, Prerequisite output:
```

The details are:

Following files are active :

```
C:\MW_130~1\oms\sysman\jlib\emCoreConsole.jar]
```

3. To ascertain if these active files are due to Oracle WebLogic Administration Server, we recommend the following steps:

1. Go to ORACLE_HOME\bin and run 'emctl stop oms -all' to stop the complete stack for OMS. Note that ORACLE_HOME is already set as Platform home.
2. Use some open source utilities like process explorer and search for file strings given as output in (2) to check if any files are still active. If so, kill the process tree of those files.
3. For each OMS instance home:

1. Navigate to ORACLE_HOME\OPatch

2. Execute the following command:

```
opatchauto checkApplicable -ph <System patch location> -oh <Platform home>
(for apply operations)
```

OR

```
run: opatchauto checkApplicable -id <list of sub-patch IDs of System patch>
-oh <Platform home> (for rollback operations)
```

Note: For a complete list of sub-patches of the System patch, refer to the Patch README.

If this command passes for each OMS instance home, it is confirmed that active files are only due to the Oracle WebLogic Administration Server being in the running home. Go ahead and apply the System patch. However, if active files still persist, Contact Oracle Support.

4. Go to ORACLE_HOME\bin and run `emctl start oms -admin_only`.

5. Run OPatchauto in non-analyze mode. For further instructions, refer to the patch README and Administrator guide.

Once the OPatchAuto is run in non-analyze mode, it will check if active files are again locked by Oracle WebLogic administration server and will provide a prompt as shown below (in silent mode it will be auto-yes):

```
Running prerequisite checks to verify if any files or services are locked by
admin server process...
Please monitor OPatch log file: c:\MW_130518\oms\cfgtoollogs\opatch\1111112_
Jun_
26_2013_08_16_19\ApplyPrereq2013-06-26_08-16-57AM_8.log
```

The details are:

```
Following files are active :
c:\MW_130518\oms\sysman\jlib\emCoreConsole.jar
```

```
Due to active files to be patched, opatchauto will stop all OMS processes so
tha
t lock on active files may be released...
Do you want to proceed? [y|n]
y
User Responded with: Y
opatchauto has stopped all OMS processes successfully.
```

If there is a failure while stopping OMS processes, OPatchauto will accordingly error out. Refer to opatchauto log file for details.

6. OPatchauto will stop the stack and then ask for a confirmation from the administrator on whether to proceed with prerequisite checks of patch binaries (in silent mode it will be auto-yes):

```
opatchauto has stopped all OMS processes successfully. Please make sure the
above listed active files are unlocked by all windows processes.
Do you want to proceed? [y|n] y
```

User Responded with: Y

Note: Administrators are again requested to check the process explorer tree and make sure that active files mentioned earlier in prerequisite log are not listed there.

Once confirmed, OPatch will run the checks, patch, and deploy the automation elements.

7. Note that OPatchauto will not attempt to re-start the stack. The administrator has to restart the stack as needed.

A complete sample trace of this case is shown below:

```
C:\MW_130518\oms\OPatch_June26>opatchauto apply ..\patches\cmdRcu\1111112
OPatch Automation Tool
Copyright (c) 2013, Oracle Corporation. All rights reserved.
```

```
OPatchauto version : 11.1.0.10.1
OUI version        : 11.1.0.11.0
Running from       : c:\MW_130518\oms
Log file location  : c:\MW_130518\oms\cfgtoollogs\opatch\opatch2013-06-26_08-16-
19AM_1.log
```

```
opatchauto log file: c:\MW_130518\oms\cfgtoollogs\opatchauto\1111112\opatch_oms_
2013-06-26_08-16-23AM_deploy.log
```

```
Please enter the WebLogic Admin Server URL for primary OMS:> t3s://example.o
racle.com:7101
```

```
Please enter the WebLogic Admin Server username for primary OMS:> weblogic
```

```
Please enter the WebLogic Admin Server password for primary OMS:>
```

```
Configuration Validation: Success
```

```
Running prerequisite checks to verify if any files or services are locked by adm
in server process...
```

```
Please monitor OPatch log file: c:\MW_130518\oms\cfgtoollogs\opatch\1111112_Jun_
26_2013_08_16_19\ApplyPrereq2013-06-26_08-16-57AM_8.log
```

```
The details are:
```

```
Following files are active :
```

```
c:\MW_130518\oms\sysman\jlib\emCoreConsole.jar
```

```
Due to active files to be patched, opatchauto will stop all OMS processes so that
lock on active files may be released...
```

```
Do you want to proceed? [y|n]
```

```
Y
```

```
User Responded with: Y
```

```
opatchauto has stopped all OMS processes successfully.
```

```
opatchauto has stopped all OMS processes successfully. Please make sure the above
listed active files are unlocked by all windows processes.
```

```
Do you want to proceed? [y|n]
```

```
Y
```

```
User Responded with: Y
```

```
Running apply prerequisite checks for patch(es) "1111112" and Oracle Home "c:\MW
_130518\oms"...
```

```
Please monitor OPatch log file: c:\MW_130518\oms\cfgtoollogs\opatch\1111112_Jun_
26_2013_09_01_33\ApplyPrereq2013-06-26_09-03-41AM_10.log
```

```
Patches "1111112" are successfully analyzed for Oracle Home "c:\MW_130518\oms"
```

```
To continue, OPatch will do the following:
```

```
[Patch and deploy patch(es) binaries] : Apply patch(es) [ 1111112 ] to Oracle
Home "c:\MW_130518\oms";
```

```
Apply RCU artifact with patch "c:\MW_1
30518\oms\patch_storage\1111112_Feb_21_2013_06_30_38\original_patch"
```

```
Do you want to proceed? [y|n]
```

```
Y
```

```
User Responded with: Y
```

```
Applying patch "1111112" to Oracle Home "c:\MW_130518\oms"...
```

```
Please monitor OPatch log file: c:\MW_130518\oms\cfgtoollogs\opatch\1111112_Jun_
26_2013_09_01_33\apply2013-06-26_09-04-17AM_12.log
```

```
Updating repository with RCU reference file "c:\MW_130518\oms\patch_storage\111
112_Feb_21_2013_06_30_38\original_patch"
```

Copying all logs to: c:\MW_130518\oms\cfgtoollogs\opatch\2013-06-26_09-01-32AM_SystemPatch_1111112_1

Patching summary:

Following patch(es) are successfully applied (Oracle home:patch list):

c:\MW_130518\oms:1111112

Log file location: c:\MW_130518\oms\cfgtoollogs\opatchauto\1111112\opatch_oms_2013-06-26_09-01-36AM_deploy.log

opatchauto succeeded.

Starting and Stopping Enterprise Manager Components

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](#)
- [Controlling JVM and ADP Engines](#)
- [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](#)

24.1 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](#)

24.1.1 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 24-1](#).

For example, to stop the Management Agent, enter the following commands:

```
$PROMPT> cd AGENT_INSTANCE_HOME/bin
$PROMPT> emctl stop agent
```

Table 24–1 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 24–1).

Example 24–1 Checking the Status of the Management Agent

```
$ ./emctl status agent
Oracle Enterprise Manager 12c Release 3 Cloud Control 12.1.0.3.0
Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved.
-----
Agent Version      : 12.1.0.2.0
OMS Version        : 12.1.0.2.0
Protocol Version   : 12.1.0.1.0
Agent Home         : /scratch/aime/work/agtbase7453/agent_inst
Agent Binaries     : /scratch/aime/work/agtbase7453/core/12.1.0.3.0
Agent Process ID   : 7834
Parent Process ID  : 7781
Agent URL          : https://slc01lwi.us.oracle.com:3872/emd/main/
Repository URL     : https://slc01lwi.us.oracle.com:4900/empbs/upload
Started at        : 2013-06-05 01:33:33
Started by user    : aime
LLast Reload       : (none)
Last successful upload          : 2013-06-20 22:17:50
Last attempted upload          : 2013-06-20 22:17:50
Total Megabytes of XML files uploaded so far : 20.08
Number of XML files pending upload          : 0
Size of XML files pending upload(MB)        : 0
Available disk space on upload filesystem   : 62.89%
Collection Status                     : Collections enabled
Heartbeat Status                       : Ok
Last attempted heartbeat to OMS         : 2013-06-20 22:17:49
Last successful heartbeat to OMS        : 2013-06-20 22:17:49
Next scheduled heartbeat to OMS        : 2013-06-20 22:18:49
-----
Agent is Running and Ready
```

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.

24.1.2 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 24.2.1](#) 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 24–2](#) describes the Windows service that you use to control the Management Agent.

Table 24–2 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.

24.1.3 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 24-1](#)).

24.1.4 Troubleshooting Management Agent Startup Errors

If the agent fails to start, see the `emctl . log` and `emagent . nohup log` files for details. The log files are saved in the `$AGENT_INSTANCE_HOME/sysman/logs` direstory. Following are common issues and troubleshooting suggestions:

24.1.4.1 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.

24.1.4.2 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 24.6.5](#).

24.1.4.3 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.

24.1.4.4 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.

24.2 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](#)

24.2.1 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 24-3](#).

For example, to stop the Management Service, enter the following commands:

```
$PROMPT> cd bin
$PROMPT> ./emctl stop oms
```

Table 24-3 Starting, Stopping, and Checking the Status of the Management Service

Command	Purpose
<code>emctl start oms</code>	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 24–3 (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 <code>emctl stop oms</code> command does not stop Fusion Middleware. Use <code>emctl stop oms -all</code> 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 <code>emctl status oms -details</code> to view information about the configuration of the management service such as ports being used and the URLs for console and upload.
emctl config oms -set_startup_mode [pbs_only console_only normal]	Starts the OMS in Console Only Mode by shutting down background services.

24.2.2 Controlling the Management Service on Windows

When you install the Oracle Management Service on a Windows system, the installation procedure creates a 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**, then **Administrative Tools** from the **Start** menu.

Note: The `emctl` utility described in [Section 24.2.1](#) 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 24–4](#) describes the Windows service that you use to control the Oracle Management Service.

Table 24–4 Service Installed and Configured When Installing the Oracle Management Service on Windows

Component	Service Name Format	Description
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.

24.2.3 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 24–5](#). The `INSTANCE_HOME` mentioned in the table is the OMS instance home and `n` is the index of the OMS server.

Table 24–5 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
Node Manager log files	\$INSTANCE_HOME/NodeManager/emnodemanager/nodemanager.log

WebTier Service Fails to Start

Check logs under <WebTier Instance Home>/diagnostics folder in case WebTier start fails.

24.3 Controlling JVMD and ADP Engines

This section contains the following:

- [Controlling JVMD Engines](#)
- [Controlling ADP Engines](#)

24.3.1 Controlling JVMD Engines

[Table 24–6](#) lists the commands used to control JVMD Engines.

Table 24–6 Commands Used to Control JVMD Engines

Command	Purpose
emctl extended oms jvmd list	Lists all the JVMD Engines.
emctl extended oms jvmd start -server=<server_name1>,<server_name2>...	Starts specified JVMD Engines. Accepts the engine names as a comma separated list.
emctl extended oms jvmd start -all	Starts all JVMD Engines.
emctl extended oms jvmd stop -server=<server_name1>,<server_name2>...	Stops specified JVMD Engines. Accepts the engine names as a comma separated list.
emctl extended oms jvmd stop -all	Stops all JVMD Engines.
emctl extended oms jvmd status -server=<server_name1>,<server_name2>...	Displays the status of specified JVMD Engines. Accepts the engine names as a comma separated list.
emctl extended oms jvmd status -all	Displays the status of all JVMD Engines.

Table 24–6 (Cont.) Commands Used to Control JVM D Engines

Command	Purpose
emctl extended oms jvmd -help	Displays the list of available commands for the JVM D verb.

24.3.2 Controlling ADP Engines

Table 24–7 lists the commands used to control ADP Engines.

Table 24–7 Commands Used to Control ADP Engines

Command	Purpose
emctl extended oms adp list	Lists all ADP Engines.
emctl extended oms adp start -server=<server_ name1>,<server_name2>...	Starts specified ADP Engines. Accepts the engine names as a comma separated list.
emctl extended oms adp start -all	Starts all ADP Engines.
emctl extended oms adp stop -server=<server_ name1>,<server_name2>...	Stops specified ADP Engines. Accepts the engine names as a comma separated list.
emctl extended oms adp stop -all	Stops all ADP Engines.
emctl extended oms adp status -server=<server_ name1>,<server_name2>...	Displays the status of specified ADP Engines. Accepts the engine names as a comma separated list.
emctl extended oms adp status -all	Displays the status of all ADP Engines.
emctl extended oms adp -help	Displays the list of available commands for the ADP verb.

24.4 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.

24.5 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.

24.5.1 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 24-5

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 24-1

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.

24.5.2 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 24-5

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 24-1

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, follow these steps:
 - 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
```

24.6 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](#)

24.6.1 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 24–8](#).

Table 24–8 *Manually Reloading and Uploading Management Data*

Command	Description
<code>emctl upload (agent)</code>	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.
<code>emctl reload (agent)</code>	<p>This command can be used to apply the changes after you have manually modified the <code>emd.properties</code> file. For example, to change the upload interval, <code>emd.properties</code> can be modified, and <code>emctl reload</code> can then be run.</p> <p>Note: Oracle does not support manual editing of the <code>targets.xml</code> files unless the procedure is explicitly documented or you are instructed to do so by Oracle Support.</p>

24.6.2 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 11g, 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 11g 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 12c Cloud Control Console.

24.6.3 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 list new monitoring credentials:

```
$PROMPT>./emctl config agent listtargets
```

[Example 24-2](#) shows the typical output of the command.

Example 24-2 Listing the Targets on a Managed Host

```
ade:[ example_name_1208_qc_ag ] [example_username@example emagent]$ emctl config
agent listtargets
Oracle Enterprise Manager 12c Release 3 Cloud Control 12.1.0.3.0
Copyright (c) 1996, 2013 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]$
```


24.6.4 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: "For more information on controlling blackouts, refer to [Chapter 9](#).

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 24–9](#).

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 24.6.3](#).

Table 24–9 Summary of Blackout Commands

Blackout Action	Command
Set an immediate blackout on a particular target or list of targets	<pre>emctl start blackout <Blackoutname> [<Target_name>[:<Target_Type>]].... [-d <Duration>]</pre> <p>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.</p> <p>The <code>-d</code> option is used to specify the duration of the blackout. Duration is specified in [days] hh:mm where:</p> <ul style="list-style-type: none"> ■ days indicates number of days, which is optional ■ hh indicates number of hours ■ mm indicates number of minutes <p>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 <code>-nodelevel</code> argument.</p> <p>If two targets of different target types share the same name, you must identify the target with its target type.</p>
Stop an immediate blackout	<pre>emctl stop blackout <Blackoutname></pre>
Set an immediate blackout for all targets on a host	<pre>emctl start blackout <Blackoutname> [-nodeLevel] [-d <Duration>]</pre> <p>The <code>-nodeLevel</code> 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 <code>-nodeLevel</code> option must follow the blackout name. If you specify any targets after the <code>-nodeLevel</code> option, the list is ignored.</p>
Check the status of a blackout	<pre>emctl status blackout [<Target_name>[:<Target_Type>]]....</pre>

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
```

24.6.5 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.

24.6.6 Reevaluating Metric Collections

Use the following command to perform an immediate reevaluation of a metric collection:

```
emctl control agent runCollection <targetName>:<targetType> <collectionItemName>
```

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/plugins/<plugin id> 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 `$INSTALL_BASE/plugins/<plugin id directory` where `$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.

24.7 emctl Commands

This section lists the `emctl` commands for the Enterprise Manager Management Agent and Oracle Management Service.

[Table 24–10](#) explains the `emctl` commands for OMS.

Table 24–10 *emctl Commands for OMS*

emctl Command	Description
<code>emctl [getversion] oms</code>	Gets the version of the Management Service. Sample output is as follows: <pre>./emctl getversion oms Oracle Enterprise Manager Cloud Control 12c Release 2 12.1.0.3.0 Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved. Enterprise Manager 12c OMS Version 12.1.0.1.0</pre>
<code>emctl [start] oms</code>	Starts the Management Service.

Table 24–10 (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	<p>-force can be used with both emctl stop oms -all and emctl stop oms.</p> <p>If the emctl stop oms commands do not shutdown the relevant processes, using -force option will forcefully stop the relevant processes.</p>
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 set property	<p>Sets the value for OMS configuration properties.</p> <p>emctl set property will set the property value by default for all the OMSes. To set the property value for a single OMS, aspecify an extra option -oms_name, which should be in format hostname.us.oracle.com:17707_Management_Service. To set the property value for the current OMS, specify -oms_name = "local_oms". You can also use the -oms_name option for remote OMS, by setting the value from any OMS, for another OMS that is running within the same domain.</p> <p>Note: From Enterprise Manager 12.1.0.2.0 onward, you can also view and edit OMS properties from the Cloud Control console as follows:</p> <ol style="list-style-type: none"> 1. From the Setup menu, select Manage Cloud Control, then select Management Services. 2. On the Management Services page, click Configuration Properties. 3. On the Configuration Properties page, you can view and edit OMS properties. <p>Note: You will need OMS Configuration Property resource privilege to navigate to this page.</p>
emctl get property	Gets the value for OMS configuration properties.
emctl list properties	<p>View properties for all OMSes.</p> <p>Use -out_file option in emctl list properties command to get the list of all the properties for all OMSes as follows:</p> <pre>emctl list properties -out_file</pre> <p>The above command enables easy comparison of configuration across two OMSes.</p>
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.

Table 24–10 (Cont.) emctl Commands for OMS

emctl Command	Description
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 24–11 explains emctl commands for Management Agents.

Table 24–11 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.
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: <pre> bash-3.00\$ emctl status agent -secure Oracle Enterprise Manager 12c Release 2 Cloud Control 12.1.0.3.0. Copyright (c) 1996, 2013 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\$ </pre>
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: <pre> bash-3.00\$ emctl status agent jobs Oracle Enterprise Manager 12c Release 2 Cloud Control 12.1.0.3.0. Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved. ----- step id typ pid stat command line ----- --- -- --- ----- Agent is Running and Ready </pre>

Table 24–11 (Cont.) emctl Commands for Management Agent

emctl Command	Description
emctl status agent target <target name>,<target type>,<metric>	<p>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.</p> <pre> bash-3.00\$ emctl status agent target database,oracle_ database Oracle Enterprise Manager 12c Release 3 Cloud Control 12.1.0.3.0. Copyright (c) 1996, 2013 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 </pre>

Table 24–11 (Cont.) emctl Commands for Management Agent

emctl Command	Description
emctl status agent mcache <target name>,<target type>,<metric>	<p>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:</p> <pre>bash-3.00\$ emctl status agent mcache example.us.oracle.com,host Oracle Enterprise Manager 12c Release 3 Cloud Control 12.1.0.3.0. Copyright (c) 1996, 2013 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</pre> <p>The metrics listed above are the ones whose values are present in the metric cache.</p>
emctl reload agent dynamicproperties [<Target_name>:<Target_ Type>]...	<p>Recomputes the dynamic properties of a target and generates the dynamic properties for the target.</p> <p>Sample output for oracle_database is as follows:</p> <pre>bash-3.00\$ emctl reload agent dynamicproperties database:oracle_database Oracle Enterprise Manager 12c Release 3 Cloud Control 12.1.0.3.0. Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved. ----- EMD recompute dynprops completed successfully</pre>
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>]	<p>Allows you to configure the console based on the configuration entries that you have mentioned in the file <fileloc>.</p> <p><EM loc> is optional and can be used to operate on a different Oracle Home.</p>
emctl config [agent] listtargets [<EM loc>]	<p>Lists all targets present in targets.xml.</p> <p><EM loc> is optional and can be used to operate on a different Oracle Home.</p>

Table 24–11 (Cont.) emctl Commands for Management Agent

emctl Command	Description
emctl control agent runCollection <target_ name>:<target_type> <metric_name>	<p>Allows to manually run the collections for a particular metric of a target. Sample output is as follows:</p> <pre>emctl control agent runCollection example.us.oracle.com:host CPUUsage Oracle Enterprise Manager 12c Release 3 Cloud Control 12.1.0.3.0. Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved. ----- EMD runCollection completed successfully</pre>
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	<p>Prints the version of the agent. Sample output is as follows:</p> <pre>./emctl getversion agent Oracle Enterprise Manager 12c Release 3 Cloud Control 12.1.0.3.0. Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved. Oracle Enterprise Manager 12c Release 1 Cloud Control Agent 12.1.0.1.0</pre>
emctl dumpstate agent <component> ...	<p>Generates the dumps for the agent. This command allow you to analyze the memory/cpu issues of the agent. Sample output is as follows:</p> <pre>./emctl dumpstate agent Oracle Enterprise Manager 12c Release 3 Cloud Control 12.1.0.3.0. Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved. Dumpstate succeeded</pre>
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	<p>Prints the agent home directory. The sample output is as follows:</p> <pre>bash-3.00\$ emctl getemhome Oracle Enterprise Manager 12c Release 3 Cloud Control 12.1.0.3.0. Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved. EMHOME=/scratch/aime/gcagent/ngagent/agent_inst</pre>
emctl start blackout <Blackoutname> [-nodeLevel] [<Target_ name>[:<Target_Type>]]... [-d <Duration>]	<p>Starts blackout on a target.</p> <p><Target_name:Target_type> defaults to local node target if not specified.</p> <p>If -nodeLevel is specified after <Blackoutname>,the blackout will be applied to all targets and any target list that follows will be ignored.</p> <p>Duration is specified in [days] hh:mm</p>

Table 24–11 (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>	<p>'agent' creates and deploys only the agent.</p> <p>[-s <password>]: Install password for securing agent.</p> <p>[-S]: Password will be provided in STDIN.</p> <p>[-o <omshostname:consoleSrvPort>]: The OMS Hostname and console servlet port. Choose the unsecured port.</p> <p><deploy-dir> : Directory to create the shared (state-only) installation port.</p> <p><deploy-hostname:port> : Host name and port of the shared (state-only) installation. Choose unused port.</p> <p><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>.</p> <p><sid>: The instance of the remote database. Only specified when deploying "dbconsole".</p>
emctl setproperty agent	<pre>./emctl setproperty agent -name emStageDir -value /tmp</pre> <p>Oracle Enterprise Manager Cloud Control 12c Release 3 Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved.</p> <p>EMD setproperty succeeded</p>
emctl getproperty agent	<pre>./emctl getproperty agent -name emStageDir</pre> <p>Oracle Enterprise Manager Cloud Control 12c Release 3 Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved.</p> <p>emStageDir=/tmp</p>
emctl clear_property agent	<pre>./emctl clear_property agent -name emStageDir</pre> <p>Oracle Enterprise Manager Cloud Control 12c Release 3 Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved.</p> <p>EMD clear_property succeeded</p>

Table 24–11 (Cont.) emctl Commands for Management Agent

emctl Command	Description
emctl status agent verify	<pre>./emctl status agent verify</pre> <p>Oracle Enterprise Manager Cloud Control 12c Release 3 Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved.</p> <p>-----</p> <p>verifykey: Successfully Completed communication with agent</p>

24.8 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/em/EMGC_OMS<n>/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:

For Unix:

```
<agent_instance_home>/bin/emctl status agent
```

For Windows:

```
<agent_instance_home>\bin\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:

For Unix:

```
<OMS_HOME>/bin/emctl status oms
```

For Windows:

```
<OMS_HOME>\bin\emctl status oms
```

Example 24–3 Sample Log Content for OMS

```
2013-06-23 22:50:25,686 [main] INFO wls.OMSController main.219 - Executing emctl
command : status
```

```
2013-06-23 22:50:26,281 [main] INFO  commands.BaseCommand printMessage.404 -
statusOMS finished with result: 0
2013-06-23 22:50:35,885 [main] INFO  wls.OMSController main.219 - Executing emctl
command : status
2013-06-23 22:50:36,464 [main] INFO  commands.BaseCommand printMessage.404 -
statusOMS finished with result: 0
```

In another example, run the following command from the Oracle home directory of the Management Agent to upload data:

For Unix:

```
<Agent_Instance_Home>/bin/emctl upload agent
```

For Windows:

```
<Agent_Instance_Home>\bin\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 25](#).

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](#)
- [Managing Saved Searches](#)
- [Locating Management Agent Log and Trace Files](#)
- [Locating and Configuring Oracle Management Service Log and Trace Files](#)
- [Monitoring Log Files](#)
- [Configuring Log Archive Locations](#)

25.1 Managing Log Files

Many Enterprise Manager components generate log files containing messages that record errors, notifications, warnings, and traces.

[Table 25-1](#) describes the columns in the Log Message table. For any given component, the optional column may not be populated in the message.

Table 25–1 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.
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 <code>oracle</code> 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.

Table 25–1 (Cont.) Message Columns

Column Name	Description
Archive	Values are Yes or No. If the checkbox is checked, the message is collected from the archive location. Otherwise, the message is collected from the live system.
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](#)

25.1.1 Viewing Log Files and Their Messages

You can use Enterprise Manager Cloud Control to view messages across log files.

In particular, when you navigate in the context of a farm or domain, then the logs that you can view and search are filtered to just those associated with that farm or domain. When you navigate to Logs by way of the Enterprise menu, you can pick and choose exactly what targets you want to view and search logs against. You could also, for example pick multiple WebLogic Server targets that span across domains/farm.

For example, to view the log files and their messages:

1. From the **Enterprise** menu, select **Monitoring**, select **Logs**, then select the target from the popup target selector.

or

The Logs menu is available at the individual target label and at the parent target level. For example, for WebLogic server and for other j2ee components, the logs menu can be accessed by choosing **Logs** from the **Targets** menu. The same is applicable for parent targets like domain and farm targets.

2. In the context of a farm or domain, expand **Selected Targets** and in the row for a particular component or application, click the **Target Log Files** icon.

When you are in the context of the Enterprise menu, add targets to the Target table and 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 target.

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 and download the log file from this page.

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. When you are in context of one domain or one farm and looking at logs, the related messages are confined to that one domain or one farm. For example, 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.

When trying to view log messages, you may see the following error:

Logging Configuration is missing or invalid for the targets (). Also, make sure that these targets are up and EM User has the CONFIGURE_TARGET privilege on the corresponding domains.

To ascertain which method to use to fix the problem, choose one of these three alternatives:

- The domain's Administration Server is down. To resolve the problem, start the Administration Server and try viewing log messages again.
- The Managed Server for which you are trying to view log messages is down. To resolve the problem, start the Managed Server and try viewing log messages again.
- The Enterprise Manager Cloud Control administrator who is trying to access log messages does not have the necessary target privileges to do so. In order to view log messages, the administrator must have been granted the target privilege "Configure target" for the corresponding WebLogic Domain target. Talk to your Oracle Enterprise Manager site administrator or super administrator regarding whether or not you have this privilege.

25.1.1.1 Restricting Access to the View Log Messages Menu Item and Functionality

You can restrict which administrators in Oracle Enterprise Manager Cloud Control have access to the View Log Messages menu item and its corresponding functionality. You can grant a target privilege labeled "Ability to view Fusion Middleware Logs" to administrators and/or roles. This target privilege is applicable to all Oracle Fusion Applications related and Oracle Fusion Middleware related target types. This target privilege is automatically included as part of the following other target privileges: Operator Fusion Middleware, Operator, and Full. Consequently, you can grant an administrator one of the following privileges in order for him/her to be able to view log messages for Oracle Fusion Applications related and Oracle Fusion Middleware related log files:

- Ability to view Fusion Middleware Logs target privilege
- Operator Fusion Middleware target privilege
- Operator target privilege
- Full target privilege

To grant the ability to an administrator to view the Fusion Middleware Logs target privilege, follow these steps:

1. Log in to the Oracle Enterprise Manager 12c Cloud Control console as a super administrator.
2. From the **Setup** menu, choose **Security**, then **Administrators**.
3. Select the appropriate administrator and click **Edit**.
4. Click **Next** twice to arrive on the Target Privileges page of the wizard.
5. Scroll down the page and click **Add** in the Target Privileges section of the page.
6. From the **Search and Add: Targets** popup dialog, select the appropriate targets for which the administrator should have access to view logs. Click **Select**.

7. From the Target Privileges section of the Target Privileges page of the wizard, select the targets to which you want to grant the “Ability to view Fusion Middleware Logs” target privilege and select **Grant to Selected**. Notice that the default target privilege automatically given for this target is View.
8. Select the **Ability to view Fusion Middleware Logs** target privilege and click **Continue**. Notice that the “Ability to view Fusion Middleware Logs” target privilege is also included as part of other target privileges (for example, Operator target privilege). So, depending on the responsibilities of the administrator, you may want to grant the Operator target privilege to the administrator.
9. Notice on the Target Privileges page of the wizard the appearance of the new privilege. Click **Review** and then **Finish** to conclude the operation.

25.1.2 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 10 minutes.

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](#)

25.1.2.1 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 **Enterprise** menu, select **Monitoring**, select **Logs**, then select the target from the popup target selector.

or

From the **Targets** menu, select **Middleware**, click a farm. From the **Farm** menu, select **Logs**, then select **View Log Messages**.

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 10 minutes.
 - **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** - You can select an operator, such as **contains** and then enter a value to be matched.

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)**.

- **Application ID** - Groups messages related to a particular application.
- **ECID + Relationship ID** - Groups messages by Execution Context (ECID) and Relationship ID (RID) which enables you to use log file entries to correlate messages from one application or across application server components. By searching related messages using the message correlation information, you can examine multiple messages and identify the component that first generated the problem.
- **Host** - Groups messages associated with a particular host.
- **Host IP Address** - Groups messages associated with a particular host IP address.
- **Incident ID**
- **Message Type** - Groups messages for each target based on the message type. It displays the total number of messages available for each message type, for example, ERROR, INCIDENT ERROR, WARNING, NOTIFICATION, TRACE, and UNKNOWN for every target.
- **Message ID** - Groups messages based on the combination of Message ID, Message Type, Target, Message Level, Component, Module, and Organization.
- **Module** - Groups the classes / modules that originated the message.
- **Target**
- **Thread ID** - Groups messages by Thread ID
- **User** - Groups all messages for a particular user. For example, all the messages for user Jones will be listed before the messages for user Smith.

25.1.2.2 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**.
- **Search Archived Logs:** Enable this check box to access the log viewer. These are the archive log file locations for multiple targets you configured on the Configure Archive Locations page.

25.1.3 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 XML (.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 XML (.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. Expand the Selected Targets section because it is hidden by default. Click **Target Log Files**.

The Log Files page is displayed.

Select one of the possibly many Target Log Files icons. Select the icon that is associated with the target type log files you want to view.

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**.

25.2 Managing Saved Searches

The following sections provide information on creating, retrieving, and managing saved searches:

- [Saving Searches](#)
- [Retrieving Saved Searches](#)
- [Managing Saved Searches](#)

25.2.1 Saving Searches

Saved searches save administrators time by not having to redefine the same search again in the future. Saved searches help you in diagnosing problems faster because you are only a few clicks away from accessing a saved search as opposed to redefining the search again and again.

Note: Saved searches are per administrator. Therefore when the administrator logs out of the console, the search is stored and is available the next time the administrator logs in. In other words, saved searches that one administrator defines are not accessible by another administrator.

Once you have specified search criteria as described in Searching Log Files, you save it by clicking **Save Search** located at the top-right of the page. The name of the search is automatically created by concatenating fields used in the search, for example, Log Messages - Saved Search: "error", Last 1 hours, Incident Error,Error,Unknown.

Note: You can change the default name using the Manage Saved Search popup. This allows you to accept the default name and change it later.

25.2.2 Retrieving Saved Searches

To retrieve a saved search, follow these steps:

1. From the **Enterprise** menu, select **Monitoring**, select **Logs**, then select the target from the popup target selector.

or

From the **Targets** menu, select **Middleware**, click a farm. From the **Farm** menu, select **Logs**, then select **View Log Messages**.

or

Access the saved search from the **Favorites** menu.

The Log Messages page appears.

2. On the Logs page, click **Saved Searches** located at the top-right of the page.
3. Choose a search.

The search results populate the Search region.

25.2.3 Managing Saved Searches

To manage a saved search, follow these steps:

1. From the **Enterprise** menu, select **Monitoring**, select **Logs**, then select the target from the popup target selector.

or

From the **Targets** menu, select **Middleware**, click a farm. From the **Farm** menu, select **Logs**, then select **View Log Messages**. You can manage the saved searches pertaining to the target context only.

or

Access the saved search from the **Favorites** menu and select **Manage Favorites**. You can manage all the log-saved searches which you have created irrespective of the context. You can see all the saved searches.

The Log Messages page appears.

2. On the Logs page, click **Saved Searches** located at the top-right of the page.
3. On the list, click **Manage Saved Searches**.

The Manage Favorites pop-up appears. You can:

- Change the name of the search.

When you select a row from the table, the name of search appears in the Name field at the bottom of the screen. You can edit the name of the search and click **OK** or you can click **Cancel**.

Note: When you click **OK**, you will only be changing the name of the search, not the saved search criteria. Once the search criteria is changed, the Save Searches button is enabled.

- Edit the search criteria.

Click the link of the saved search. The Log Viewer screen appears in the context of the saved search. Make the changes and click **Save**.

- Delete a search

Choose a search and click **Remove Selected**.

25.3 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](#)

25.3.1 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 metadata log file (`gcagent.log`)
This log file contains trace, debug, information, error, or warning messages from the agent.
- Oracle Management Agent fetchlet trace 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 metadata 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.

25.3.1.1 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
```

25.3.2 Locating the Management Agent Log and Trace Files

The log and trace files for the Management 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.

25.3.3 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.

The logging configuration syntax uses the concept of handlers (appenders in log4j terms) and loggers. A handler defines a single output file and how the file is to be managed (maximum file size, number of segments, and so on). Note that there is a default logging prefix oracle.sysman that is used for all handlers that does not specify any logging prefix. The logging properties uses the `Logger . prefix` for agent (log4j) logging configuration and `ODLLogger . prefix` for the ODL (which is based on `java.util.logger.*`) logging configuration. Beside the prefix, both systems share the same syntax. The configuration full syntax (without a `Logger` or `ODLLogger` prefix) is the following:

Table 25-2

Property Name	Description	Mandatory	Default Value
directory=<directory>	Defines the logging system (log4j or ODL) logging directory. Specifying a directory for one system does not affect the other system (setting <code>Logger . directory</code> will only affect the <code>Logger . configuration</code> but not <code>ODLLogger .</code>)	No	\$EMSTATE/sysman/log
<handler>.filename=<filename>	The filename to use for the handler. If the filename is relative it will be relative to the logging directory (see <code>directory</code> property above). An absolute file name will be used as is.	Yes	
<handler>.level=<level>	The default logging level for the handler. Possible levels are: DEBUG, INFO, WARN, ERROR, FATAL	Yes	
<handler>.totalSize=<size>	The total size in MB for all the handler file segments.	No	No limit

Table 25–2 (Cont.)

Property Name	Description	Mandatory	Default Value
<handler>.segment.count=<count>	The number of segments to use for the handler.	No	1
<handler>.logger=<logger names>	A comma delimited list of logger names that will use this handler.	No	When not specified, the default logger is used.
level.<logger name>=<level>	Set a specific logging level to the logger and all its descendants. Possible levels are: DEBUG, INFO, WARN, ERROR, FATAL	No	
additivity.<logger name>=<true or false>	If set to false, only handlers that are configured for the specific logger name will be used. Otherwise, handlers that are configured for the logger parent name will also be used.	No	true

An example of the syntax is as follows:

```
# logging properties
Logger.log.filename=gcagent.log
Logger.log.level=INFO
Logger.log.totalSize=100
Logger.log.segment.count=20

ODLLogger.wsm.level=ERROR
ODLLogger.wsm.totalSize=5
ODLLogger.wsm.segment.count=5
ODLLogger.wsm.filename=gcagent_wsm.log
```

The above log configuration sets up a handler (log) that creates a gcagent.log file (in the default logging directory) with a default logging level of INFO, total size of 100MB, uses up to 20 segments, and is configured to be used by the default logger (oracle.sysman).

25.3.3.1 Modifying the Default Logging Level

To enable DEBUG level logging for the Management Agent, set the log handler level to DEBUG (see below). And then reload the agent.

```
Logger.log.level=DEBUG
```

Alternatively, use emctl setproperty agent command as follows:

```
$ emctl setproperty agent -name 'Logger.log.level' -value 'DEBUG'
```

25.3.3.2 Setting gcagent.log

The gcagent.log is the agent main log that contain log entries from all the agent core code. The following is gcagent.log configuration:

```
Logger.log.filename=gcagent.log
Logger.log.level=DEBUG
Logger.log.totalSize=100
Logger.log.segment.count=20
```

25.3.3.3 Setting gcagent_error.log

The gcagent_errors.log is a subset of the gcagent.log and contains log messages of ERROR and FATAL levels. The logging configuration for gcagent_errors.log is specified in emd.properties. Following are the settings for gcagent_errors.log:

```
Logger.err.filename=gcagent_errors.log
Logger.err.level=ERROR
Logger.err.totalSize=100
Logger.err.segment.count=5
```

25.3.3.4 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.level.oracle.sysman.gcagent.metadata.impl.ChainedClassLoader=INFO
Logger.level.oracle.sysman.gcagent.metadata.impl.ReverseDelegationClassLoader=INFO
Logger.level.oracle.sysman.gcagent.metadata.impl.PluginLibraryClassLoader=INFO
Logger.level.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.level.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.

25.3.3.5 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 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.mdu.filename=gcagent_mdu.log
Logger.mdu.level=INFO
Logger.mdu.totalSize=100
Logger.mdu.segment.count=5
Logger.mdu.logger=Mdu

```

Note: Change the filename and logger settings only if asked by Support.

25.3.3.6 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.

25.4 Locating and Configuring Oracle Management Service Log and Trace Files

The following sections describe how to locate and configure the OMS log files:

- [About the Oracle Management Service Log and Trace 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](#)

25.4.1 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*.
- Enterprise Manager Control message file (`emctl.msg`)
This file is created by the HealthMonitor thread of the OMS when it restarts the OMS because of a critical error. This file is used for troubleshooting the OMS restart problem. It provides information such as the exact time when the OMS is restarted and which module has caused the crash.

25.4.2 Locating Oracle Management Service Log and Trace Files

The OMS Instance Base directory is `gc_inst` in the Oracle Middleware Home (middleware home). This directory stores all log and trace files related to OMS 12c.

You can choose to change this, if you want, in the installer.

For example, if the Middleware home is `/u01/app/Oracle/Middleware/`, then the instance base location is `/u01/app/Oracle/gc_inst`. You can choose to change this, if you want, in the installer. However, you can change it for only advanced installation and not for simple installation.

However, you can change it for only advanced installation and not for simple installation.

25.4.3 Controlling the Size and Number of Oracle Management Service Log and Trace Files

OMS log and trace files increase 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 25-3](#):

```
emctl set property -name <property> -value <property value> -module logging
```

The above command will set the property for all OMSes. If you want to set it for a single OMS, then specify an extra option `-oms_name` as follows:

```
emctl set property -name <name> -value <value> -module logging -oms_name
example.us.oracle.com:portnumber_Management_Service
```

To set it for the current OMS, use the property `-oms_name local_oms`. To set it for any other OMS, you can provide the name of that OMS. The OMS name has to be similar to `example.us.oracle.com:portnumber_Management_Service`.

Note: In Oracle Enterprise Manager Cloud Control 12c, you do not have to restart OMS for the changes to take effect.

Note: In Oracle Enterprise Manager Cloud Control 12c, `emctl set property` by default sets the logging properties for all the OMS. To set the property for only one OMS, use the `-oms_name` option.

Table 25–3 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 <i>emoms.log</i> log file. The size of the log is specified in units of bytes. This property is also applicable for <i>emoms_pbs.log</i> .	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 <i>emoms_pbs.log</i> . 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 <i>emoms.trc</i> log file. This property is also applicable for <i>emoms_pbs.trc</i> .	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 <i>emoms_pbs.trc</i> .	log4j.appender.emtrcAppender. MaxBackupIndex=10

25.4.4 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
```

The above command will change the log level for all OMS, unless *-oms_name* option is specified.

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.

25.4.5 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 25–4](#) lists the location of the log files for some components.

Table 25–4 Component Log File Location

Component	Location
Oracle HTTP Server (OHS)	<code><EM_INSTANCE_BASE>/<webtier_instance_name>/diagnostics/logs/OHS/<ohs_name></code> For example, <code>/u01/app/Oracle/gc_inst/WebTierIH1/diagnostics/logs/OHS/ohs1</code>
OPMN	<code><EM_INSTANCE_BASE>/<webtier_instance_name>/diagnostics/logs/OPMN/<opmn_name></code> For example, <code>/u01/app/Oracle/gc_inst/WebTierIH1/diagnostics/logs/OPMN/opmn</code>

Table 25–4 (Cont.) Component Log File Location

Component	Location
Oracle WebLogic	<p>The log data from WebLogic will be at:</p> <pre><EM_INSTANCE_BASE>/user_projects/domains/<domain_name>/servers/<SERVER_NAME>/logs/<SERVER_NAME>.log</pre> <p>This log can be restricted, rotated by size, time, and other conditions from the WebLogic Console. The default settings are:</p> <ul style="list-style-type: none"> ■ In production mode, they are rotated at a default of 5MB. ■ The log level is WARNING. ■ The number files are restricted to 10. <p>For example,</p> <pre>/u01/app/Oracle/gc_inst/user_projects/domains/GCDomain/servers/EMGC_OMS1/logs/EMGC_OMS1.log</pre> <p>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:</p> <pre><EM_INSTANCE_BASE>/user_projects/domains/<domain_name>/servers/<SERVER_NAME>/logs/<SERVER_NAME>.out</pre> <p>For example,</p> <pre>/u01/app/Oracle/gc_inst/user_projects/domains/GCDomain/servers/EMGC_OMS1/logs/EMGC_OMS1.out</pre> <p>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.</p>

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*.

25.5 Monitoring Log Files

You can use Log File Monitoring to monitor WebLogic Server and Application Deployment log files for specific patterns. You can set up Cloud Control to receive alert notifications in the context of targets when patterns are found. This allows you to be more proactive and learn of problems as an administrator before end users discover them.

Use the following topics to learn how to set up and use Log File Monitoring:

- [About Log Viewer](#)
- [Overview of WebLogic Server and Application Deployment Log File Monitoring](#)
- [Enabling Log File Monitoring](#)
- [Configuring Log File Monitoring](#)
- [Viewing Alerts from Log File Monitoring](#)

25.5.1 About Log Viewer

Log Viewer enables administrators to view, search, and download middleware-related log files regardless of where the files reside on disk. Complex search criteria can be specified and saved for future reference in order to help administrators quickly diagnose performance problems across multiple middleware components spanning multiple Fusion Middleware Farms and WebLogic Domains.

When accessing Log Viewer, default search criteria is specified for the selected target type. The administrator can then refine the search criteria based on diagnostic requirements for the particular Fusion Middleware Farm. By using the Add Fields button, you can refine the search criteria to include:

- Selecting one or more member targets of the Fusion Middleware Farm
- Specifying the date range
- Selecting the message types
- Specifying the messages to be searched
- Specifying the ECIDs to be searched
- Specifying the application name
- Specifying the user name

Once the search criteria has been defined, the administrator clicks on the search button.

The administrator modifies the search as needed and clicks the **Save Search** button on the Log Viewer.

The search criteria specified, including the targets against which the search was performed, is then saved to the Management Repository for the currently logged in administrator.

You can click on the **Saved Searches** button to retrieve and apply a previously stored Search Criteria.

You can click on the Manage Saved Searches and bring up a pop-up to edit or delete the previously Saved Search Criteria.

25.5.2 Overview of WebLogic Server and Application Deployment Log File Monitoring

You can use Log File Monitoring to monitor WebLogic Server and Application Deployment log files for specific patterns and thereby reduce troubleshooting time. You can set up Cloud Control to receive alert notifications in context of targets when patterns are found.

The Log File Monitoring metric, Log File Pattern Matched Line Count for WebLogic Server and Application Deployment target types allows you to monitor one or more log files for the occurrence of one or more perl patterns. In addition, you can specify a perl pattern to be ignored for the log file. Periodic scanning, which occurs by default every 60 minutes, is performed against any new content added since the last scan. Lines matching the ignore pattern are ignored first, then lines matching specified match patterns result in one record being uploaded to the repository for each pattern. You can set a threshold against the number of lines matching the given pattern. File rotation will be handled within the given file.

You can also use the monitoring templates functionality, which allows an administrator to configure a metric once in a template and then apply the template to several WebLogic Server or Application Deployment targets at once, rather than having to configure each WebLogic Server log file monitoring metric individually.

If you are currently using log file monitoring via the Host target type, you should configure log file monitoring via the Fusion Middleware related target type instead so you can see alerts in context of a Fusion Middleware target.

Prerequisites to Use Log File Monitoring

Log File Monitoring requires a local Management Agent monitoring target. In other words, the host on which the log files you want to monitor reside must have an agent installed and running. The OS user who installed the agent must have read access to directories where the monitored log files reside. Log file monitoring is disabled by default. You must enable it in order to use this feature.

25.5.3 Enabling Log File Monitoring

Log File Monitoring is disabled by default. To enable Log File Monitoring, follow these steps:

1. From the target menu, select **Monitoring**.
2. Choose **Metric and Collection Settings**.
3. Under the Log File Monitoring row, click the **Disabled** link to change the setting to Enabled.
4. On the Edit Collection Settings: Log File Monitoring page, click **Enable** in the Collection Schedule section.
5. Click **Continue**.

Once you enable and configure Log File Monitoring, the default collection schedule is set for every 60 minutes.

25.5.4 Configuring Log File Monitoring

To configure Log File Monitoring, follow these steps:

1. From the target menu, choose **Monitoring**.
2. From the Monitoring menu, select **Metric and Collection Settings**, then choose the **Log File Pattern Matched Line Count** metric.
3. Click the **Edit** icon on the right.

Cloud Control displays the Edit Advanced Settings:Log File Pattern Matched Line Count page.

4. Click **Add** to add new object(s) in order to specify settings for the log files to be monitored.

Add objects to the Monitored Objects table. The table lists all Log File Name/Match Pattern in Perl/Ignore Pattern in Perl objects monitored for this metric. You can specify different threshold settings for each Log File Name/Match Pattern in Perl/Ignore Pattern in Perl objects. The Reorder button specifies which log file to scan first.

In past releases, if % was provided, the text was not ignored and all lines read from the file were included for pattern matching. However, this behavior has been updated wherein only "" (an empty string) is the prescribed method not to ignore any lines. However for backward compatibility "%" will still be considered as the equivalent to the "" string.

Good examples:

```
/u01/middleware/user_projects/domains/riddles_
domain/servers/ManagedServer_1/logs/access.log
```

```
C:\u01\middleware\user_projects\domains\riddles_
domain\severs\ManagedServer_1\logs\access.log
```

Bad example:

```
/u01/middleware/user_projects/domains/riddles_  
domain/servers/ManagedServer_1/logs/%.log
```

The Match Pattern in Perl value specifies the pattern that should be monitored in the log file. Perl expressions are supported in this field, and case is ignored.

Examples:

- FATAL - This pattern will be true for any lines containing fatal
- *fatal.*critical.* - This pattern will be true for any lines containing fatal and critical

The Ignore Pattern in Perl value specifies the pattern that should be ignored in the log file. If the Ignore Pattern in Perl field has a default value of % in the field, you should remove the default value if nothing should be ignored. Perl expressions are supported in this field, and case is ignored.

The Warning Threshold and Critical Threshold values should be set to a number such that if the pattern occurs in the log file the specified number of times within the collection schedule, then an alert will be triggered. If the number of occurrences is specified in the advanced settings, then this factors into when alert is raised.

For example, if you set the critical threshold to 1 (if pattern found more than 1 time in log file, it is critical alert) and the number of occurrences to 2, then a critical alert is raised only when the pattern is found more than once in the log file within 2 consecutive collections.

Once log file monitoring is enabled and configured, you can include the 'Log File Pattern Matched Line Count' metric as part of a Monitoring Template. Log file locations must be the same across targets to which the template is applied. You can apply the template to multiple WebLogic servers or Application Deployment targets at once rather than setting monitoring settings individually on a per-target basis.

If after configuring the Log File Monitoring metric the log file contains the specified patterns but the alerts are not generated in the OMS, you should do the following:

- Check whether the log file name contains a perl pattern.
- Check whether the ignore pattern contains an asterisk (*). Providing an asterisk in the ignore pattern field will also ignore all the lines which include the matched patterns.

Configuration Issues

If an error message displays indicating that logging configuration is missing or invalid for certain targets, you can try the following options.

First, the WebLogic Domain that you are accessing may not be Oracle JRF (Java Required Files) enabled. Oracle JRF consists of components not included in the Oracle WebLogic Server installation and that provide common functionality for Oracle business applications and application frameworks. To view log messages, the target must be Oracle JRF enabled. To check to see if your WebLogic Domain, for example, is Oracle JRF enabled, perform the following steps:

1. From the WebLogic Domain menu, select Target Setup submenu and then Monitoring Configuration.
2. On the Monitoring Configuration page for the domain, look for the property labeled "Can Apply JRF". The value for this property could be true or false. If the value is false, then the domain is not Oracle JRF enabled.

If the value of the “Can Apply JRF” property is true for the domain, this does not necessarily mean that all managed servers within the domain are Oracle JRF enabled. If you are unable to access log messages in the context of a specific managed server, then navigate to the relevant managed server’s Monitoring Configuration page. From the Monitoring Configuration page, look for the property “Is JRF Enabled”. The value for this property could be true or false. If the value is false, then the managed server is not Oracle JRF enabled.

Second, the Enterprise Manager Cloud Control administrator who is trying to access log messages does not have the necessary target privileges to do so. In order to view log messages, the administrator must have been granted the target privilege “Ability to view Fusion Middleware Logs” for the corresponding target. Talk to your Oracle Enterprise Manager’s site administrator or super administrator regarding whether you have this privilege or not. Refer to later questions in this document for additional details on this target privilege and granting the privilege to administrators.

25.5.5 Viewing Alerts from Log File Monitoring

Alerts generated from the Log File Pattern Matched Line Count metric appear on the home page of the target or the Alert History page.

Triggered alerts must be manually cleared.

25.6 Configuring Log Archive Locations

You can configure the host, its credentials, and archive location information for a WebLogic domain and for all targets under the domain. You can either configure everything collectively under the target at the same time, or you can configure the targets individually.

To configure all of the targets at the same time, follow these steps:

1. From the WebLogic domain home page, select **Logs** from the WebLogic Domain menu, then select **Configure Archive Locations**.

The Configure Archive Locations page appears.

2. Select the WebLogic domain in the table, then click **Assign Host Credentials**.

An Assign Host Credentials pop-up appears.

3. Provide the requisite information and make sure that the Apply Above Host Credentials to Child Targets check box is enabled, then click **OK**.

The host name you selected now appears in the Host column of the Configure Archive Locations page, and the column also displays this host for all of the child targets.

4. Click **Assign Archive Location**.

A Remote File Browser pop-up appears.

5. Double-click a directory name to enter in the host name field, then repeat this process for each sub-directory that you want to in the field. Click **OK** when you have finished.

The directory location you selected now appears in the Archive Location column of the Configure Archive Locations page, and the column also displays this location for all of the child targets.

To configure the targets separately, follow the procedure above, except select a particular target rather than the WebLogic domain.

Part VII

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.

- [High Availability Solutions](#)
- [Enterprise Manager High Availability](#)
- [Enterprise Manager Disaster Recovery](#)
- [Backing Up and Recovering Enterprise Manager](#)

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 monitored by the agent no longer being monitored by Enterprise Manager.
- Management Service failure
Results in downtime for Enterprise Manager.
- Management Repository failure
Results in downtime for Enterprise Manager.
- Software Library Failure
Results in a sub-set of Enterprise Manager operations being unavailable. These operations include self-update and provisioning and patching operations including Agent deployment.

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.

26.1 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>
- Enterprise Manager 12c Framework and Infrastructure Web site
<http://www.oracle.com/technetwork/oem/frmwrk-infra-496656.html>
- Support Note 330072.1: "How To Configure Enterprise Manager for High Availability "

26.2 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 availability 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.

26.2.1 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 four example high availability levels for Oracle Enterprise Manager installations as well as general resource requirements.

Table 26–1 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 26–1 (Cont.) Enterprise Manager High Availability Levels

Level	Description	Minimum Number of Nodes	Recommended Number of Nodes	Load Balancer Requirements
Level 4	<p>OMS on the primary site in Active/Active Configuration. Repository deployed using Oracle RAC.</p> <p>Duplicate hardware deployed at the standby site.</p> <p>DR for OMS and Software Library using Storage Replication between primary and standby sites.</p> <p>Database DR using Oracle Data Guard.</p> <p>Note: Level 4 is a MAA Best Practice, achieving highest availability in the most cost effective, simple architecture.</p>	4	8	<p>Required: Local Load Balancer for each site.</p> <p>Optional: Global Load Balancer</p>

26.3 Comparing Availability Levels

The following tables compare the protection levels and recovery times for the various high-availability levels.

Table 26–2 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 26–3 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.

26.4 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	<i>Oracle Enterprise Manager Basic Installation Guide</i> and the <i>Oracle Enterprise Manager Advanced Installation and Configuration Guide</i>
Level 2	<i>Oracle Enterprise Manager Basic Installation Guide</i> and the <i>Oracle Enterprise Manager Advanced Installation and Configuration Guide</i> PLUS <ul style="list-style-type: none">■ 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	<i>Oracle Enterprise Manager Basic Installation Guide</i> and the <i>Oracle Enterprise Manager Advanced Installation and Configuration Guide</i> PLUS <ul style="list-style-type: none">■ Oracle Management Service 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	<i>Oracle Enterprise Manager Basic Installation Guide</i> and the <i>Oracle Enterprise Manager Advanced Installation and Configuration Guide</i> PLUS <ul style="list-style-type: none">■ Configuring a Standby Database for the Management Repository■ Management Service Disaster Recovery

Enterprise Manager High Availability

This chapter discusses best practices for installation and configuration of each Cloud Control component and covers the following topics:

- [Agent High Availability](#)
- [Repository High Availability](#)
- [Oracle Management Service High Availability](#)

27.1 Agent High Availability

The following sections discuss best practices for installation and configuration of the Management Agent.

27.1.1 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.

27.1.2 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.

27.1.3 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.

27.2 Repository High Availability

The following sections document best practices for repository configuration.

27.2.1 General Best Practice for 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.

- Choose Automatic Storage Management (ASM) as the underlying storage technology.
- 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.

27.2.2 Configuring RAC for the Management Repository

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 Management Service. If you have multiple Management Services configured, this command must be run on each 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 descriptor as above>
```

27.3 Oracle Management Service High Availability

The following sections document the configuring the OMS for high availability.

27.3.1 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.

27.3.1.1 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).

27.3.1.2 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.

27.3.1.3 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.

27.3.1.4 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 supported.**

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.

27.3.1.5 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

27.3.1.6 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)
```

27.3.1.7 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
```

27.3.1.8 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
```

You can also provide the ORACLE_HOSTNAME when prompted for this information from in Enterprise Manager runInstaller UI.
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.

27.3.1.9 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 ["Performing Switchover and Failover Operations"](#) on page 28-14.

27.3.1.10 Repository Host Name Change

In general Enterprise Manager does not store IP address lookups. So if the IP address of the repository host changes but the hostname remains the same, no configuration change is required in Enterprise Manager. Ensure that the OMS(s) hosts can communicate with the repository host with the new IP address.

However if the Repository hostname changes, please use the following procedure to make appropriate configuration changes so that the OMS(s) can communicate with the repository with new hostname.

1. Stop all OMS instances using the command

```
$emctl stop oms -all
```
2. Stop the Agent running on the database host.

```
$emctl stop agent
```
3. Perform host name change for the machine either by local resolution /etc/hosts or DNS based resolution or a combination of both.
4. Follow Oracle Database Administrator guide to configure the database for hostname change.
5. Correct the connect descriptor for the repository by running the following commands:

On the first Oracle Management Service

```
emctl start oms -admin-only
```

On each Oracle Management Service

```
$emctl config oms -store_repos_details -repos_conndesc  
<connect descriptor> -repos_user sysman
```

Where `<connect descriptor>` is the repository connect descriptor using the new hostname.

Note: If Single Client Access Name (SCAN) is configured on the Repository database, it may be possible to change the repository hostname without needing to change the connect descriptor of clients like OMS. Please consult Oracle Database Administrator Guide.

6. Stop and Start OMS instances using the command:

```
$emctl stop oms -all
```

```
$emctl start oms
```

7. The agent installed on the repository host using the old hostname can no longer be used. Install new agent on database host, if RAC database is configured then install new agent on each node. Follow below step to install new agent <include agent installation url for reference>

1. Remove the agent_inst directory .

2. Configure the agent 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>
```

8. Relocate the Repository database target to the new Management Agent configured using the following commands:

```
<OMS_HOME>/bin/emcli login -username=sysman
```

```
<OMS_HOME>/bin/emcli sync
```

```
<OMS_HOME>/bin/emctl config repos -agent <new agent on host,
e.g myNewReposHost.example.com:3872>
```

9. Change the monitoring configuration for the OMS and Repository target: by running the following command from the OMS:

```
$emctl config emrep -conn_desc
```

10. Verify that the site is fully operational.

11. Log in to Cloud Control. Navigate Setup, Manage Cloud Control and click on Health Overview.

12. Delete old agent target entry.

13. Verify that the site is fully operational.

27.3.1.11 Multiple OMS, SLB, First OMS Host Name Change

In this scenario, you have multiple OMSs fronted by Server Load Balancer. OMS configuration backed up using the emctl exportconfig oms command on the primary OMS running the WebLogic AdminServer.

1. Copy the emkey to the Management Repository by running the following command on the OMS:

```
$<ORACLE_HOME>/bin/emctl config emkey -copy_to_repos
```

2. Export the configuration from the OMS.

```
$<ORACLE_HOME>/bin/emctl exportconfig oms -dir <location for the export file>
```

After the configuration is exported, do not make any configuration changes to the OMS until hostname change is configured.

3. Deconfigure and delete the OMS

1. Stop the OMS \$<ORACLE_HOME>/bin/emctl stop oms -all

2. Stop the agent running \$<AGENT_HOME>/bin/emctl stop agent

3. Delete the OMS `<OMS_HOME>/bin/omsca delete -full`

For example:

```
/u01/app/Oracle/Middleware/oms/bin/omsca delete -full
```

Note : You are prompted to confirm your action, and furnish the AdminServer credentials and the repository database details such as the database host name, listener port, SID, and password. Once you provide the required details, the command automatically stops the OMS, Oracle WebLogic Server, and also Oracle WebTier.

4. Delete following file from the Oracle Home

Also, delete Instance home directory `gc_inst`, if it exists.

5. Make sure no process running on the OMS host.

```
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

6. Perform hostname change for the machine either by local resolution `/etc/hosts` or DNS based resolution or a combination of both.
7. 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> -EM_INSTANCE_HOST <new hostname> -AS_
HOST <new hostname>
```

Note: The `-backup_file` to be passed must be the latest file generated from `emctl exportconfig oms` command in Step 2.

Provide new hostname for `-EM_INSTANCE_HOST` and `-AS_HOST`

8. Add the new hostname to the SLB virtual server pools and remove the old OMS entry.
9. Start the OMS

```
<OMS_HOME>/bin/emctl start oms
```

10. Configure the Management Agent.

1. Remove the `agent_inst` directory .
2. Configure the agent 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>
```

11. The additional OMS must be re-enrolled so that they can communicate with the admin server running on the new host. Do the following on each additional management service.

```
enroll oms -as_host <newhostname> -as_port <AdminServer_ssl_
port>
```

Note: You will be prompted to provide password for admin and node manager.

12. Login to EM console. Navigate to Setup, Provisioning and Patching and click on Software library. If software library location is configured, verify it is accessible. If not accessible then follow software library configuration document to either migrate or configure new.

<swlib configuration location>

13. Verify that the site is fully operational.

27.3.1.12 Multiple OMS, SLB, Additional OMS Host Name Change

In this scenario, you have multiple OMS sites where the OMS instances are fronted by an SLB. OMS configuration backed up using the `emctl exportconfig oms` command on the first OMS. Hostname change is required on Additional OMS host.

1. Stop the OMS using below command

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

2. Stop the agent running on the host

```
$<AGENT_HOME>/bin/emctl stop agent
```

3. Deconfigure and delete the OMS:

```
Delete the OMS $<OMS_HOME>/bin/omsca delete -OMSNAME  
<Additional OMS Name>
```

For example:

```
/u01/app/Oracle/Middleware/oms/bin/omsca delete -OMSNAME  
EMGC_OMS2
```

Note: You are prompted to confirm your action, and furnish the AdminServer credentials and the repository database details such as the database host name, listener port, SID, and password. Once you provide the required details, the command automatically stops the OMS, Oracle WebLogic Server, and also Oracle WebTier.

4. Delete following file from the Oracle Home

Delete the Instance home directory `gc_inst`, if it exists.

5. Make sure no process running on the OMS host.

```
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

6. Perform hostname change for the machine either by local resolution `/etc/hosts` or DNS based resolution or a combination of both.
7. Ensure that shared software library locations are accessible.
8. Install an Management Agent on the required host using new hostname.

1. Remove the `agent_inst` directory .

2. Configure the agent 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>
```

9. Use the Additional OMS deployment procedure to configure a new additional OMS. See the Oracle® Enterprise Manager Cloud Control Basic Installation Guide for more information.
10. Verify that the site is fully operational.

27.3.2 Configuring the Cloud Control OMS for High Availability Failover

The following sections discuss how to configure the OMS for high availability failover in an Active/Active environment using a Server Load Balancer.

27.3.2.1 Configuring the Software Library

The software library location must be accessible by all active Management Services. If the software library is not configured during installation, it 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.

27.3.2.2 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.

27.3.2.2.1 SLB Setup

Use the following table as reference for setting up the SLB with Cloud Control Management Services.

Table 27–1 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 27–2 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 27–3 Virtual Servers

Virtual Server Name	Usage	Virtual Server Port	Pool
vs_gcsu4900	Secure upload	4900	pool_gcsu4900

Table 27–3 (Cont.) Virtual Servers

Virtual Server Name	Usage	Virtual Server Port	Pool
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 27–4 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/consoleStatus.jsp Receive String: Enterprise Manager Console is UP	HostA:7799 HostB:7799
mon_gcuc7788 (optional)	Type: http Interval: 5 Timeout: 16 Send String: GET /em/consoleStatus.jsp Receive String: Enterprise Manager Console is UP	HostA:7788 HostB:7788

Note: Some Load Balancers require <CR><LF> characters to be added explicitly to the Send String using literal "\r\n". This is vendor-specific. Refer to your SLB documentation for details.

27.3.2.2.2 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
-host slb.example.com
-secure_port 4900
-slb_port 4900
-slb_console_port 443
-console
[-lock] [-lock_console]
```

Output:

```
Oracle Enterprise Manager Cloud Control 12c Release 3
Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved.
Securing OMS... Started.
Enter Enterprise Manager Root (SYSMAN) Password :
Enter Agent Registration Password :
Securing OMS... Successful
Restart OMS
```

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
```

27.3.2.3 Installing Additional Management Services

There are two ways to install additional Management Services:

- Using the Add Additional OMS Deployment Procedure (preferred method). For more information about using this deployment procedure, see the chapter on Adding Additional Oracle Management Service in the *Oracle® Enterprise Manager Cloud Control Basic Installation Guide*.
- Installing Additional Oracle Management Service in Silent Mode (alternative method). For more information about silent mode installation, see the Oracle® Enterprise Manager Cloud Control Advanced Installation and Configuration guide.

Enterprise Manager Disaster Recovery

While high availability typically protects against local outages such as application failures or system-level problems, disaster tolerance protects against larger outages such as catastrophic data center failure due to natural disasters, fire, electrical failure, evacuation, or pervasive sabotage. For Maximum Availability, the loss of a site cannot be the cause for outage of the management tool that handles your enterprise.

Maximum Availability Architecture for Enterprise Manager involves deploying a remote failover architecture that allows a secondary data center to take over the management infrastructure in the event that disaster strikes the primary management infrastructure.

Note: The procedure for creating a standby OMS differs according to the version of Enterprise Manager. The

12.1.0.2 and earlier - Standby OMSs using Standby WebLogic Domain. See [Appendix D, "Standby OMSs using Standby WebLogic Domain"](#) for more information

12.1.0.3 and later - Standby OMSs using Storage Replication. Discussed in this chapter.

This chapter covers the following topics:

- [Disaster Recover Overview and Topology](#)
- [Design Considerations](#)
- [Setting Up Management Repository Disaster Recovery](#)
- [Setting Up Management Service and Software Library Disaster Recovery](#)
- [Performing Switchover and Failover Operations](#)
- [Keeping the Standby Site in Sync with the Primary](#)

28.1 Disaster Recover Overview and Topology

The Disaster Recovery solution for a Cloud Control deployment involves replicating the OMS, Software Library and Repository components to a standby site. This solution can be combined with the high availability solution described earlier to ensure that failures ranging from component failure to a complete site outage can be recovered from with minimal disruption to the availability of Cloud Control.

Some advantages to using a replicated standby configuration are:

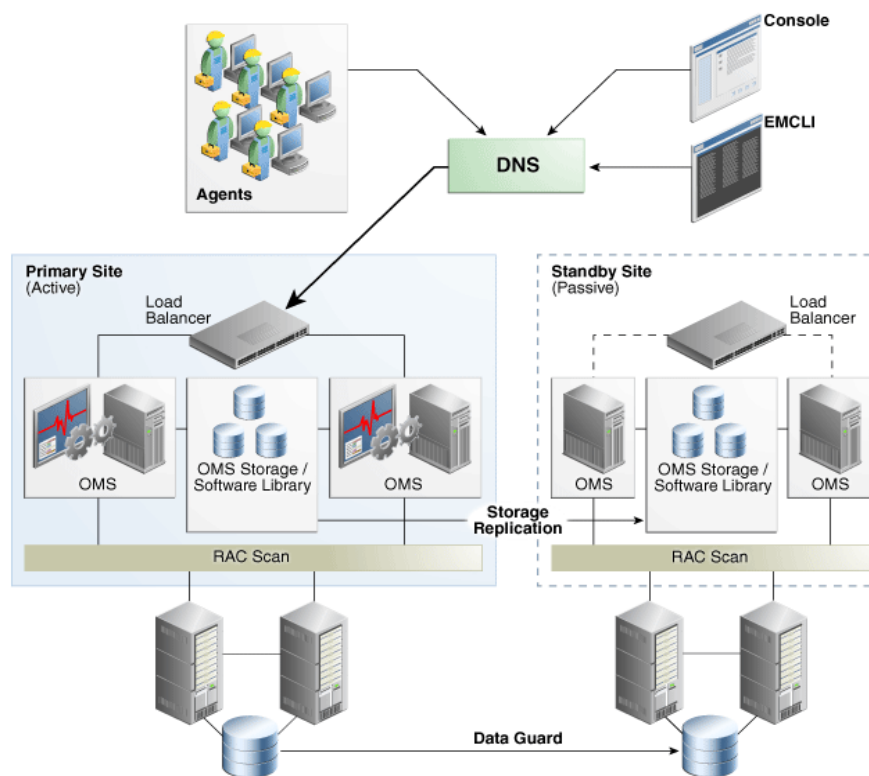
- OMS patching and upgrade only needs to be performed at one site.

- Plug-ins only need to be installed at one site.
- Keeping primary and standby sites synchronized is greatly simplified.

Only one site can be active at any given time. In this configuration, the application traffic is directed to the appropriate site by a DNS entry that is updated when the standby site is activated. 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.

A complete implementation of the Enterprise Manager Cloud Control Disaster Recovery solution is shown in the following figure.

Figure 28–1 Disaster Recovery Topology



Key aspects of the DR solution shown in the figure are:

- The solution has two sites. The current production site is running and active, while the second site is serving as a standby site and is in passive mode.
- Oracle Data Guard Physical Standby is used to replicate Repository transaction from primary site to standby site.
- The primary OMS hostnames should be resolvable to the IP addresses of the corresponding standby hosts at the standby site.
- OMS hosts on each site have mount points defined for accessing the shared storage system for the site.
- On the primary site, OMS software binaries and configuration files, Oracle Inventory, Software library and Agent binaries and configuration files for all OMS(s) are on the shared storage.

- Storage replication technology is used to copy the OMS file systems from the production site's shared storage to the standby site's shared storage.
- It is not necessary to perform any Oracle software installations at the standby OMS site hosts. When the production site storage is replicated at the standby site storage, the equivalent Oracle home directories and data are written to the standby site storage.
- Schedule incremental replications at a specified interval. The minimum recommended interval is once a day for the OMS filesystem as the configuration does not change very often. Additionally, you should force a manual synchronization whenever you make a change to the middle tier configuration at the production site (for example, if you apply a patch on the production site). The software library on the other hand requires a more frequent replication schedule, as frequent as the network infrastructure would allow. Oracle recommends continuous replication or a minimum frequency of 1 hour.
- Before forcing a manual synchronization, you should take a snapshot of the site to capture its current state. This ensures that the snapshot gets replicated to the standby site storage and can be used to roll back the standby site to a previous synchronization state, if desired. Recovery to the point of the previously successful replication (for which a snapshot was created) is possible when a replication fails.
- Client access to Cloud Control is through a virtual hostname. This is accomplished via a DNS entry or a Global Load Balancer that is configured to direct Cloud Control users and Agents to the currently active site
- When there is a failure or planned outage of the production site, you perform the following steps to enable the standby site to assume the production role in the topology:
 - Stop OMSs at the primary site
 - Failover/switchover of the database to the standby site
 - Activate OMS and Software Library replicated storage at standby site
 - Reverse storage replication direction
 - Start OMSs at standby site
 - Update DNS or global load balancer to re-route user requests to the standby site. At this point, the standby site has assumed the production role.

28.2 Design Considerations

This section describes design considerations for an Cloud Control Disaster Recovery solution for an enterprise deployment.

The following topics are covered:

- [Network Considerations](#)
- [Storage Considerations](#)
- [Database Considerations](#)
- [Starting Points](#)

28.2.1 Network Considerations

The following sections discuss network considerations that must be taken into account when implementing standby Management Services using storage replication

28.2.1.1 Planning Host Names

In a Disaster Recovery topology, the production site host names must be resolvable to the IP addresses of the corresponding peer systems at the standby site. Therefore, it is important to plan the host names for the production site and standby site. After switchover or failover from a primary site to a standby site, it should be possible to start applications on the standby hosts without requiring you to change the hostname for hosts on the standby site.

This can be achieved in either of the following ways:

- Option 1: Physical host names on primary site and alias on standby site: OMSs at the primary site are configured using physical host names and aliases for these host names are configured on the corresponding hosts at the standby site.
- Option 2: Alias host names on both sites: OMSs at the primary site are configured using an alias host name that can be configured at both the primary and standby sites.

The choice between these options would depend on your network infrastructure and corporate policies. From a setup procedure perspective, Option 1 is easier to implement if you have an existing single site Cloud Control installation which uses the physical host names as it does not require any transformation of your existing site to setup DR. Option 2 is easier to implement if you are setting up a new Cloud Control installation and start with alias host names or you have an existing Cloud Control installation using alias host names.

Host name resolution at each site can be done using either local resolution (/etc/hosts) or DNS based resolution or a combination of both. The following examples use these physical host names and IP addresses:

HOSTNAME	IP ADDRESS	DESCRIPTION
oms1-p.example.com	123.1.2.111	Physical host for OMS1 on Primary site
oms2-p.example.com	123.1.2.112	Physical host for OMS2 on Primary site
oms1-s.example.com	123.2.2.111	Physical host for OMS1 on Standby site
oms2-s.example.com	123.2.2.112	Physical host for OMS2 on Standby site

Example for Option 1: /etc/hosts configurations when OMSs are installed at primary site using primary site physical host names (oms1-p.example.com and oms2-p.example.com):

Primary Site

```
127.0.0.1    localhost.localdomain  localhost
123.1.2.111  oms1-p.example.com        oms1-p #OMS1
123.1.2.112  oms2-p.example.com        oms2-p #OMS2
```

Standby Site

```
127.0.0.1    localhost.localdomain  localhost
123.2.2.111  oms1-s.example.com        oms1-s      oms1-p.example.com #OMS1
123.2.2.112  oms2-s.example.com        oms2-s      oms2-p.example.com #OMS2
```

If the network has been configured correctly, a ping of the OMS host name from the primary site should result in a reply from the primary host, and a ping of the OMS host name from the standby site should result in a reply from the standby host.

Ping results from primary site (reply from primary site):

```
[oracle@oms1-p ~]$ ping oms1-p.example.com
PING oms1-p.example.com (123.1.2.111) 56(84) bytes of data.
64 bytes from oms1-p.example.com (123.1.2.111): icmp_seq=1 ttl=64 time=0.018 ms
```



```
64 bytes from oms1-p.example.com (123.1.2.111): icmp_seq=2 ttl=64 time=0.020 ms
64 bytes from oms1-p.example.com (123.1.2.111): icmp_seq=3 ttl=64 time=0.022 ms
```

Ping results from standby site (reply from standby site)

```
[oracle@oms1-s ~]$ ping oms1-p.example.com
PING oms1-s.example.com (123.2.2.111) 56(84) bytes of data.
64 bytes from oms1-s.example.com (123.2.2.111): icmp_seq=1 ttl=64 time=0.018 ms
64 bytes from oms1-s.example.com (123.2.2.111): icmp_seq=2 ttl=64 time=0.020 ms
64 bytes from oms1-s.example.com (123.2.2.111): icmp_seq=3 ttl=64 time=0.022 ms
```

Example for Option 2: /etc/hosts configuration when OMSs are installed using alias host names (oms1.example.com and oms2.example.com):

Primary Site

127.0.0.1	localhost.localdomain	localhost	
123.1.2.111	oms1-p.example.com	oms1-p	oms1.example.com #OMS1
123.1.2.112	oms2-p.example.com	oms2-p	oms2.example.com #OMS2

Standby Site

127.0.0.1	localhost.localdomain	localhost	
123.2.2.111	oms1-s.example.com	oms1-s	oms1.example.com #OMS1
123.2.2.112	oms2-s.example.com	oms2-s	oms2.example.com #OMS2

If the network has been configured correctly, a ping of the OMS host name from the primary site should result in a reply from the primary host, and a ping of the OMS host name from the standby site should result in a reply from the standby host.

Example:

Ping results from primary site (reply from primary site):

```
[oracle@oms1-p ~]$ ping oms1.example.com
PING oms1-p.example.com (123.1.2.111) 56(84) bytes of data.
64 bytes from oms1-p.example.com (123.1.2.111): icmp_seq=1 ttl=64 time=0.018 ms
64 bytes from oms1-p.example.com (123.1.2.111): icmp_seq=2 ttl=64 time=0.020 ms
64 bytes from oms1-p.example.com (123.1.2.111): icmp_seq=3 ttl=64 time=0.022 ms
```

Ping results from standby site (reply from standby site)

```
[oracle@oms1-s ~]$ ping oms1.example.com
PING oms1-s.example.com (123.2.2.111) 56(84) bytes of data.
64 bytes from oms1-s.example.com (123.2.2.111): icmp_seq=1 ttl=64 time=0.018 ms
64 bytes from oms1-s.example.com (123.2.2.111): icmp_seq=2 ttl=64 time=0.020 ms
64 bytes from oms1-s.example.com (123.2.2.111): icmp_seq=3 ttl=64 time=0.022 ms
```

28.2.1.2 Load Balancers Consideration

A multiple OMS site requires a server load balancer. Both Primary and Standby Sites must be fronted by their own server load balancer. See ["Configuring a Load Balancer"](#) on page 27-10. The SLB pools on the standby site will reference the IP addresses of the standby OMS hosts.

28.2.1.3 Application Virtual Host Name Consideration

A hostname through which the Cloud Control clients (agents and users) should access Cloud Control is required. When the primary site is active, this hostname should be configured in DNS to resolve to the IP address hosted by the primary site SLB. When the standby site is activated, the DNS entry should be updated so that the hostname resolves to the IP address hosted by the standby site SLB.

The DNS configuration for the Cloud Control application hostname is shown in the table below:

Table 28–1 DNS Configuration

DNS NAME	DNS RECORD TYPE	VALUE	COMMENTS
em.example.com	CNAME	slb_primary.example.com	Virtual Hostname used by Cloud Control clients to communicate with Management Service. Should point to SLB of currently active site
slb_primary.example.com	A	123.1.2.110	Primary Site SLB address
slb_standby.example.com	A	123.2.2.110	Standby Site SLB address

The DNS switchover can be accomplished by either using a global load balancer or manually changing DNS names.

- A global load balancer can provide authoritative DNS name server equivalent capabilities. One advantage of using a global load balancer is that the time for a new name-to-IP mapping to take effect can be almost immediate. The downside is that an additional investment must be made for the global load balancer
- Manually changing the DNS names. To ensure that DNS records cached by the Cloud Control clients are updated in a timely fashion after an update, it is recommended to set the TTL for the em.example.com CNAME to a low value such as 60 seconds. This will ensure that DNS changes will quickly propagate to all clients. However due to the shortened caching period, an increase in DNS requests can be observed.

28.2.2 Storage Considerations

The Disaster Recovery solution for a Cloud Control deployment involves installing the Software Library, OMS installation, Agent installation and Oracle inventory on replicated storage.

Storage Replication Requirements

Your chosen method of storage replication should support the following:

- Snapshots and consistent filesystem copies
- Ability to perform scheduled and on-demand replication between sites

The following section details the storage structure recommended by Oracle.

- Create one volume per OMS host.
- Mount the above volumes to each OMS host using the same mount point e.g. /em. On each host, this volume would contain the OMS installation, Agent installation and Oracle inventory.
- Create a consistency group for the above volumes so that consistent replication can be done for all the volumes.
- Create one volume for the software library. This volume must be mounted simultaneously to all the OMS hosts using the same mount point e.g. /swlib.
- Decide on appropriate replication frequency for the OMS filesystems and software library based on your infrastructure. Oracle recommends a minimum frequency of

24 hours for the OMS filesystem and continuous or hourly replication for the software library.

Example: The following table shows an example configuration.

Table 28–2 Storage Configuration

Volume	Mounted on Host	Mount Point	Comments
VOLOMS1	oms1-p.example.com	/em	Installation of Enterprise Manager on Primary Site OMS1
VOLOMS2	oms2-p.example.com	/em	Installation of Enterprise Manager on Primary Site OMS2
VOLSWLIB	oms1-p.example.com and oms2-p.example.com	/swlib	Software library on Primary Site OMS1 and OMS2

28.2.3 Database Considerations

This section provides the recommendations and considerations for setting up Repository databases for Disaster Recovery.

- Oracle recommends creating Real Application Cluster databases on both the production site and standby site.
- The Oracle Data Guard configuration used should be decided based on the data loss requirements of the database as well as the network considerations such as the available bandwidth and latency when compared to the redo generation. Make sure that this is determined correctly before setting up the Oracle Data Guard configuration.
- 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.
- It is strongly recommended to force Data Guard to perform manual database synchronization whenever middle tier synchronization is performed. This is especially true for components that store configuration data in the metadata repositories.
- It is strongly recommended to set up aliases for the database host names on both the production and standby sites. This enables seamless switchovers, switchbacks and failovers. For example:

```

Site      Repository Host Names      Repository Connect String
Primary   repos1-p.example.com,           (DESCRIPTION= (ADDRESS_
LIST= (FAILOVER=ON) (ADDRESS= (PROTOCOL=TCP) (HOST=repos1-p.example.com) (PORT=1521)
) (ADDRESS= (PROTOCOL=TCP) (HOST=repos2-p.example.com) (PORT=1521))) (CONNECT_
DATA= (SERVICE_NAME=EMREP)))
repos2-p.example.com

Standby   repos1-s.example.com,           (DESCRIPTION= (ADDRESS_
LIST= (FAILOVER=ON) (ADDRESS= (PROTOCOL=TCP) (HOST=repos1-s.example.com) (PORT=1521)
) (ADDRESS= (PROTOCOL=TCP) (HOST=repos2-s.example.com) (PORT=1521))) (CONNECT_
DATA= (SERVICE_NAME=EMREP)))
repos2-s.example.com

```

In the above example, after a failover or switchover operation, the OMS on the standby site must be switched to use the standby repository connection string. You can avoid changing of connect strings by optionally setting up a hostname alias for the repository database hosts. For example:

Site	Repository Host Names	Host Name Alias
Primary	repos1-p.example.com,	repos1.example.com
	repos2-p.example.com	repos2.example.com
Standby	repos1-s.example.com,	repos1.example.com
	repos2-s.example.com	repos2.example.com

Thus the connect string on each site can be the same, alleviating the need to do a change during failover or switchover.

```
(DESCRIPTION=(ADDRESS_
LIST=(FAILOVER=ON)(ADDRESS=(PROTOCOL=TCP)(HOST=repos1.example.com)(
PORT=1521))(ADDRESS=(PROTOCOL=TCP)(HOST=repos2.example.com)(PORT=152
1)))(CONNECT_DATA=(SERVICE_NAME=EMREP))).
```

28.2.4 Starting Points

Before setting up the standby site, the administrator must evaluate the starting point of the project. The starting point for designing an Enterprise Manager Cloud Control Disaster Recovery topology is usually one of the following:

- The primary site is already created, standby site is being planned
- The primary site is already created, standby site is already created using the deprecated "Standby WLS Domain" method
- No installation exists, both primary and standby sites are being planned

28.2.4.1 The primary site is already created, standby site is being planned

When the starting point is an existing primary site, the OMS installation for the primary site already exist on the file system. Also, the host names, ports, and user accounts are already defined. The following procedure must be used to transform the site and prepare it for Disaster Recovery topology.

1. Review the Network Considerations and plan your host names

If using option 1, no host name changes are required on the primary site. Prepare your standby site hosts by adding appropriate alias host names.

If using option 2, change the OMS host name (see ["Multiple OMS, SLB, First OMS Host Name Change"](#) on page 27-7 to move your existing OMS installation to use alias host names. Prepare your standby site hosts by adding the appropriate alias host names.

2. Review the Storage Considerations and move your OMS installation to shared storage

Migrate the primary site to shared storage. See ["Migrating an Existing Site to Shared Storage"](#) on page 28-14.

3. Review the Database considerations and plan your repository host names and connect descriptors

To achieve seamless failover/switchover consider if you want to use hostname alias for the repository database. If so, migrate your repository database to use alias hostname. See ["Repository Host Name Change"](#) on page 27-6.

4. Now that your primary site is ready, use the procedures in [""Setting Up Management Repository Disaster Recovery"](#) on page 28-9 and ["Setting Up Management Service and Software Library Disaster Recovery"](#) on page 28-10 to complete the DR setup.

28.2.4.2 The primary site is already created, standby site is already created using the deprecated "Standby WLS Domain" method.

1. Use the deleting standby OMS procedure to delete the Standby OMS. See *Removing Additional Standby OMS Instances in the Enterprise Manager Advanced Installation and Configuration Guide*.
2. Use the 28.2.4.1 procedure documented in ["The primary site is already created, standby site is being planned"](#) on page 28-8.

28.2.4.3 No installation exists, both primary and standby sites are being planned

When you are designing a new primary site (not using a pre-existing primary site), its easier as the site planning can be done before starting the installation of software.

1. Review the Network Considerations and plan your host names.
2. Review the Storage Considerations and prepare your storage volumes.
3. Review the Database Considerations and prepare your repository host names.
4. Now do your primary site installation using the procedures in [Chapter 27, "Enterprise Manager High Availability,"](#) taking care to using the correct host names and installing on the shared storage.
5. Now that your primary site is ready, use the "Setup Standby Repository" and "Setup Standby OMS and Software Library using storage replication" procedures to complete the DR setup.

28.3 Setting Up Management Repository Disaster Recovery

The Management Repository should use Data Guard as a Disaster Recovery solution.

28.3.1 Configuring a Standby Database for the Management Repository

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.

28.4 Setting Up Management Service and Software Library Disaster Recovery

The Disaster Recovery solution for a Cloud Control deployment involves installing the Software Library, OMS installation, Agent installation and Oracle inventory on replicated filesystem.

Standby OMSs implemented using Standby WebLogic Domain are still supported but have been deprecated and may be desupported in a future release (see My Oracle Support Note 1563541.1 for details). The recommended method for creating Standby OMSs is to use storage replication as documented in this chapter. Creating standby OMSs using a Standby WebLogic Domain is documented in [Appendix D, "Standby OMSs using Standby WebLogic Domain."](#)

Storage Replication Requirements

Your chosen method of storage replication should support the following:

- Snapshots and consistent filesystem copies
- Ability to perform an on-demand replication between sites

28.4.1 Management Service Disaster Recovery

1. Ensure that the primary OMS host names are resolvable to the IP addresses of the corresponding standby hosts at the standby site. This can be achieved in either of the following ways:
 - By installing OMSs at the primary site using physical host names and configuring aliases for these host names on the corresponding hosts at the standby site.
 - By installing each OMS using an alias host name that can be configured at both the primary and standby sites.

Host name resolution at each site can be done using either local resolution (/etc/hosts) or DNS based resolution or a combination of both.

Example /etc/hosts configurations when OMSs are installed at primary site using primary site physical host names (oms1-p.example.com and oms2-p.example.com):

Primary Site

```
127.0.0.1      localhost.localdomain
123.1.2.111    oms1-p.example.com  oms1-p #OMS1
123.1.2.112    oms2-p.example.com  oms2-p #
```

Standby Site

```
127.0.0.1      localhost.localdomain
123.2.2.111    oms1-s.example.com  oms1-s oms1-p.example.com #OMS1
123.2.2.112    oms2-s.example.com  oms2-s oms2-p.example.com #OMS2
```

Example /etc/hosts configuration when OMSs are installed using alias host names (oms1.example.com and oms2.example.com):

Primary Site

```
127.0.0.1      localhost.localdomain
123.1.2.111    oms1-p.example.com  oms1-p oms1.example.com #OMS1
123.1.2.112    oms2-p.example.com  oms2-p oms2.example.com #OMS2
```

Standby Site

```
127.0.0.1      localhost.localdomain
123.2.2.111    oms1-s.example.com  oms1-s oms1.example.com #OMS1
123.2.2.112    oms2-s.example.com  oms2-s oms2.example.com #OMS2
```

If the network has been configured correctly, a ping of the OMS host name from the primary site should result in a reply from the primary host, and a ping of the OMS host name from the standby site should result in a reply from the standby host.

Example

Ping results from primary site (reply from primary site):

```
[oracle@oms1-p ~]$ ping oms1-p.example.com
```

```
PING oms1-p.example.com (123.1.2.111) 56(84) bytes of data.  
64 bytes from oms1-p.example.com (123.1.2.111): icmp_seq=1 ttl=64 time=0.018 ms  
64 bytes from oms1-p.example.com (123.1.2.111): icmp_seq=2 ttl=64 time=0.020 ms  
64 bytes from oms1-p.example.com (123.1.2.111): icmp_seq=3 ttl=64 time=0.022 ms
```

Ping results from standby site (reply from standby site)

```
[oracle@oms1-s ~]$ ping oms1-p.example.com  
PING oms1-s.example.com (123.2.2.111) 56(84) bytes of data.  
64 bytes from oms1-s.example.com (123.2.2.111): icmp_seq=1 ttl=64 time=0.018 ms  
64 bytes from oms1-s.example.com (123.2.2.111): icmp_seq=2 ttl=64 time=0.020 ms  
64 bytes from oms1-s.example.com (123.2.2.111): icmp_seq=3 ttl=64 time=0.022 ms
```

2. Ensure that the OMS installation, Agent Installation and Oracle Inventory for each OMS at the primary site is placed on replicated storage. This can either be done by specifying replicated storage during OMS installation or by moving these components onto replicated storage after installation.

Note: If the components are moved to shared storage after installation they must retain their original pathnames.

3. Configure an application virtual host name in DNS to point to Primary site.
 - If there is a single OMS at the primary site the DNS entry for the application virtual host name should point to this OMS.
 - If there are multiple OMSs at the primary site the DNS entry for the application virtual host name should point to the SLB.
 - This host name should be configured with a short TTL value (30-60 seconds) so that it will not be cached by clients for extended periods.
4. Configure SLB at the standby site (only required if multiple OMSs are required at the standby site). See ["Configuring a Load Balancer"](#) on page 27-10 for more information. The SLB pools on the standby site will reference the IP addresses of the standby OMS hosts.
5. Resecure all Agents and OMSs using application virtual host name.

Examples

For OMS

```
emctl secure oms -sysman_pwd <sysman_pwd>  
-reg_pwd <agent_reg_password>  
-host em.example.com  
-secure_port 4900  
-slb_port 4900  
-slb_console_port 443  
-console  
-lock_upload -lock_console
```

For Agent

```
emctl secure agent -emdWalletSrcUrl https://em.example.com:4901/em
```

6. Configure the storage replication schedule for as frequently as the network infrastructure will allow (minimum every 24 hours).

Note: Refer to your storage/network documentation to determine a replication schedule that maximizes the resource utilization performance of your network infrastructure.

7. Move HTTP Lock files to local filesystem. See the *Enterprise Manager Cloud Control Advanced Installation and Configuration Guide* for more information.

28.4.2 Monitoring Standby OMS Hosts

Monitoring the availability of the standby OMS hosts is necessary to ensure that they are ready for switchover/failover operations. In order to monitor these hosts, Agents should be deployed to local filesystems on each standby OMS host.

To avoid conflicts with the components that will be started on the standby site after a switchover/failover, when deploying Agents on the standby OMS hosts the following points should be considered:

- The Agents deployed to the standby OMS hosts should not use the replicated Oracle Inventory. They should be installed using a local inventory that does not include the replicated OMS and Agent installs.
- The Agents deployed to the standby OMS hosts should be deployed on a different port to that used by the replicated Agents. This will avoid port conflicts when the replicated OMS and Agent are started on the standby OMS host.
- Regardless of which network topology is used (aliases at both sites or aliases only at the standby site), these Agents should be deployed using the physical hostnames of the standby OMS hosts.

To specify an inventory location for Agent installation, an inventory pointer file can be created and the `-invPtrLoc` flag can be used during installation.

The following example shows an inventory pointer file that specifies the inventory location as `/u01/oraInventory_standby`

Example 28–1 Inventory Pointer File

```
more /u01/oraInst_standby.loc

inventory_loc=/u01/oraInventory_standby
inst_group=dba
```

The `-invPtrLoc` flag can then be passed during Agent installation.

ORACLE Enterprise Manager Cloud Control 12c Help Log Out

Add Target

Host and Platform **Installation Details** Review

Add Host Targets : Installation Details Back Step 2 of 3 Next Cancel

On this screen, select each row from the following table and provide the installation details in the Installation Details section.

Deployment Type	Platform	Agent Software Version	Hosts	Mandatory Inputs
	Linux x86-64	12.1.0.3.0		

Linux x86-64 : Installation Details

* Installation Base Directory: /u01/agent12c

* Instance Directory: /u01/agent12c/agent_inst

* Named Credential: NC_HOST_2013-06-18-095529

Privileged Delegation Setting: /usr/bin/sudo -u %RUNAS% %COMMAND%

Port: 3972

Optional Details

Preinstallation Script: ☐ Run as Root ☐ Script on OMS

Postinstallation Script: ☐ Run as Root ☐ Script on OMS

Additional Parameters: -invPrtLoc /u01/orainst_standby.loc

28.4.3 Software Library Disaster Recovery

1. The Software Library should be located on a filesystem that is replicated using storage replication. If the Software Library is currently located on another filesystem it can be migrated using the 'Migrate and Remove' option in the Software Library Administration page.

See [Chapter 11, "Configuring a Software Library"](#) for more information.

2. Configure the storage replication schedule for as frequently as the network infrastructure as the network infrastructure will allow. Oracle recommends continuous replication occur every 2 hours (minimum).

28.4.4 Migrating an Existing Site to Shared Storage

- - Use file system backups to move existing OMS and agent installations to shared storage.
- - Use the following guidelines to migrate from local file system to shared storage
 - All backups must be offline backups, i.e. OMS and agent processes on a host must be shut down completed before backing up and restoring.
 - The backups must be performed as root user and permissions must be preserved.
 - The directory paths for Middleware Home and Instance Home must not change. If required use symbolic links to keep the paths the same.
 - The directory paths for Middleware Home and Instance Home must not change. If required use symbolic links to keep the paths the same.
 - The migration can be done in a rolling fashion to avoid complete downtime of Cloud Control.
- Use "Migrate Swlib Storage Location" [Link to Software library chapter] to move the software library to shared storage.

28.5 Performing Switchover and Failover Operations

Activating the standby site can take place either by using a switchover or a failover. These are used in different situations as described below:

- **Switchover** - A pre-planned role reversal of the primary and standby sites. In a switchover, functionality is transferred from the primary site to a standby site in an orderly, coordinated operation. As such, both sites must be available for a switchover to complete. Switchover is usually performed for testing and validation of Disaster Recovery (DR) scenarios and for planned maintenance activities on the primary infrastructure. A switchover is the preferred method of activating the standby site as the primary.
- **Failover** - Activation of the standby site as the primary site when the original primary site becomes unavailable.

28.5.1 Switchover Procedure

This section describes the steps to switchover to the standby site. The same procedure is applied to switchover in either direction.

1. Shut down all OMS components at the primary site.
2. Shut down all virtual Management Agents at the primary site.
3. Unmount the OMS filesystem and the software library filesystems from OMS hosts at the primary site.
4. Perform on-demand replication of OMS and software library filesystems.

Note: Refer to your storage documentation for steps required to perform an on-demand replication.

5. Update DNS entry for the application virtual hostname.
6. Switchover Oracle Database using Data Guard switchover.

Use DGMRGL 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>;
```

7. Perform role reversal of the Software Library and OMS storage (refer to your storage documentation for instructions).
8. Re-enable replication schedules for SWLIB and OMS storage
9. Mount OMS and Software Library filesystems on OMS hosts at Standby site
10. Start the first OMS Admin Server at the standby site
11. Point OMS to new Primary Repository Database using the following command:

```
emctl config oms -store_repos_details -repos_conndesc <connect descriptor>
-repos_user <username>
```

Example

```
emctl config oms -store repos_details -repos_conndesc '(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=newscan.domain) (PORT=1521))) (CONNECT_
DATA=(SERVICE_NAME=enreps.domain)))' -repos_user SYSMAN
```

12. Start the OMS at the standby site.
13. Start the Management Agents at the standby site using the following command:

```
emctl start agent
```

14. Relocate Management Services and Repository target using the following command:

```
emctl config emrep -agent <agent name> -conn_desc <repository connection>
```

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.

15. Update the WebLogic Admin Server URL. This is done by navigating to the target homepage for GC Domain and selecting **WebLogic Domain-->Target Setup-->Monitoring Configuration from within Cloud Control**.

28.5.2 Failover Procedure

This section describes the steps to failover to the standby site, recover the Enterprise Manager application state by resynchronizing the Management Repository database with all Management Agents, and finally enabling the original primary database

1. Shut down all OMS components at the primary site if running.
2. Shut down all virtual agents at primary site if running.
3. Unmount OMS and Software Library filesystems from OMS hosts at primary site.
4. Perform on-demand replication of the OMS and Software Library filesystems.

Note: Refer to your storage documentation for steps required to perform an on-demand replication.

5. Update the DNS entry for the application virtual hostname.
6. Failover Oracle Database using Data Guard failover.
7. Perform role reversal of Software Library and OMS storage.
8. Re-enable replication schedules for SWLIB and OMS storage
9. Mount the OMS and Software Library filesystems on OMS hosts at the standby site
10. Start first OMS admin server
11. Point the OMS to new Primary Repository Database using the

```
emctl config oms -store_repos_details -repos_conndesc <connect descriptor>
```

```
-repos_user <username>
```

Example

```
emctl config oms -store repos_details -repos_conn_desc '(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=newscan.domain) (PORT=1521))) (CONNECT_
DATA=(SERVICE_NAME=emreps.domain)))' -repos_user SYSMAN
```

12. Start the OMSs.

13. Resync the Agents. You can perform this operation either through the Cloud Control console (Management Agent homepage) or using EM CLI. If the Agent is not blocked, then the resync should be performed using the following command:

```
emcli resync agent <agenthost:port>
```

14. Start the Agents.

15. Relocate Management Services and Repository target using the following command:

```
emctl config emrep -agent <agent name> -conn_desc
```

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.

16. Update WebLogic Admin Server URL. This is done by navigating to the target homepage for GC Domain and selecting **WebLogic Domain-->Target Setup-->Monitoring Configuration from within Cloud Control**.

28.6 Keeping the Standby Site in Sync with the Primary

The standby site will be kept in sync with the primary automatically through the combination of Data Guard and storage replication.

The administrator should ensure that an on-demand replication to the standby site takes place before and after the following operations on the OMS or the agent:

- Plug-in deployment/undeployment, or existing plug-in upgrade
- Upgrade
- Patch
- emctl commands (other than lifecycle verbs (start/stop/status oms))
- Configuration of ADP/JVMD/BI Publisher

Note: Refer to your storage documentation for steps required to perform an on-demand replication.

Backing Up and Recovering 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](#)
- [Software Library Backup](#)
- [Management Repository Backup](#)
- [Oracle Management Service Backup](#)
- [Management Agent Backup](#)
- [Recovery of Failed Enterprise Manager Components](#)
- [Recovering from a Simultaneous OMS-Management Repository Failure](#)

29.1 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
- Software Library

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.

29.2 Software Library Backup

The software library is a centralized media storage for Enterprise Manager software entities such as software patches, virtual appliance images, reference gold images, application software, and their associated directive scripts. The software library is an essential part of Enterprise Manager framework and is required by many Enterprise

Manager features in order to function properly. The software library storage locations should be backed up periodically using file system backup. Oracle recommends the backup be performed at a frequency of 1 to 24 hours.

29.3 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 Repository 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.

29.4 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 changes 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).

29.5 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.

29.6 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
- Software Library

29.6.1 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:

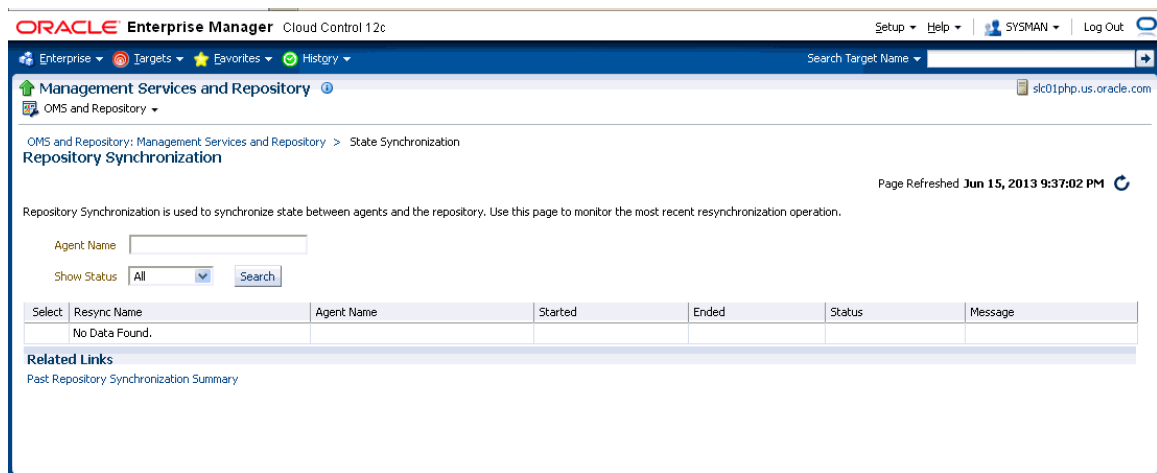
- **Full Recovery:** No special consideration is required for Enterprise Manager.
- **Point-in-Time/Incomplete Recovery:** Recovered repository may be out of sync with Agents because of lost transactions. In this situation, some metrics may show up incorrectly in the Cloud Control console unless the 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 resynchronization from the Enterprise Manager console's Repository Resynchronization page, as shown in the following figure.

Figure 29–1 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

29.6.2 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.

29.6.2.1 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 -all`.
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.

29.6.2.2 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 -all`.
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 in to Cloud Control. From the Setup menu, select Manage Cloud Control, and then Health Overview. The Management Services and Repository page displays.
6. From the **OMS and Repository** menu, select **Repository Synchronization**.
7. Verify that the site is fully operational.

29.6.2.3 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. Stop the OMS using the following command:

```
emctl stop oms -all
```

5. Start the OMS instances using the command

```
emctl start oms.
```

6. 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.
-

7. 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>"
```

8. Verify that the site is fully operational.

29.6.2.4 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
```

This commands will prompt you to stop and start the oms.

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. From the Setup menu, select Manage Cloud Control, and then select Health Overview.
9. From the OMS and Repository menu, select **Repository Synchronization**. Monitor the status of resync jobs. Resubmit failed jobs, if any, after fixing the error mentioned.
10. Verify that the site is fully operational.

29.6.3 Recovering the OMS

If an Oracle Management Service instance is lost, recovering it essentially consists of three steps: Recovering the Software Homes, configuring the Instance Home and recovering the Software Library if configured on same host as Enterprise Manager.

29.6.3.1 Recovering the Software Homes

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.

1. Connect to the Management Repository as SYSMAN and run the following SQL query to retrieve a list of installed plug-ins:

```
SELECT epv.display_name, epv.plugin_id, epv.version, epv.rev_version,
decode(su.aru_file, null, 'Media/External',
'https://updates.oracle.com/Orion/Services/download/'||aru_file||'?aru='||aru_
id||chr(38)||'patch_file='||aru_file) URL
FROM em_plugin_version epv, em_current_deployed_plugin ecp, em_su_entities su
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 su.entity_id = epv.su_entity_id;
```

The above query returns the list of plug-ins along with the URLs to download them if they were downloaded through self update. If plug-ins are present in the install media or are third party plug-ins not available through Self Update, the URLs are marked as "Media/Unknown".

2. Download the additional plug-ins, if any, from the URLs in the list returned by the query in step 1 and place them in a single directory. Change the filename extension from *.zip* to *.opar*.
3. Invoke the installer and select the Software-Only option to install the Middleware and OMS Oracle Home.
4. To install the required plug-ins, you must then run the PluginInstall.sh script (OMS_HOME/sysman/install/PluginInstall.sh) with the -pluginLocation <absolute path to plugin dir> specifying the path to the directory where downloaded plugins are kept. When asked to select plugins, make sure you select the same plugins as were listed in the SQL query.

Note: Recovery will fail if all required plug-ins have not been installed.

After the 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 SQL 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.

29.6.3.2 Recreating the OMS

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.

29.6.4 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.

29.6.4.1 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 and valid. If a Software library is accessible but corrupt, it will affect OMSCA recovery.
3. Restore the Software Homes. See [Section 29.6.3.1, "Recovering the Software Homes"](#) for more information.

If restoring from a filesystem backup, delete the following file:

`OMS_HOME/sysman/config/emInstanceMapping.properties`

In addition, delete any `gc_inst` directories that may have been restored, if they exist.

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 may block the Management Agent. Synchronize the agent with repository using the following command:

```
<OMS_HOME>/bin/emcli resyncAgent -agent=<agent target name myhost.example.com:3872>
```


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.

29.6.4.2 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". See [Section 29.6.3.1, "Recovering the Software Homes"](#) for more information.
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>
```

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>
```

Note: If you run `emctl config emrep -agent` and set the flag `-ignore_timeskew`, there may be a loss of monitoring data as the availability of monitored targets may be affected when the Management Services and Repository target is moved to the new Agent.

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.

29.6.4.3 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 the software library location is accessible from Host "B".
2. Restore the software homes on Host B. See [Section 29.6.3.1, "Recovering the Software Homes"](#) for more information.
3. 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>
```

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>
-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.

5. Start the OMS.

```
<OMS_HOME>/bin/emctl start oms
```

6. 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>
```

The OMS may block the Management Agent. Synchronize the Agent with repository using the following command:

```
<OMS_HOME>/bin/emcli resyncAgent -agent=<agent target name
myhost.example.com:3872>
```

7. Verify that the site is fully operational.

29.6.4.4 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. See [Section 29.6.3.1, "Recovering the Software Homes"](#) for more information.

If restoring from a filesystem backup, delete the following file:

<OMS_HOME>/sysman/config/emInstanceMapping.properties

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 may block the Management Agent. Synchronize the Agent with the repository using the following command:

```
<OMS_HOME>/bin/emcli resyncAgent -agent=<agent target name
e.g. myhost.example.com:3872>
```

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. Verify that the site is fully operational.

29.6.4.5 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". See [Section 29.6.3.1, "Recovering the Software Homes"](#) for more information.
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>
```

If any non-default plug-ins were previously deployed on the failed agent, they must be re-deployed after recovery of the Agent. Note that this pertains to plug-ins that existed on the recovering Agent before it failed (that are not related to the OMS/Repository target), and any plug-ins for additional targets the OMS Agent happened to be also monitoring. To re-deploy the plug-ins, run the following command (not as part of `config emrep`, or manually):

```
emcli relocate_targets
```

7. Additional Management Services, if any, must be re-enrolled with the Admin Server that is now running on host B. To re-enroll the Management Services, run the following command on each additional OMS:

```
<OMS-HOME>/bin/emctl enroll oms -as_host <new Admin Server host, i.e. host B> -as_port <admin server port>
```

8. Add the new OMS to the SLB virtual server pools and remove the old OMS.
9. 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>
```

Note: If you run `emctl config emrep -agent` and set the flag `-ignore_timeskew`, there may be a loss of monitoring data as the availability of monitored targets may be affected when the Management Services and Repository target is moved to the new Agent.

10. 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.
11. 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".
12. 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.
13. All other OMSs in the system must re-enroll with the newly recovered OMS using the following command:

```
emctl enroll oms -as_host <new OMS host> -as_port <port #, default 7101>
```

14. Verify that the site is fully operational.

29.6.4.6 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. See the Oracle® Enterprise Manager Cloud Control Basic Installation Guide for more information.
5. Verify that the site is fully operational.

29.6.5 Recovering the Software Library

If the software library is lost, it should be restored from the last available backup. After restoring the backup, the following commands must be run to verify and re-import missing entities:

1. `emcli verify_swlib` - This command verifies the accessibility of the software library storage locations and reports if entities are missing any files on the file system.
2. `emcli reimport_swlib_metadata` - This command re-imports all Oracle-supplied entities that are shipped along with the product. If you have a recent backup, this should not be required. Run `emcli reimport_swlib_metadata` if the `emcli verify_swlib` command reports Oracle-owned entities with files missing from the filesystem.
3. `emcli verify_updates` - This command verifies whether entities downloaded by Self Update are missing from the software library. For each missing entity, the command also displays the instructions to re-import the entity into the software library.

29.6.6 Recovering Management Agents

If a 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 a Management Agent install because 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: 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 using the `emcli resyncAgent <agent target name>` command. Resynchronization pushes all targets from the Management Repository to the Management Agent and then unblocks the Agent.

29.6.7 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.

29.6.7.1 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 through 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 using the following command:

```
emcli resyncAgent -agent="Agent Host:Port"
```

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.

29.6.7.2 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 using the following command:

```
emcli resyncAgent -agent="Agent Host:Port"
```

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.

29.7 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.

29.7.1 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 29.6.4.1, "Single OMS, No Server Load Balancer \(SLB\), OMS Restored on the same Host."](#)
3. Verify that the site is fully operational.

29.7.2 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 29.6.2.2, "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 29.6.4.5, "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 29.6.4.5, "Multiple OMS, Server Load Balancer Configured, Primary OMS Recovered on a Different Host."](#)

3. Follow the OMS recovery procedure shown in [Section 29.6.4.6, "Multiple OMS, SLB configured, additional OMS recovered on same or different host."](#)
4. Verify that the site is fully operational.

Interpreting Variables of the Enterprise Manager 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. This appendix provides information to help you interpret the variables of the private Oracle Enterprise Manager MIB.

For information about the format of MIB variable descriptions, see [Section 4.6.3, "Reading the MIB Variable Descriptions"](#).

This appendix contains the following sections:

- [oraEMNGEvent](#)
- [oraEM4AlertTable](#)
- [oraEM4JobAlertTable](#)

Note: SNMP trap notification methods created from the Cloud Control 12c console send the new oraEMNGEvent trap.

SNMP trap notification methods created before Enterprise Manager Release 12c send oraEM4Alert and oraEM4JobAlert traps. After upgrading to Enterprise Manager Release 12c, existing methods from earlier releases continue to deliver the old traps for backward compatibility.

oraEMNGEvent

The following sections describe the SNMP traps sent from SNMP trap notification methods created from the Cloud Control 12c console.

Table A-1 *oraEMNGEvent Variables and Corresponding Object IDs*

Variable Name	Object ID
oraEMNGEventIndex	1.3.6.1.4.1.111.15.3.1.1.1.1
oraEMNGEventNotifType	1.3.6.1.4.1.111.15.3.1.1.2.1
oraEMNGEventMessage	1.3.6.1.4.1.111.15.3.1.1.3.1
oraEMNGEventMessageURL	1.3.6.1.4.1.111.15.3.1.1.4.1
oraEMNGEventSeverity	1.3.6.1.4.1.111.15.3.1.1.5.1
oraEMNGEventSeverityCode	1.3.6.1.4.1.111.15.3.1.1.6.1

Table A-1 (Cont.) oraEMNGEvent Variables and Corresponding Object IDs

Variable Name	Object ID
oraEMNGEventRepeatCount	1.3.6.1.4.1.111.15.3.1.1.7.1
oraEMNGEventActionMsg	1.3.6.1.4.1.111.15.3.1.1.8.1
oraEMNGEventOccurrenceTime	1.3.6.1.4.1.111.15.3.1.1.9.1
oraEMNGEventReportedTime	1.3.6.1.4.1.111.15.3.1.1.10.1
oraEMNGEventCategories	1.3.6.1.4.1.111.15.3.1.1.11.1
oraEMNGEventCategoryCodes	1.3.6.1.4.1.111.15.3.1.1.12.1
oraEMNGEventType	1.3.6.1.4.1.111.15.3.1.1.13.1
oraEMNGEventName	1.3.6.1.4.1.111.15.3.1.1.14.1
oraEMNGAssocIncidentId	1.3.6.1.4.1.111.15.3.1.1.15.1
oraEMNGAssocIncidentOwner	1.3.6.1.4.1.111.15.3.1.1.16.1
oraEMNGAssocIncidentAked	1.3.6.1.4.1.111.15.3.1.1.17.1
oraEMNGAssocIncidentStatus	1.3.6.1.4.1.111.15.3.1.1.18.1
oraEMNGAssocIncidentPriority	1.3.6.1.4.1.111.15.3.1.1.19.1
oraEMNGAssocIncidentEscLevel	1.3.6.1.4.1.111.15.3.1.1.20.1
oraEMNGEventTargetName	1.3.6.1.4.1.111.15.3.1.1.21.1
oraEMNGEventTargetNameURL	1.3.6.1.4.1.111.15.3.1.1.22.1
oraEMNGEventTargetType	1.3.6.1.4.1.111.15.3.1.1.23.1
oraEMNGEventHostName	1.3.6.1.4.1.111.15.3.1.1.24.1
oraEMNGEventTargetOwner	1.3.6.1.4.1.111.15.3.1.1.25.1
oraEMNGEventTgtLifeCycleStatus	1.3.6.1.4.1.111.15.3.1.1.26.1
oraEMNGEventTargetVersion	1.3.6.1.4.1.111.15.3.1.1.27.1
oraEMNGEventUserDefinedTgtProp	1.3.6.1.4.1.111.15.3.1.1.28.1
oraEMNGEventSourceObjName	1.3.6.1.4.1.111.15.3.1.1.29.1
oraEMNGEventSourceObjNameURL	1.3.6.1.4.1.111.15.3.1.1.30.1
oraEMNGEventSourceObjType	1.3.6.1.4.1.111.15.3.1.1.31.1
oraEMNGEventSourceObjSubType	1.3.6.1.4.1.111.15.3.1.1.32.1
oraEMNGEventSourceObjOwner	1.3.6.1.4.1.111.15.3.1.1.33.1
oraEMNGEventCAJobName	1.3.6.1.4.1.111.15.3.1.1.34.1
oraEMNGEventCAJobStatus	1.3.6.1.4.1.111.15.3.1.1.35.1
oraEMNGEventCAJobOwner	1.3.6.1.4.1.111.15.3.1.1.36.1
oraEMNGEventCAJobStepOutput	1.3.6.1.4.1.111.15.3.1.1.37.1
oraEMNGEventCAJobType	1.3.6.1.4.1.111.15.3.1.1.38.1
oraEMNGEventRuleSetName	1.3.6.1.4.1.111.15.3.1.1.39.1
oraEMNGEventRuleName	1.3.6.1.4.1.111.15.3.1.1.40.1
oraEMNGEventRuleOwner	1.3.6.1.4.1.111.15.3.1.1.41.1
oraEMNGEventSequenceId	1.3.6.1.4.1.111.15.3.1.1.42.1
oraEMNGEventRCADetails	1.3.6.1.4.1.111.15.3.1.1.43.1

Table A–1 (Cont.) oraEMNGEvent Variables and Corresponding Object IDs

Variable Name	Object ID
oraEMNGEventContextAttrs	1.3.6.1.4.1.111.15.3.1.1.44.1
oraEMNGEventUserComments	1.3.6.1.4.1.111.15.3.1.1.45.1
oraEMNGEventUpdates	1.3.6.1.4.1.111.15.3.1.1.46.1
oraEMNGEventTypeAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1
oraEMNGEventTypeAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1
oraEMNGEventTypeAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1
oraEMNGEventTypeAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1
oraEMNGEventTypeAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1
oraEMNGEventTypeAttr6	1.3.6.1.4.1.111.15.3.1.1.66.1
oraEMNGEventTypeAttr7	1.3.6.1.4.1.111.15.3.1.1.67.1
oraEMNGEventTypeAttr8	1.3.6.1.4.1.111.15.3.1.1.68.1
oraEMNGEventTypeAttr9	1.3.6.1.4.1.111.15.3.1.1.69.1
oraEMNGEventTypeAttr10	1.3.6.1.4.1.111.15.3.1.1.70.1
oraEMNGEventTypeAttr11	1.3.6.1.4.1.111.15.3.1.1.71.1
oraEMNGEventTypeAttr12	1.3.6.1.4.1.111.15.3.1.1.72.1
oraEMNGEventTypeAttr13	1.3.6.1.4.1.111.15.3.1.1.73.1
oraEMNGEventTypeAttr14	1.3.6.1.4.1.111.15.3.1.1.74.1
oraEMNGEventTypeAttr15	1.3.6.1.4.1.111.15.3.1.1.75.1
oraEMNGEventTypeAttr16	1.3.6.1.4.1.111.15.3.1.1.76.1
oraEMNGEventTypeAttr17	1.3.6.1.4.1.111.15.3.1.1.77.1
oraEMNGEventTypeAttr18	1.3.6.1.4.1.111.15.3.1.1.78.1
oraEMNGEventTypeAttr19	1.3.6.1.4.1.111.15.3.1.1.79.1
oraEMNGEventTypeAttr20	1.3.6.1.4.1.111.15.3.1.1.80.1
oraEMNGEventTypeAttr21	1.3.6.1.4.1.111.15.3.1.1.81.1
oraEMNGEventTypeAttr22	1.3.6.1.4.1.111.15.3.1.1.82.1
oraEMNGEventTypeAttr23	1.3.6.1.4.1.111.15.3.1.1.83.1
oraEMNGEventTypeAttr24	1.3.6.1.4.1.111.15.3.1.1.84.1
oraEMNGEventTypeAttr25	1.3.6.1.4.1.111.15.3.1.1.85.1

oraEMNGEventIndex

Syntax

Integer

Access

Read-only

Status

Mandatory

Description

The index of a particular event, unique only at the moment an event is generated.

oraEMNGEventNotifType**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The notification type. Possible values are:

- NOTIF_NORMAL
- NOTIF_RETRY
- NOTIF_DURATION
- NOTIF_REPEAT
- NOTIF_CA
- NOTIF_RCA

oraEMNGEventMessage**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The message associated with this event.

oraEMNGEventMessageURL

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The Enterprise Manager Cloud Control console URL for the event message. It is populated for events with severities other than INFORMATIONAL. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventSeverity**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The severity of the event, such as Fatal, Critical, Warning, Advisory, Information, or Clear.

oraEMNGEventSeverityCode**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The internal code of the severity, such as Fatal, Critical, Warning, Advisory, Informational, or Clear.

oraEMNGEventRepeatCount

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The repeat notification counter for the event.

oraEMNGEventActionMsg**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The action message for this event.

oraEMNGEventOccurrenceTime**Syntax**

DisplayString

Access

read-only

Status

Mandatory

Explanation

The time when this event occurred (optional). This is only populated for events that have an occurrence time.

oraEMNGEventReportedTime**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The time when this event was reported.

oraEMNGEventCategories**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The list of categories to which this event belongs. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventCategoryCodes**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The list of internal category codes to which this event belongs. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventType**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the event type to which this event belongs.

Available Event Types

- metric_alert
- target_availability
- job_status_change
- metric_error
- user_reported
- cs_core
- sla_alert
- next_update
- selfupdate
- cs_rule_violation

oraEMNGEventName**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of this event.

oraEMNGAssocIncidentId**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The ID of the associated incident with the event (optional).

oraEMNGAssocIncidentOwner**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Owner of the associated incident with the event (optional).

oraEMNGAssocIncidentAcked**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Acknowledged status of the associated incident with the event. 1 indicates acknowledged. 0 indicates unacknowledged.

oraEMNGAssocIncidentStatus**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The status of the associated incident with the event.

oraEMNGAssocIncidentPriority**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The priority of the associated incident with the event.

oraEMNGAssocIncidentEscLevel**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The escalation level of the associated incident with the event.

oraEMNGEventTargetName**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the target to which this event applies. Populated for events that are about a target only.

oraEMNGEventTargetNameURL

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The Enterprise Manager Console URL of the target to which this event applies. Populated for events that are about a target only. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventTargetType

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The type of the target to which this event applies. Populated for events that are about a target only.

oraEMNGEventHostName

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the host on which this event originated. Populated for events that are about a target only.

oraEMNGEventTargetOwner

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The primary administrator of the target on which this event originated. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventTgtLifeCycleStatus

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The life cycle status of the target on which this event originated.

oraEMNGEventTargetVersion

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The version of the target on which this event originated.

oraEMNGEventUserDefinedTgtProp

Syntax

DisplayString

Access

read-only

Status

Mandatory

Explanation

The user defined target properties [name,value pair list] of the associated target with this event. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventSourceObjName**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the source object to which this event belongs to. Populated for events that are about a non-target object only, such as Jobs.

oraEMNGEventSourceObjNameURL**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Enterprise Manager Console URL for the source object to which this event belongs. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventSourceObjType

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The type of the source object to which this event belongs.

oraEMNGEventSourceObjSubType**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The subtype of the source object to which this event belongs. (Optional). This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventSourceObjOwner**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The primary administrator of the source object to which this event belongs. (Optional). This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventCAJobName

Syntax

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the corrective action job associated with this event.

oraEMNGEventCAJobStatus**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The status of the corrective action job associated with this event.

oraEMNGEventCAJobOwner**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The owner of the corrective action job associated with this event.

oraEMNGEventCAJobStepOutput**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The job step output from the corrective action job associated with this event.

oraEMNGEventCAJobType**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The job type from the corrective action job associated with this event.

oraEMNGEventRuleSetName**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the ruleset that caused this notification. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventRuleName**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The name of the rule within the ruleset that caused this notification.

oraEMNGEventRuleOwner**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The owner of the ruleset that caused this notification.

oraEMNGEventSequenceld**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

An Enterprise Manager-generated identifier that uniquely identifies the current issue until it is cleared.

oraEMNGEventRCADetails**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

Root Cause Analysis information associated with this event if it exists.

oraEMNGEventContextAttrs**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The context attributes associated with this event. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventUserComments**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The user comments associated with this event. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventUpdates**Syntax**

DisplayString

Access

Read-only

Status

Mandatory

Explanation

The updates associated with this event. This variable is empty if the trap size exceeds the configured SNMP packet size.

oraEMNGEventAttr(1-71)

The following tables list oraEMNGEvent MIB variables 1 through 71. Each table categorizes the MIB variables by specific event type.

Table A-2 Metric Alert Event Type

Variable Name	OID Number	Event Type	
		Attribute	Description
oraEMNGEventAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Metric GUID	A unique ID for the metric.
oraEMNGEventAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Severity GUID	A unique ID for the alert record.
oraEMNGEventAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Cycle GUID	A unique ID for the alert cycle.
oraEMNGEventAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Collection Name	The name of the collection collecting the metric.
oraEMNGEventAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1	Metric Group	The name of the metric.
oraEMNGEventAttr6	1.3.6.1.4.1.111.15.3.1.1.66.1	Metric	The name of the metric column.
oraEMNGEventAttr7	1.3.6.1.4.1.111.15.3.1.1.67.1	Metric Description	A brief description of the metric.
oraEMNGEventAttr8	1.3.6.1.4.1.111.15.3.1.1.68.1	Metric value	The value of the metric when the event triggered.
oraEMNGEventAttr9	1.3.6.1.4.1.111.15.3.1.1.69.1	Key Value	The monitored object for the metric corresponding to the Metric Alert event.
oraEMNGEventAttr10	1.3.6.1.4.1.111.15.3.1.1.70.1	Key Column 1	Key Column 1
oraEMNGEventAttr11	1.3.6.1.4.1.111.15.3.1.1.71.1	Key Column 1 Value	The value of Key Column 1.
oraEMNGEventAttr12	1.3.6.1.4.1.111.15.3.1.1.72.1	Key Column 2	Key Column 2
oraEMNGEventAttr13	1.3.6.1.4.1.111.15.3.1.1.73.1	Key Column 2 Value	The value of Key Column 2.
oraEMNGEventAttr14	1.3.6.1.4.1.111.15.3.1.1.74.1	Key Column 3	Key Column 3
oraEMNGEventAttr15	1.3.6.1.4.1.111.15.3.1.1.75.1	Key Column 3 Value	The value of Key Column 3.
oraEMNGEventAttr16	1.3.6.1.4.1.111.15.3.1.1.76.1	Key Column 4	Key Column 4
oraEMNGEventAttr17	1.3.6.1.4.1.111.15.3.1.1.77.1	Key Column 4 Value	The value of Key Column 4.
oraEMNGEventAttr18	1.3.6.1.4.1.111.15.3.1.1.78.1	Key Column 5	Key Column 5
oraEMNGEventAttr19	1.3.6.1.4.1.111.15.3.1.1.79.1	Key Column 5 Value	The value of Key Column 5.
oraEMNGEventAttr20	1.3.6.1.4.1.111.15.3.1.1.80.1	Key Column 6	Key Column 6
oraEMNGEventAttr21	1.3.6.1.4.1.111.15.3.1.1.81.1	Key Column 6 Value	The value of Key Column 6.
oraEMNGEventAttr22	1.3.6.1.4.1.111.15.3.1.1.82.1	Key Column 7	Key Column 7
oraEMNGEventAttr23	1.3.6.1.4.1.111.15.3.1.1.83.1	Key Column 7 Value	The value of Key Column 7.
oraEMNGEventAttr24	1.3.6.1.4.1.111.15.3.1.1.84.1	Number of keys	The number of key metric columns in the metric.

Table A–3 Target Availability Event Type

Variable Name	OID Number	Event Type	
		Attribute	Description
oraEMNGEventAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Availability status	The current availability status of the target.
oraEMNGEventAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Severity GUID	The GUID of the severity record associated with this availability status.
oraEMNGEventAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Availability Sub-status	The sub-status of a target for the current status.
oraEMNGEventAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Transition Severity	The severity that resulted in the target's status change to the current availability status.
oraEMNGEventAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1	Response metric GUID	The Metric GUID of response metric.
oraEMNGEventAttr6	1.3.6.1.4.1.111.15.3.1.1.66.1	Severity GUID of the first severity in the availability cycle	The GUID of the first severity record in this availability cycle.

Table A–4 Job Status Change Event Type

Variable Name	OID Number	Event Type	
		Attribute	Description
oraEMNGEventAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Execution ID	The unique ID of the job execution.
oraEMNGEventAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Job Status	The status of the job execution.
oraEMNGEventAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Execution Log	The job output of the last step executed.
oraEMNGEventAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Job Status Code	The execution status code of job execution.
oraEMNGEventAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1	State Change ID	The unique ID of the last status change.

Table A–5 Compliance Standard Rule Violation Event Type

Variable Name	OID Number	Event Type	
		Attribute	Description
oraEMNGEventAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Root Compliance Standard	The root compliance standard node display name.
oraEMNGEventAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Root Compliance Standard Version	The root compliance standard version.
oraEMNGEventAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Root Compliance Standard Author	The author of the root compliance standard.
oraEMNGEventAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Parent Compliance Standard	The parent compliance standard node display name.

Table A–5 (Cont.) Compliance Standard Rule Violation Event Type

Variable Name	OID Number	Event Type	
		Attribute	Description
oraEMNGEventAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1	Compliance Standard Version	The compliance standard version.
oraEMNGEventAttr6	1.3.6.1.4.1.111.15.3.1.1.66.1	Parent Compliance Standard Author	The author of a parent compliance standard.
oraEMNGEventAttr7	1.3.6.1.4.1.111.15.3.1.1.67.1	Root Target Name	The root target name.
oraEMNGEventAttr8	1.3.6.1.4.1.111.15.3.1.1.68.1	Root Target Type	The root target type.
oraEMNGEventAttr9	1.3.6.1.4.1.111.15.3.1.1.69.1	Rule Name	The rule display name

Table A–6 Compliance Standard Score Event Type

Variable Name	OID Number	Event Type	
		Attribute	Description
oraEMNGEventAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Root Compliance Standard	The root compliance standard node display name.
oraEMNGEventAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Root Compliance Standard Author	The author of the root compliance standard.
oraEMNGEventAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Root Compliance Standard Version	The version of the root compliance standard.
oraEMNGEventAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Compliance Standard	The compliance standard node display name.
oraEMNGEventAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1	Compliance Standard Version	The version of a compliance standard.
oraEMNGEventAttr6	1.3.6.1.4.1.111.15.3.1.1.66.1	Compliance Standard Author	The author of a compliance standard.
oraEMNGEventAttr7	1.3.6.1.4.1.111.15.3.1.1.67.1	Root Target Name	The root target name.
oraEMNGEventAttr8	1.3.6.1.4.1.111.15.3.1.1.68.1	Root Target Type	The root target type.
oraEMNGEventAttr10	1.3.6.1.4.1.111.15.3.1.1.70.1	Warning Threshold	The warning threshold of a compliance score.
oraEMNGEventAttr9	1.3.6.1.4.1.111.15.3.1.1.69.1	Compliance Score	The compliance score.
oraEMNGEventAttr11	1.3.6.1.4.1.111.15.3.1.1.71.1	Critical Threshold	The critical threshold of a compliance score.

Table A–7 Metric Error Event Type

Variable Name	OID Number	Event Type	
		Attribute	Description
oraEMNGEventAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Metric Group	The name of the metric.
oraEMNGEventAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Collection Name	The name of the collection collecting the metric.

Table A–8 Metric Extension Event Type

Variable Name	OID Number	Event Type	
		Attribute	Description
oraEMNGEventAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Metric Extension Version attribute	The version of the metric extension.

Table A–9 Self-update Event Type

Variable Name	OID Number	Event Type	
		Attribute	Description
oraEMNGEventAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Type	Type
oraEMNGEventAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Description	Description
oraEMNGEventAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Version	Version
oraEMNGEventAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Status	Status

Table A–10 Service Level Agreement Alert Event Type

Variable Name	OID Number	Event Type	
		Attribute	Description
oraEMNGEventAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Service Level Agreement Name	Service Level Agreement Name
oraEMNGEventAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	Service Level Objective Name	Service Level Objective Name
oraEMNGEventAttr3	1.3.6.1.4.1.111.15.3.1.1.63.1	Service Level Objective Type	The type of the Service Level Objective which will be either Performance or Availability
oraEMNGEventAttr4	1.3.6.1.4.1.111.15.3.1.1.64.1	Value at Event Triggered	The value at Event Triggered
oraEMNGEventAttr5	1.3.6.1.4.1.111.15.3.1.1.65.1	Customer Name	Customer Name

Table A–11 User-reported Event Type

Variable Name	OID Number	Event Type	
		Attribute	Description
oraEMNGEventAttr1	1.3.6.1.4.1.111.15.3.1.1.61.1	Name	The name describing the nature of the issue.
oraEMNGEventAttr2	1.3.6.1.4.1.111.15.3.1.1.62.1	key	The optional key describing a sub-component within the target that this event is about.

oraEM4AlertTable

The oraEM4AlertTable describes the SNMP traps sent from Oracle Enterprise Manager for both metric severity alerts and policy violations.

[Table A–12](#) lists the variables of the oraEM4AlertTable and their corresponding Object IDs.

Table A–12 oraEM4AlertTable Variables and Corresponding Object IDs

Variable Name	Object ID
oraEM4AlertTargetName	1.3.6.1.4.1.111.15.1.1.1.2.1
oraEM4AlertTargetType	1.3.6.1.4.1.111.15.1.1.1.3.1
oraEM4AlertHostName	1.3.6.1.4.1.111.15.1.1.1.4.1
oraEM4AlertMetricName	1.3.6.1.4.1.111.15.1.1.1.5.1
oraEM4AlertKeyName	1.3.6.1.4.1.111.15.1.1.1.6.1
oraEM4AlertKeyValue	1.3.6.1.4.1.111.15.1.1.1.7.1
oraEM4AlertTimeStamp	1.3.6.1.4.1.111.15.1.1.1.8.1
oraEM4AlertSeverity	1.3.6.1.4.1.111.15.1.1.1.9.1
oraEM4AlertMessage	1.3.6.1.4.1.111.15.1.1.1.10.1
oraEM4AlertRuleName	1.3.6.1.4.1.111.15.1.1.1.11.1
oraEM4AlertRuleOwner	1.3.6.1.4.1.111.15.1.1.1.12.1
oraEM4AlertMetricValue	1.3.6.1.4.1.111.15.1.1.1.13.1
oraEM4AlertContext	1.3.6.1.4.1.111.15.1.1.1.14.1
oraEM4AlertCycleGuid	1.3.6.1.4.1.111.15.1.1.1.15.1
oraEM4AlertRepeatCount	1.3.6.1.4.1.111.15.1.1.1.16.1
oraEM4AlertUDTargetProperties	1.3.6.1.4.1.111.15.1.1.1.17.1
oraEM4AlertAck	1.3.6.1.4.1.111.15.1.1.1.18.1
oraEM4AlertAckBy	1.3.6.1.4.1.111.15.1.1.1.19.1
oraEM4AlertNotifType	1.3.6.1.4.1.111.15.1.1.1.20.1
oraEM4AlertViolationGuid	1.3.6.1.4.1.111.15.1.1.1.21.1

A description of each of these variables follows.

oraEM4AlertTargetName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the target to which this alert applies.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertTargetType**Syntax**

DisplayString

Max-Access

read-only

Status

Mandatory

Explanation

The type of the target to which this alert applies.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertHostName**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the host on which this alert originated.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertMetricName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the metric or policy which generated this alert.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertKeyName**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the key-column, if present, for the metric which generated this alert.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertKeyValue**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The value of the key-column, if present, for the metric which generated this alert.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertTimeStamp**Syntax**

DisplayString

Max-Access

read-only

Status

Mandatory

Explanation

The time at which this alert was generated.

Typical Range

Not applicable

Significance

Important

oraEM4AlertSeverity**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The severity of the alert (for example, Clear, Informational, Warning, Critical, Unreachable Start, Blackout End, Blackout Start, Metric Error Clear, Metric Error Start, Status Pending).

Typical Range

Critical, warning, clear

Significance

Very important

oraEM4AlertMessage**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The message associated with the alert.

Typical Range

Not applicable

Significance

Very important

oraEM4AlertRuleName**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the notification rule that caused this notification.

Typical Range

Not applicable

Significance

Important

oraEM4AlertRuleOwner

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The owner of the notification rule that caused this notification.

Typical Range

Not applicable

Significance

Important

oraEM4AlertMetricValue

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The value of the metric which caused this alert to be generated.

Typical Range

Not applicable

Significance

Important

oraEM4AlertContext

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

A comma separated list of metric column names and values associated with the metric that caused this alert to be generated.

Typical Range

Not applicable

Significance

Important

oraEM4AlertCycleGuid**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

An Enterprise Manager-generated identifier that is unique for the lifecycle of an alert.

Typical Range

Not applicable

Significance

Important

oraEM4AlertRepeatCount**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The repeat notification counter for the alert.

Typical Range

Not applicable

Significance

Important

oraEM4AlertUDTargetProperties**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

User-defined target properties associated with the target.

Typical Range

Not applicable

Significance

Important

oraEM4AlertAck**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

Acknowledged status flag associated with the alert. 1 indicates acknowledged. 0 indicates unacknowledged.

Typical Range

Not applicable

Significance

Important

oraEM4AlertAckBy**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

Acknowledged By value associated with the alert.

Typical Range

Not applicable

Significance

Important

oraEM4AlertNotifType**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

Notification type.

Possible values:

- 1 - Normal
- 4 - Repeat
- 9 - Duration

Typical Range

Not applicable

Significance

Important

oraEM4AlertViolationGuid**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

An Enterprise Manager-generated identifier that identifies a particular alert.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertTable

The oraEM4JobAlertTable describes changes in the status of either a Job or a Corrective Action that is running as part of the Oracle Enterprise Manager Job system.

[Table A-13](#) lists the variables of the oraEM4JobAlertTable and their corresponding Object IDs.

Table A-13 oraEM4JobAlertTable Variables and Corresponding Object IDs

Variable Name	Object ID
oraEM4JobAlertJobName	1.3.6.1.4.1.111.15.1.2.1.2.1
oraEM4JobAlertJobOwner	1.3.6.1.4.1.111.15.1.2.1.3.1
oraEM4JobAlertJobType	1.3.6.1.4.1.111.15.1.2.1.4.1
oraEM4JobAlertJobStatus	1.3.6.1.4.1.111.15.1.2.1.5.1
oraEM4JobAlertTargets	1.3.6.1.4.1.111.15.1.2.1.6.1
oraEM4JobAlertTimeStamp	1.3.6.1.4.1.111.15.1.2.1.7.1
oraEM4JobAlertRuleName	1.3.6.1.4.1.111.15.1.2.1.8.1
oraEM4JobAlertRuleOwner	1.3.6.1.4.1.111.15.1.2.1.9.1
oraEM4JobAlertMetricName	1.3.6.1.4.1.111.15.1.2.1.10.1
oraEM4JobAlertMetricValue	1.3.6.1.4.1.111.15.1.2.1.11.1
oraEM4JobAlertContext	1.3.6.1.4.1.111.15.1.2.1.12.1
oraEM4JobAlertKeyName	1.3.6.1.4.1.111.15.1.2.1.13.1

Table A–13 (Cont.) oraEM4JobAlertTable Variables and Corresponding Object IDs

Variable Name	Object ID
oraEM4JobAlertKeyValue	1.3.6.1.4.1.111.15.1.2.1.14.1
oraEM4JobAlertSeverity	1.3.6.1.4.1.111.15.1.2.1.15.1
oraEM4JobAlertJobId	1.3.6.1.4.1.111.15.1.2.1.16.1
oraEM4JobAlertJobExecId	1.3.6.1.4.1.111.15.1.2.1.17.1

A description of each of these variables follows.

oraEM4JobAlertJobName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the job to which this alert applies.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertJobOwner

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The owner of the job to which this alert applies.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertJobType**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The type of the job to which this alert applies.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertJobStatus**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The status of the job to which this alert applies.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertTargets

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

A comma separated list of target to which this alert applies.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertTimeStamp**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The time at which this job status changed causing this alert.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertRuleName**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the notification rule that caused this notification.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertRuleOwner**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The owner of the notification rule that caused this notification.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertMetricName**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the metric or policy which caused the Corrective Action to run that caused this alert.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertMetricValue**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The value of the metric which caused the Corrective Action to run that caused this alert.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertContext**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

A comma separated list of metric column names and values associated with the metric which caused the Corrective Action to run that caused this alert.

Typical Range

Not applicable

Significance

Important

oraEM4JobAlertKeyName

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The name of the key-column, if present, for the metric which caused the Corrective Action to run that generated this alert.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertKeyValue

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The value of the key-column, if present, for the metric which caused the Corrective Action to run that generated this alert.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertSeverity

Syntax

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The severity of the metric which caused the Corrective Action to run that generated this alert (for example, Critical).

Typical Range

Critical, warning, clear

Significance

Very important

oraEM4JobAlertJobId**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The job ID of the Enterprise Manager job that triggered this notification.

Typical Range

Not applicable

Significance

Very important

oraEM4JobAlertJobExecId**Syntax**

DisplayString

Max-Access

Read-only

Status

Mandatory

Explanation

The job execution ID of the Enterprise Manager job that triggered this notification.

Typical Range

Not applicable

Significance

Very important

Enterprise Manager MIB Definition

The following MIB definition is the latest version at the time of publication. For the most recent version of the Enterprise Manager 12c MIB definition, view your installation MIB definition file at:

OMS_HOME/network/doc/omstrap.v1

B.1 MIB Definition

```

ORACLE-ENTERPRISE-MANAGER-4-MIB DEFINITIONS ::= BEGIN

IMPORTS
TRAP-TYPE
FROM RFC-1215
DisplayString
FROM RFC1213-MIB
OBJECT-TYPE
FROM RFC-1212
enterprises
FROM RFC1155-SMI;

oracle OBJECT IDENTIFIER ::= { enterprises 111 }

oraEM4 OBJECT IDENTIFIER ::= { oracle 15 }

oraEM4Objects OBJECT IDENTIFIER ::= { oraEM4 1 }

oraEM4AlertTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OraEM4AlertEntry
    ACCESS not-accessible
    STATUS mandatory
    DESCRIPTION
        "Information on alerts generated by Oracle Enterprise Manager. This table is
not queryable; it exists only to document the variables included in the
oraEM4Alert trap. Each trap contains a single instance of each variable in the
table."
    ::= { oraEM4Objects 1 }

oraEM4AlertEntry OBJECT-TYPE
    SYNTAX OraEM4AlertEntry
    ACCESS not-accessible
    STATUS mandatory
    DESCRIPTION
        "Information about a particular Oracle Enterprise Manager alert."

```

```
INDEX    { oraEM4AlertIndex }
 ::= { oraEM4AlertTable 1 }
OraEM4AlertEntry ::=
  SEQUENCE {
    oraEM4AlertIndex
      INTEGER,

    oraEM4AlertTargetName
      DisplayString,

    oraEM4AlertTargetType
      DisplayString,

    oraEM4AlertHostName
      DisplayString,

    oraEM4AlertMetricName
      DisplayString,

    oraEM4AlertKeyName
      DisplayString,

    oraEM4AlertKeyValue
      DisplayString,

    oraEM4AlertTimeStamp
      DisplayString,

    oraEM4AlertSeverity
      DisplayString,

    oraEM4AlertMessage
      DisplayString,

    oraEM4AlertRuleName
      DisplayString,

    oraEM4AlertRuleOwner
      DisplayString,

    oraEM4AlertMetricValue
      DisplayString,

    oraEM4AlertContext
      DisplayString,

    oraEM4AlertCycleGuid
      DisplayString,

    oraEM4AlertRepeatCount
      DisplayString,

    oraEM4AlertUDTargetProperties
      DisplayString,

    oraEM4AlertAck
      DisplayString,

    oraEM4AlertAckBy
      DisplayString,
```

```

        oraEM4AlertNotifType
            DisplayString,
        oraEM4AlertViolationGuid
            DisplayString
    }

oraEM4AlertIndex OBJECT-TYPE
    SYNTAX  INTEGER (0..2147483647)
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "Index of a particular alert, unique only at the moment an alert is
        generated."
    ::= { oraEM4AlertEntry 1 }

oraEM4AlertTargetName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The name of the target to which this alert applies."
    ::= { oraEM4AlertEntry 2 }

oraEM4AlertTargetType OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The type of the target to which this alert applies."
    ::= { oraEM4AlertEntry 3 }

oraEM4AlertHostName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The name of the host on which this alert originated."
    ::= { oraEM4AlertEntry 4 }

oraEM4AlertMetricName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The name of the metric or policy which generated this alert."
    ::= { oraEM4AlertEntry 5 }

oraEM4AlertKeyName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The name of the key-column, if present, for the metric which generated this
        alert."
    ::= { oraEM4AlertEntry 6 }

oraEM4AlertKeyValue OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only

```

```
STATUS mandatory
DESCRIPTION
    "The value of the key-column, if present, for the metric which generated this
    alert."
 ::= { oraEM4AlertEntry 7 }

oraEM4AlertTimeStamp OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The time at which this alert was generated."
    ::= { oraEM4AlertEntry 8 }

oraEM4AlertSeverity OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The severity of the alert e.g. Critical."
    ::= { oraEM4AlertEntry 9 }

oraEM4AlertMessage OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The message associated with the alert."
    ::= { oraEM4AlertEntry 10 }

oraEM4AlertRuleName OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The name of the notification rule that caused this notification."
    ::= { oraEM4AlertEntry 11 }

oraEM4AlertRuleOwner OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The owner of the notification rule that caused this notification."
    ::= { oraEM4AlertEntry 12 }

oraEM4AlertMetricValue OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The value of the metric which caused this alert to be generated."
    ::= { oraEM4AlertEntry 13 }

oraEM4AlertContext OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "A comma separated list of metric column names and values associated with the
```

```

metric that caused this alert to be generated."
 ::= { oraEM4AlertEntry 14 }

oraEM4AlertCycleGuid OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "An EM generated identifier that is unique for the lifecycle of an alert."
    ::= { oraEM4AlertEntry 15 }

oraEM4AlertRepeatCount OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The repeat notification counter for the alert."
    ::= { oraEM4AlertEntry 16 }

oraEM4AlertUDTargetProperties OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "User-defined target properties associated with the target."
    ::= { oraEM4AlertEntry 17 }

oraEM4AlertAck OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "Acknowledged status flag associated with the alert. 1 indicates
        acknowledged, 0 indicates unacknowledged."
    ::= { oraEM4AlertEntry 18 }

oraEM4AlertAckBy OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "Acknowledged By value associated with the alert."
    ::= { oraEM4AlertEntry 19 }

oraEM4AlertNotifType OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "Notification Type. 1 - Normal, 4 - Repeat, 9 - Duration"
    ::= { oraEM4AlertEntry 20 }

oraEM4AlertViolationGuid OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "An EM generated identifier that identifies a particular alert."
    ::= { oraEM4AlertEntry 21 }

```

```

oraEM4Traps OBJECT IDENTIFIER ::= { oraEM4 2 }

oraEM4Alert TRAP-TYPE
    ENTERPRISE  oraEM4Traps
    VARIABLES   { oraEM4AlertTargetName, oraEM4AlertTargetType,
                  oraEM4AlertHostName, oraEM4AlertMetricName,
                  oraEM4AlertKeyName, oraEM4AlertKeyValue, oraEM4AlertTimeStamp,
                  oraEM4AlertSeverity, oraEM4AlertMessage,
                  oraEM4AlertRuleName, oraEM4AlertRuleOwner,
                  oraEM4AlertMetricValue, oraEM4AlertContext,
    oraEM4AlertCycleGuid,
                  oraEM4AlertRepeatCount,
                  oraEM4AlertUDTargetProperties, oraEM4AlertAck, oraEM4AlertAckBy,
                  oraEM4AlertNotifType, oraEM4AlertViolationGuid }

    DESCRIPTION
        "The variables included in the oraEM4Alert trap."
    ::= 1

oraEM4JobAlertTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF OraEM4JobAlertEntry
    ACCESS      not-accessible
    STATUS      mandatory
    DESCRIPTION
        "Information on alerts generated by Oracle Enterprise Manager. This table is
    not queryable; it exists only to document the variables included in the
    oraEM4JobAlert trap. Each trap contains a single instance of each variable in the
    table."
    ::= { oraEM4Objects 2 }

oraEM4JobAlertEntry OBJECT-TYPE
    SYNTAX      OraEM4JobAlertEntry
    ACCESS      not-accessible
    STATUS      mandatory
    DESCRIPTION
        "Information about a particular Oracle Enterprise Manager alert."
    INDEX       { oraEM4JobAlertIndex }
    ::= { oraEM4JobAlertTable 1 }

OraEM4JobAlertEntry ::=
    SEQUENCE {
        oraEM4JobAlertIndex
            INTEGER,

        oraEM4JobAlertJobName
            DisplayString,

        oraEM4JobAlertJobOwner
            DisplayString,

        oraEM4JobAlertJobType
            DisplayString,

        oraEM4JobAlertJobStatus
            DisplayString,

        oraEM4JobAlertTargets
            DisplayString,

        oraEM4JobAlertTimeStamp
            DisplayString,
    }

```



```

        oraEM4JobAlertRuleName
        DisplayString,
        oraEM4JobAlertRuleOwner
        DisplayString,

oraEM4JobAlertMetricName
        DisplayString,

oraEM4JobAlertMetricValue
        DisplayString,

        oraEM4JobAlertContext
        DisplayString,

        oraEM4JobAlertKeyName
        DisplayString,

        oraEM4JobAlertKeyValue
        DisplayString,

        oraEM4JobAlertSeverity
        DisplayString,

        oraEM4JobAlertJobId
        DisplayString,

        oraEM4JobAlertJobExecId
        DisplayString
    }

oraEM4JobAlertIndex OBJECT-TYPE
    SYNTAX  INTEGER (0..2147483647)
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "Index of a particular alert, unique only at the moment an alert is
        generated."
    ::= { oraEM4JobAlertEntry 1 }

oraEM4JobAlertJobName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The name of the job to which this alert applies."
    ::= { oraEM4JobAlertEntry 2 }

oraEM4JobAlertJobOwner OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The owner of the job to which this alert applies."
    ::= { oraEM4JobAlertEntry 3 }

oraEM4JobAlertJobType OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory

```

```
DESCRIPTION
    "The type of the job to which this alert applies."
 ::= { oraEM4JobAlertEntry 4 }

oraEM4JobAlertJobStatus OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The status of the job to which this alert applies."
 ::= { oraEM4JobAlertEntry 5 }

oraEM4JobAlertTargets OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "A comma separated list of target to which this alert applies."
 ::= { oraEM4JobAlertEntry 6 }

oraEM4JobAlertTimeStamp OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The time at which this job status changed causing this alert."
 ::= { oraEM4JobAlertEntry 7 }

oraEM4JobAlertRuleName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The name of the notification rule that caused this notification."
 ::= { oraEM4JobAlertEntry 8 }

oraEM4JobAlertRuleOwner OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The owner of the notification rule that caused this notification."
 ::= { oraEM4JobAlertEntry 9 }

oraEM4JobAlertMetricName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The name of the metric or policy which caused the Corrective Action to run
that caused this alert."
 ::= { oraEM4JobAlertEntry 10 }

oraEM4JobAlertMetricValue OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The value of the metric which caused the Corrective Action to run that
caused this alert."
```

```

        ::= { oraEM4JobAlertEntry 11 }

oraEM4JobAlertContext OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "A comma separated list of metric column names and values associated with the
        metric which caused the Corrective Action to run that caused this alert."
    ::= { oraEM4JobAlertEntry 12 }

oraEM4JobAlertKeyName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The name of the key-column, if present, for the metric which caused the
        Corrective Action to run that generated this alert."
    ::= { oraEM4JobAlertEntry 13 }

oraEM4JobAlertKeyValue OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The value of the key-column, if present, for the metric which caused the
        Corrective Action to run that generated this alert."
    ::= { oraEM4JobAlertEntry 14 }

oraEM4JobAlertSeverity OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The severity of the metric which caused the Corrective Action to run that
        generated this alert e.g. Critical."
    ::= { oraEM4JobAlertEntry 15 }

oraEM4JobAlertJobId OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The Job Id of the EM Job that triggered this notification."
    ::= { oraEM4JobAlertEntry 16 }

oraEM4JobAlertJobExecId OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
        "The Job Execution Id of the EM Job that triggered this notification."
    ::= { oraEM4JobAlertEntry 17 }

oraEM4JobAlert TRAP-TYPE
    ENTERPRISE  oraEM4Traps
    VARIABLES   { oraEM4JobAlertJobName, oraEM4JobAlertJobOwner,
                  oraEM4JobAlertJobType, oraEM4JobAlertJobStatus,
                  oraEM4JobAlertTargets, oraEM4JobAlertTimeStamp,
                  oraEM4JobAlertRuleName, oraEM4JobAlertRuleOwner,

```

```

        oraEM4JobAlertMetricName, oraEM4JobAlertMetricValue,
        oraEM4JobAlertContext, oraEM4JobAlertKeyName,
        oraEM4JobAlertKeyValue, oraEM4JobAlertSeverity,
        oraEM4JobAlertJobId, oraEM4JobAlertJobExecId }
DESCRIPTION
    "The variables included in the oraEM4JobAlert trap."
    ::= 2

oraEMNGObjects OBJECT IDENTIFIER ::= { oraEM4 3 }

oraEMNGEventTable OBJECT-TYPE
    SYNTAX  SEQUENCE OF OraEMNGEventEntry
    ACCESS  not-accessible
    STATUS  mandatory
    DESCRIPTION
        "Information on events published to Oracle Enterprise Manager. This table is
        not queryable; it exists only to document the variables included in the
        oraEMNGEventTrap trap. Each trap can contain a single instance of each variable
        in the table."
    ::= { oraEMNGObjects 1 }

oraEMNGEventEntry OBJECT-TYPE
    SYNTAX  OraEMNGEventEntry
    ACCESS  not-accessible
    STATUS  mandatory
    DESCRIPTION
        "Information about a particular Oracle Enterprise Manager event."
    INDEX   { oraEMNGEventIndex }
    ::= { oraEMNGEventTable 1 }

OraEMNGEventEntry ::=
    SEQUENCE {
        oraEMNGEventIndex
            INTEGER,

        oraEMNGEventNotifType
            DisplayString,

        oraEMNGEventMessage
            DisplayString,

        oraEMNGEventMessageURL
            DisplayString,

        oraEMNGEventSeverity
            DisplayString,

        oraEMNGEventSeverityCode
            DisplayString,

        oraEMNGEventRepeatCount
            DisplayString,

        oraEMNGEventActionMsg
            DisplayString,

        oraEMNGEventOccurrenceTime
            DisplayString,

        oraEMNGEventReportedTime

```

```
        DisplayString,

oraEMNGEventCategories
        DisplayString,

oraEMNGEventCategoryCodes
        DisplayString,

oraEMNGEventType
        DisplayString,

oraEMNGEventName
        DisplayString,

oraEMNGAssocIncidentId
        DisplayString,

oraEMNGAssocIncidentOwner
        DisplayString,

oraEMNGAssocIncidentAcked
        DisplayString,

oraEMNGAssocIncidentStatus
        DisplayString,

oraEMNGAssocIncidentPriority
        DisplayString,

oraEMNGAssocIncidentEscLevel
        DisplayString,

oraEMNGEventTargetName
        DisplayString,

oraEMNGEventTargetNameURL
        DisplayString,

oraEMNGEventTargetType
        DisplayString,

oraEMNGEventHostName
        DisplayString,

oraEMNGEventTargetOwner
        DisplayString,

oraEMNGEventTgtLifeCycleStatus
        DisplayString,

oraEMNGEventTargetVersion
        DisplayString,

oraEMNGEventUserDefinedTgtProp
        DisplayString,

oraEMNGEventSourceObjName
        DisplayString,

oraEMNGEventSourceObjNameURL
```

```
        DisplayString,

oraEMNGEventSourceObjType
        DisplayString,

oraEMNGEventSourceObjSubType
        DisplayString,

oraEMNGEventSourceObjOwner
        DisplayString,

oraEMNGEventCAJobName
        DisplayString,

oraEMNGEventCAJobStatus
        DisplayString,

oraEMNGEventCAJobOwner
        DisplayString,

oraEMNGEventCAJobStepOutput
        DisplayString,

oraEMNGEventCAJobType
        DisplayString,

oraEMNGEventRuleSetName
        DisplayString,

oraEMNGEventRuleName
        DisplayString,

oraEMNGEventRuleOwner
        DisplayString,

oraEMNGEventSequenceId
        DisplayString,

oraEMNGEventRCADetails
        DisplayString,

oraEMNGEventContextAttrs
        DisplayString,

oraEMNGEventUserComments
        DisplayString,

oraEMNGEventUpdates
        DisplayString,

oraEMNGEventTypeAttr1
        DisplayString,

oraEMNGEventTypeAttr2
        DisplayString,

oraEMNGEventTypeAttr3
        DisplayString,

oraEMNGEventTypeAttr4
```

```
        DisplayString,

oraEMNGEEventTypeAttr5
        DisplayString,

oraEMNGEEventTypeAttr6
        DisplayString,

oraEMNGEEventTypeAttr7
        DisplayString,

oraEMNGEEventTypeAttr8
        DisplayString,

oraEMNGEEventTypeAttr9
        DisplayString,

oraEMNGEEventTypeAttr10
        DisplayString,

oraEMNGEEventTypeAttr11
        DisplayString,

oraEMNGEEventTypeAttr12
        DisplayString,

oraEMNGEEventTypeAttr13
        DisplayString,

oraEMNGEEventTypeAttr14
        DisplayString,

oraEMNGEEventTypeAttr15
        DisplayString,

oraEMNGEEventTypeAttr16
        DisplayString,

oraEMNGEEventTypeAttr17
        DisplayString,

oraEMNGEEventTypeAttr18
        DisplayString,

oraEMNGEEventTypeAttr19
        DisplayString,

oraEMNGEEventTypeAttr20
        DisplayString,

oraEMNGEEventTypeAttr21
        DisplayString,

oraEMNGEEventTypeAttr22
        DisplayString,

oraEMNGEEventTypeAttr23
        DisplayString,

oraEMNGEEventTypeAttr24
```

```
        DisplayString,

        oraEMNGEventTypeAttr25
        DisplayString
    }

oraEMNGEventIndex OBJECT-TYPE
    SYNTAX INTEGER (0..2147483647)
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Index of a particular event, unique only at the moment an event is
generated."
    ::= { oraEMNGEventEntry 1 }

oraEMNGEventNotifType OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Notification Type. NOTIF_NORMAL, NOTIF_RETRY, NOTIF_DURATION, NOTIF_REPEAT,
NOTIF_CA, NOTIF_RCA"
    ::= { oraEMNGEventEntry 2 }

oraEMNGEventMessage OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The message associated with this event."
    ::= { oraEMNGEventEntry 3 }

oraEMNGEventMessageURL OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "EM Console URL for the event message. Populated for events with severity
other than INFORMATIONAL. Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry 4 }

oraEMNGEventSeverity OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The severity of the event e.g. Critical."
    ::= { oraEMNGEventEntry 5 }

oraEMNGEventSeverityCode OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Internal code of the severity: FATAL, WARNING, CRITICAL, MINOR_WARNING,
INFORMATIONAL, CLEAR."
    ::= { oraEMNGEventEntry 6 }

oraEMNGEventRepeatCount OBJECT-TYPE
    SYNTAX DisplayString
```



```
ACCESS read-only
STATUS mandatory
DESCRIPTION
    "The repeat notification counter for the event."
::= { oraEMNGEventEntry 7 }

oraEMNGEventActionMsg OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The action message for this event."
    ::= { oraEMNGEventEntry 8 }

oraEMNGEventOccurrenceTime OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The time when this event occurred (optional), this is only populated for
events that have occurrence time."
    ::= { oraEMNGEventEntry 9 }

oraEMNGEventReportedTime OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The time when this event was reported."
    ::= { oraEMNGEventEntry 10 }

oraEMNGEventCategories OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The list of categories to which this event belongs to. Empty if trap size
exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry 11 }

oraEMNGEventCategoryCodes OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The list of internal category codes to which this event belongs to. Empty if
trap size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry 12 }

oraEMNGEventType OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The name of the event type to which this event belongs to."
    ::= { oraEMNGEventEntry 13 }

oraEMNGEventName OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
```

```
STATUS mandatory
DESCRIPTION
    "The name of this event."
 ::= { oraEMNGEventEntry 14 }

oraEMNGAssocIncidentId OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "ID of the associated incident with the event (optional)."
```

```
 ::= { oraEMNGEventEntry 15 }

oraEMNGAssocIncidentOwner OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Owner of the associated incident with the event (optional)."
```

```
 ::= { oraEMNGEventEntry 16 }

oraEMNGAssocIncidentAked OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Acknowledged status of the associated incident with the event. 1 indicates
        acknowledged, 0 indicates unacknowledged."
```

```
 ::= { oraEMNGEventEntry 17 }

oraEMNGAssocIncidentStatus OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The status of the associated incident with the event."
```

```
 ::= { oraEMNGEventEntry 18 }

oraEMNGAssocIncidentPriority OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The proirity of the associated incident with the event."
```

```
 ::= { oraEMNGEventEntry 19 }

oraEMNGAssocIncidentEscLevel OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The Escalation Level of the associated incident with the event."
```

```
 ::= { oraEMNGEventEntry 20 }

oraEMNGEventTargetName OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The name of the target to which this event applies. Populated for events
```

```

that are about a target only."
 ::= { oraEMNGEventEntry 21 }

oraEMNGEventTargetNameURL OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "EM Console URL of the target to which this event applies. Populated for
        events that are about a target only. Empty if trap size exceeds configured snmp
        packet size."
    ::= { oraEMNGEventEntry 22 }

oraEMNGEventTargetType OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The type of the target to which this event applies. Populated for events
        that are about a target only."
    ::= { oraEMNGEventEntry 23 }

oraEMNGEventHostName OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The name of the host on which this event originated. Populated for events
        that are about a target only."
    ::= { oraEMNGEventEntry 24 }

oraEMNGEventTargetOwner OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The primary administrator of the target on which this event originated.
        Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry 25 }

oraEMNGEventTgtLifeCycleStatus OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The life cycle status of the target on which this event originated."
    ::= { oraEMNGEventEntry 26 }

oraEMNGEventTargetVersion OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The version of the target on which this event originated."
    ::= { oraEMNGEventEntry 27 }

oraEMNGEventUserDefinedTgtProp OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory

```

```
DESCRIPTION
    "The user defined target properties [name,value pair list] of the associated
    target with this event. Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry 28 }

oraEMNGEventSourceObjName OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The name of the source object to which this event belongs to. Populated for
        events that are about a non-target object only, such as Job."
    ::= { oraEMNGEventEntry 29 }

oraEMNGEventSourceObjNameURL OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "EM Console URL for the source object to which this event belongs to. Empty
        if trap size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry 30 }

oraEMNGEventSourceObjType OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The type of the source object to which this event belongs to."
    ::= { oraEMNGEventEntry 31 }

oraEMNGEventSourceObjSubType OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The sub type of the source object to which this event belongs to (Optional
        property). Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry 32 }

oraEMNGEventSourceObjOwner OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The primary administrator of the source object to which this event belongs
        to. (Optional property). Empty if trap size exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry 33 }

oraEMNGEventCAJobName OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The Name of the Corrective Action Job associated with this event."
    ::= { oraEMNGEventEntry 34 }

oraEMNGEventCAJobStatus OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
```

```

STATUS mandatory
DESCRIPTION
    "The Status of the Corrective Action Job associated with this event."
 ::= { oraEMNGEventEntry 35 }

oraEMNGEventCAJobOwner OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The Owner of the Corrective Action Job associated with this event."
    ::= { oraEMNGEventEntry 36 }

oraEMNGEventCAJobStepOutput OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The job step output from the Corrective Action Job associated with this
event."
    ::= { oraEMNGEventEntry 37 }

oraEMNGEventCAJobType OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The job type from the Corrective Action Job associated with this event."
    ::= { oraEMNGEventEntry 38 }

oraEMNGEventRuleSetName OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The name of the ruleset that caused this notification. Empty if trap size
exceeds configured snmp packet size."
    ::= { oraEMNGEventEntry 39 }

oraEMNGEventRuleName OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The name of the rule within the ruleset that caused this notification."
    ::= { oraEMNGEventEntry 40 }

oraEMNGEventRuleOwner OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "The owner of the ruleset that caused this notification."
    ::= { oraEMNGEventEntry 41 }

oraEMNGEventSequenceId OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION

```

"An EM generated identifier that uniquely identifies current issue until it is cleared."

::= { oraEMNGEventEntry 42 }

oraEMNGEventRCADetails OBJECT-TYPE

SYNTAX DisplayString

ACCESS read-only

STATUS mandatory

DESCRIPTION

"Root Cause Analysis details associated with this event if it exists."

::= { oraEMNGEventEntry 43 }

oraEMNGEventContextAttrs OBJECT-TYPE

SYNTAX DisplayString

ACCESS read-only

STATUS mandatory

DESCRIPTION

"The context attributes associated with this event. Empty if trap size exceeds configured snmp packet size."

::= { oraEMNGEventEntry 44 }

oraEMNGEventUserComments OBJECT-TYPE

SYNTAX DisplayString

ACCESS read-only

STATUS mandatory

DESCRIPTION

"The user comments associated with this event. Empty if trap size exceeds configured snmp packet size."

::= { oraEMNGEventEntry 45 }

oraEMNGEventUpdates OBJECT-TYPE

SYNTAX DisplayString

ACCESS read-only

STATUS mandatory

DESCRIPTION

"The updates associated with this event. Empty if trap size exceeds configured snmp packet size."

::= { oraEMNGEventEntry 46 }

oraEMNGEventTypeAttr1 OBJECT-TYPE

SYNTAX DisplayString

ACCESS read-only

STATUS mandatory

DESCRIPTION

"Name and value pair as name=value for event type specific attribute#1."

::= { oraEMNGEventEntry 61 }

oraEMNGEventTypeAttr2 OBJECT-TYPE

SYNTAX DisplayString

ACCESS read-only

STATUS mandatory

DESCRIPTION

"Name and value pair as name=value for event type specific attribute#2."

::= { oraEMNGEventEntry 62 }

oraEMNGEventTypeAttr3 OBJECT-TYPE

SYNTAX DisplayString

ACCESS read-only

STATUS mandatory

DESCRIPTION

```
        "Name and value pair as name=value for event type specific attribute#3."
        ::= { oraEMNGEventEntry 63 }

oraEMNGEventTypeAttr4 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#4."
        ::= { oraEMNGEventEntry 64 }

oraEMNGEventTypeAttr5 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#5."
        ::= { oraEMNGEventEntry 65 }

oraEMNGEventTypeAttr6 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#6."
        ::= { oraEMNGEventEntry 66 }

oraEMNGEventTypeAttr7 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#7."
        ::= { oraEMNGEventEntry 67 }

oraEMNGEventTypeAttr8 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#8."
        ::= { oraEMNGEventEntry 68 }

oraEMNGEventTypeAttr9 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#9."
        ::= { oraEMNGEventEntry 69 }

oraEMNGEventTypeAttr10 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#10."
        ::= { oraEMNGEventEntry 70 }

oraEMNGEventTypeAttr11 OBJECT-TYPE
```

```
SYNTAX DisplayString
ACCESS read-only
STATUS mandatory
DESCRIPTION
    "Name and value pair as name=value for event type specific attribute#11."
 ::= { oraEMNGEventEntry 71 }

oraEMNGEventTypeAttr12 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#12."
    ::= { oraEMNGEventEntry 72 }

oraEMNGEventTypeAttr13 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#13."
    ::= { oraEMNGEventEntry 73 }

oraEMNGEventTypeAttr14 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#14."
    ::= { oraEMNGEventEntry 74 }

oraEMNGEventTypeAttr15 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#15."
    ::= { oraEMNGEventEntry 75 }

oraEMNGEventTypeAttr16 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#16."
    ::= { oraEMNGEventEntry 76 }

oraEMNGEventTypeAttr17 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#17."
    ::= { oraEMNGEventEntry 77 }

oraEMNGEventTypeAttr18 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
```



```

        "Name and value pair as name=value for event type specific attribute#18."
        ::= { oraEMNGEventEntry 78 }

oraEMNGEventTypeAttr19 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#19."
        ::= { oraEMNGEventEntry 79 }

oraEMNGEventTypeAttr20 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#20."
        ::= { oraEMNGEventEntry 80 }

oraEMNGEventTypeAttr21 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#21."
        ::= { oraEMNGEventEntry 81 }

oraEMNGEventTypeAttr22 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#22."
        ::= { oraEMNGEventEntry 82 }

oraEMNGEventTypeAttr23 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#23."
        ::= { oraEMNGEventEntry 83 }

oraEMNGEventTypeAttr24 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#24."
        ::= { oraEMNGEventEntry 84 }

oraEMNGEventTypeAttr25 OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "Name and value pair as name=value for event type specific attribute#25."
        ::= { oraEMNGEventEntry 85 }

oraEMNGEvent TRAP-TYPE

```

```

ENTERPRISE  oraEM4Traps
VARIABLES  {
    oraEMNGEventIndex, oraEMNGEventNotifType,
    oraEMNGEventMessage, oraEMNGEventMessageURL,
    oraEMNGEventSeverity, oraEMNGEventSeverityCode,
    oraEMNGEventRepeatCount, oraEMNGEventActionMsg,
    oraEMNGEventOccurrenceTime, oraEMNGEventReportedTime,
    oraEMNGEventCategories, oraEMNGEventCategoryCodes,
    oraEMNGEventType, oraEMNGEventName,
    oraEMNGAssocIncidentId, oraEMNGAssocIncidentOwner,
    oraEMNGAssocIncidentAked, oraEMNGAssocIncidentStatus,
    oraEMNGAssocIncidentPriority, oraEMNGAssocIncidentEscLevel,
    oraEMNGEventTargetName, oraEMNGEventTargetNameURL,
    oraEMNGEventTargetType, oraEMNGEventHostName,
    oraEMNGEventTargetOwner, oraEMNGEventTgtLifeCycleStatus,
    oraEMNGEventTargetVersion, oraEMNGEventUserDefinedTgtProp,
    oraEMNGEventSourceObjName, oraEMNGEventSourceObjNameURL,
    oraEMNGEventSourceObjType, oraEMNGEventSourceObjSubType,
    oraEMNGEventSourceObjOwner, oraEMNGEventCAJobName,
    oraEMNGEventCAJobStatus, oraEMNGEventCAJobOwner,
    oraEMNGEventCAJobStepOutput, oraEMNGEventCAJobType,
    oraEMNGEventRuleSetName, oraEMNGEventRuleName,
    oraEMNGEventRuleOwner, oraEMNGEventSequenceId,
    oraEMNGEventRCADetails, oraEMNGEventContextAttrs,
    oraEMNGEventUserComments, oraEMNGEventUpdates,
    oraEMNGEventTypeAttr1, oraEMNGEventTypeAttr2,
    oraEMNGEventTypeAttr3, oraEMNGEventTypeAttr4,
    oraEMNGEventTypeAttr5, oraEMNGEventTypeAttr6,
    oraEMNGEventTypeAttr7, oraEMNGEventTypeAttr8,
    oraEMNGEventTypeAttr9, oraEMNGEventTypeAttr10,
    oraEMNGEventTypeAttr11, oraEMNGEventTypeAttr12,
    oraEMNGEventTypeAttr13, oraEMNGEventTypeAttr14,
    oraEMNGEventTypeAttr15, oraEMNGEventTypeAttr16,
    oraEMNGEventTypeAttr17, oraEMNGEventTypeAttr18,
    oraEMNGEventTypeAttr19, oraEMNGEventTypeAttr20,
    oraEMNGEventTypeAttr21, oraEMNGEventTypeAttr22,
    oraEMNGEventTypeAttr23, oraEMNGEventTypeAttr24,
    oraEMNGEventTypeAttr25
}
DESCRIPTION
    "The variables included in the oraEMNGAlert trap."
 ::= 3
END

```

SNMP Trap Mappings

The following tables list SNMP trap mappings between Enterprise Manager 12c and previous releases.

C.1 Pre-12c Enterprise Manager Metric Alerts

Prior to Enterprise Manager 12c, metric alerts were sent using the oraEM4Alert trap type. In 12c, the event type corresponding to these alerts is metric alert. The value for oraEMNGEvent type in an Enterprise Manager 12c SNMP trap would be set to 'Metric Alert'.

Table C-1 Metric Alert Mappings

Pre-12C OID Number	Pre-12C OID Name	12C OID Number	12C OID Name
1.3.6.1.4.1.111.15.1.1.1.2.1	oraEM4AlertTargetName	1.3.6.1.4.1.111.15.3.1.1.21.1	oraEMNGEventTargetName
1.3.6.1.4.1.111.15.1.1.1.3.1	oraEM4AlertTargetType	1.3.6.1.4.1.111.15.3.1.1.23.1	oraEMNGEventTargetType
1.3.6.1.4.1.111.15.1.1.1.4.1	oraEM4AlertHostName	1.3.6.1.4.1.111.15.3.1.1.24.1	oraEMNGEventHostName
1.3.6.1.4.1.111.15.1.1.1.5.1	oraEM4AlertMetricName	1.3.6.1.4.1.111.15.3.1.1.65.1	oraEMNGEventTypeAttr5
1.3.6.1.4.1.111.15.1.1.1.6.1	oraEM4AlertKeyName	1.3.6.1.4.1.111.15.3.1.1.66.1	oraEMNGEventTypeAttr6
1.3.6.1.4.1.111.15.1.1.1.7.1	oraEM4AlertKeyValue	* See the note below for details	* See the note below for details
1.3.6.1.4.1.111.15.1.1.1.8.1	oraEM4AlertTimeStamp	1.3.6.1.4.1.111.15.3.1.1.10.1	oraEMNGEventReportedTime
1.3.6.1.4.1.111.15.1.1.1.9.1	oraEM4AlertSeverity	1.3.6.1.4.1.111.15.3.1.1.5.1	oraEMNGEventSeverity
1.3.6.1.4.1.111.15.1.1.1.10.1	oraEM4AlertMessage	1.3.6.1.4.1.111.15.3.1.1.3.1	oraEMNGEventMessage
1.3.6.1.4.1.111.15.1.1.1.11.1	oraEM4AlertRuleName	1.3.6.1.4.1.111.15.3.1.1.39.1	oraEMNGEventRuleSetName
1.3.6.1.4.1.111.15.1.1.1.12.1	oraEM4AlertRuleOwner	1.3.6.1.4.1.111.15.3.1.1.41.1	oraEMNGEventRuleOwner
1.3.6.1.4.1.111.15.1.1.1.13.1	oraEM4AlertMetricValue	1.3.6.1.4.1.111.15.3.1.1.68.1	oraEMNGEventTypeAttr8
1.3.6.1.4.1.111.15.1.1.1.14.1	oraEM4AlertContext	1.3.6.1.4.1.111.15.3.1.1.44.1	oraEMNGEventContextAttrs
1.3.6.1.4.1.111.15.1.1.1.15.1	oraEM4AlertCycleGuid	1.3.6.1.4.1.111.15.3.1.1.70.1	oraEMNGEventTypeAttr3
1.3.6.1.4.1.111.15.1.1.1.16.1	oraEM4AlertRepeatCount	1.3.6.1.4.1.111.15.3.1.1.7.1	oraEMNGEventRepeatCount
1.3.6.1.4.1.111.15.1.1.1.17.1	oraEM4AlertUDTargetProperties	1.3.6.1.4.1.111.15.3.1.1.28.1	oraEMNGEventUserDefinedTgtProp
1.3.6.1.4.1.111.15.1.1.1.18.1	oraEM4AlertAck	1.3.6.1.4.1.111.15.3.1.1.17.1	oraEMNGAssocIncidentAcked

Table C–1 (Cont.) Metric Alert Mappings

Pre-12C OID Number	Pre-12C OID Name	12C OID Number	12C OID Name
1.3.6.1.4.1.111.15.1.1.1.19.1	oraEM4AlertAckBy	1.3.6.1.4.1.111.15.3.1.1.16.1	oraEMNGAssocIncidentOwner
1.3.6.1.4.1.111.15.1.1.1.20.1	oraEM4AlertNotifType	1.3.6.1.4.1.111.15.3.1.1.2.1	oraEMNGEventNotifType
1.3.6.1.4.1.111.15.1.1.1.21.1	oraEM4AlertViolationGui d	1.3.6.1.4.1.111.15.3.1.1.42.1	oraEMNGEventSequenceId

C.2 Pre-12C Target Availability Alerts

Prior to Enterprise Manager 12c, target availability alerts were sent using oraEM4Alert SNMP trap type. In 12c, the event type corresponding to these alerts is target_availability. Value for oraEMNGEventType in the 12C trap would be set to 'Target Availability'.

Table C–2 Target Availability Alert Mappings

Pre-12C OID Number	Pre-12C OID Name	12C OID Number	12C OID Name
1.3.6.1.4.1.111.15.1.1.1.2.1	oraEM4AlertTargetName	1.3.6.1.4.1.111.15.3.1.1.21.1	oraEMNGEventTargetName
1.3.6.1.4.1.111.15.1.1.1.3.1	oraEM4AlertTargetType	1.3.6.1.4.1.111.15.3.1.1.23.1	oraEMNGEventTargetType
1.3.6.1.4.1.111.15.1.1.1.4.1	oraEM4AlertHostName	1.3.6.1.4.1.111.15.3.1.1.24.1	oraEMNGEventHostName
1.3.6.1.4.1.111.15.1.1.1.5.1	oraEM4AlertMetricName	N/A	N/A
1.3.6.1.4.1.111.15.1.1.1.6.1	oraEM4AlertKeyName	// deprecated in 12C, was always null in 11GC	// deprecated in 12C, was always null in 11GC
1.3.6.1.4.1.111.15.1.1.1.7.1	oraEM4AlertKeyValue	N/A	N/A
1.3.6.1.4.1.111.15.1.1.1.8.1	oraEM4AlertTimeStamp	1.3.6.1.4.1.111.15.3.1.1.10.1	oraEMNGEventReportedTime
1.3.6.1.4.1.111.15.1.1.1.9.1	oraEM4AlertSeverity	1.3.6.1.4.1.111.15.3.1.1.61.1	oraEMNGEventTypeAttr1 //target_status
1.3.6.1.4.1.111.15.1.1.1.10.1	oraEM4AlertMessage	1.3.6.1.4.1.111.15.3.1.1.3.1	oraEMNGEventMessage
1.3.6.1.4.1.111.15.1.1.1.11.1	oraEM4AlertRuleName	1.3.6.1.4.1.111.15.3.1.1.39.1	oraEMNGEventRuleSetName
1.3.6.1.4.1.111.15.1.1.1.12.1	oraEM4AlertRuleOwner	1.3.6.1.4.1.111.15.3.1.1.41.1	oraEMNGEventRuleOwner
1.3.6.1.4.1.111.15.1.1.1.13.1	oraEM4AlertMetricValue	N/A	N/A
1.3.6.1.4.1.111.15.1.1.1.14.1	oraEM4AlertContext	1.3.6.1.4.1.111.15.3.1.1.44.1	oraEMNGEventContextAttrs
1.3.6.1.4.1.111.15.1.1.1.15.1	oraEM4AlertCycleGuid	1.3.6.1.4.1.111.15.3.1.1.66.1	oraEMNGEventTypeAttr6
1.3.6.1.4.1.111.15.1.1.1.16.1	oraEM4AlertRepeatCount	1.3.6.1.4.1.111.15.3.1.1.7.1	oraEMNGEventRepeatCount
1.3.6.1.4.1.111.15.1.1.1.17.1	oraEM4AlertUDTargetPro perties	1.3.6.1.4.1.111.15.3.1.1.28.1	oraEMNGEventUserDefinedTg tProp
1.3.6.1.4.1.111.15.1.1.1.18.1	oraEM4AlertAck	1.3.6.1.4.1.111.15.3.1.1.17.1	oraEMNGAssocIncidentAcked
1.3.6.1.4.1.111.15.1.1.1.19.1	oraEM4AlertAckBy	1.3.6.1.4.1.111.15.3.1.1.16.1	oraEMNGAssocIncidentOwner
1.3.6.1.4.1.111.15.1.1.1.20.1	oraEM4AlertNotifType	1.3.6.1.4.1.111.15.3.1.1.2.1	oraEMNGEventNotifType
1.3.6.1.4.1.111.15.1.1.1.21.1	oraEM4AlertViolationGui d	1.3.6.1.4.1.111.15.3.1.1.42.1	oraEMNGEventSequenceId

C.3 Pre-12C Corrective Action Results for Metric Alerts

Prior to Enterprise Manager 12c, corrective action results for metric alerts were sent using the oraEM4JobAlert trap type. For Enterprise Manager 12c, the event type

corresponding to these alerts is metric alert. The value for oraEMNGEvent in the 12c trap would be set to 'Metric Alert'.

Table C-3 Corrective Action Results for Metric Alert Mappings

Pre-12C OID Number	Pre-12C OID Name	12C OID Number	12C OID Name
1.3.6.1.4.1.111.15.1.2.1.2.1	oraEM4JobAlertJobName	1.3.6.1.4.1.111.15.3.1.1.34.1	oraEMNGEventCAJobName
1.3.6.1.4.1.111.15.1.2.1.3.1	oraEM4JobAlertJobOwner	1.3.6.1.4.1.111.15.3.1.1.36.1	oraEMNGEventCAJobOwner
1.3.6.1.4.1.111.15.1.2.1.4.1	oraEM4JobAlertJobType	1.3.6.1.4.1.111.15.3.1.1.38.1	oraEMNGEventCAJobType
1.3.6.1.4.1.111.15.1.2.1.5.1	oraEM4JobAlertJobStatus	1.3.6.1.4.1.111.15.3.1.1.35.1	oraEMNGEventCAJobStatus
1.3.6.1.4.1.111.15.1.2.1.6.1	oraEM4JobAlertTargets	1.3.6.1.4.1.111.15.3.1.1.23.1	oraEMNGEventTargetType
		1.3.6.1.4.1.111.15.3.1.1.21.1	oraEMNGEventTargetName
1.3.6.1.4.1.111.15.1.2.1.7.1	oraEM4JobAlertTimeStam P	1.3.6.1.4.1.111.15.3.1.1.10.1	oraEMNGEventReportedTime
1.3.6.1.4.1.111.15.1.2.1.8.1	oraEM4JobAlertRuleName	1.3.6.1.4.1.111.15.3.1.1.39.1	oraEMNGEventRuleSetName
1.3.6.1.4.1.111.15.1.2.1.9.1	oraEM4JobAlertRuleOwne r	1.3.6.1.4.1.111.15.3.1.1.41.1	oraEMNGEventRuleOwner
1.3.6.1.4.1.111.15.1.2.1.10.1	oraEM4JobAlertMetricNa me	1.3.6.1.4.1.111.15.3.1.1.65.1	oraEMNGEventTypeAttr5
1.3.6.1.4.1.111.15.1.2.1.11.1	oraEM4JobAlertMetricVal ue	1.3.6.1.4.1.111.15.3.1.1.68.1	oraEMNGEventTypeAttr8
1.3.6.1.4.1.111.15.1.2.1.12.1	oraEM4JobAlertContext	1.3.6.1.4.1.111.15.3.1.1.44.1	oraEMNGEventContextAttrs
1.3.6.1.4.1.111.15.1.2.1.13.1	oraEM4JobAlertKeyName	1.3.6.1.4.1.111.15.3.1.1.66.1	oraEMNGEventTypeAttr6
1.3.6.1.4.1.111.15.1.2.1.14.1	oraEM4JobAlertKeyValue	1.3.6.1.4.1.111.15.3.1.1.69.1	oraEMNGEventTypeAttr9
1.3.6.1.4.1.111.15.1.2.1.15.1	oraEM4JobAlertSeverity	1.3.6.1.4.1.111.15.3.1.1.5.1	oraEMNGEventSeverity
1.3.6.1.4.1.111.15.1.2.1.16.1	oraEM4JobAlertJobId	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.17.1	oraEM4JobAlertJobExecId	N/A	N/A

C.4 Corrective Action Results for Target Availability

Prior to Enterprise Manager 12c, corrective action results for target availability alerts were sent using the oraEM4JobAlert trap type. In 12c, the event type corresponding to these alerts is target_availability alert. The value for oraEMNGEvent in the 12c trap would be set to 'Metric Alert'.

Table C-4 Target Availability Mappings

Pre-12C OID Number	Pre-12C OID Name	12C OID Number	12C OID Name
1.3.6.1.4.1.111.15.1.2.1.2.1	oraEM4JobAlertJobName	1.3.6.1.4.1.111.15.3.1.1.34.1	oraEMNGEventCAJobName
1.3.6.1.4.1.111.15.1.2.1.3.1	oraEM4JobAlertJobOwner	1.3.6.1.4.1.111.15.3.1.1.36.1	oraEMNGEventCAJobOwner
1.3.6.1.4.1.111.15.1.2.1.4.1	oraEM4JobAlertJobType	1.3.6.1.4.1.111.15.3.1.1.38.1	oraEMNGEventCAJobType
1.3.6.1.4.1.111.15.1.2.1.5.1	oraEM4JobAlertJobStatus	1.3.6.1.4.1.111.15.3.1.1.35.1	oraEMNGEventCAJobStatus
1.3.6.1.4.1.111.15.1.2.1.6.1	oraEM4JobAlertTargets	1.3.6.1.4.1.111.15.3.1.1.23.1	oraEMNGEventTargetType and
		1.3.6.1.4.1.111.15.3.1.1.21.1	oraEMNGEventTargetName
1.3.6.1.4.1.111.15.1.2.1.7.1	oraEM4JobAlertTimeStam P	1.3.6.1.4.1.111.15.3.1.1.10.1	oraEMNGEventReportedTime

Table C–4 (Cont.) Target Availability Mappings

Pre-12C OID Number	Pre-12C OID Name	12C OID Number	12C OID Name
1.3.6.1.4.1.111.15.1.2.1.8.1	oraEM4JobAlertRuleName	1.3.6.1.4.1.111.15.3.1.1.39.1	oraEMNGEventRuleSetName
1.3.6.1.4.1.111.15.1.2.1.9.1	oraEM4JobAlertRuleOwner	1.3.6.1.4.1.111.15.3.1.1.41.1	oraEMNGEventRuleOwner
1.3.6.1.4.1.111.15.1.2.1.10.1	oraEM4JobAlertMetricName	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.11.1	oraEM4JobAlertMetricValue	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.12.1	oraEM4JobAlertContext	1.3.6.1.4.1.111.15.3.1.1.44.1	oraEMNGEventContextAttrs
1.3.6.1.4.1.111.15.1.2.1.13.1	oraEM4JobAlertKeyName	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.14.1	oraEM4JobAlertKeyValue	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.15.1	oraEM4JobAlertSeverity	1.3.6.1.4.1.111.15.3.1.1.61.1	oraEMNGEventTypeAttr5 // target_status
1.3.6.1.4.1.111.15.1.2.1.16.1	oraEM4JobAlertJobId	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.17.1	oraEM4JobAlertJobExecId	N/A	N/A

C.5 Job Status Change

Prior to Enterprise Manager 12c , job status change was sent using oraEM4JobAlert trap type. For 12c, the event type corresponding to these alerts is the job_status_change alert. The value for the oraEMNGEventType in the 12c trap would be set to 'Job Status Change'.

Table C–5 Job Status Change Mappings

Pre-12c OID Number	Pre-12c OID Name	12c OID Number	12c OID Name
1.3.6.1.4.1.111.15.1.2.1.2.1	oraEM4JobAlertJobName	1.3.6.1.4.1.111.15.3.1.1.29.1	oraEMNGEventSourceObjName
1.3.6.1.4.1.111.15.1.2.1.3.1	oraEM4JobAlertJobOwner	1.3.6.1.4.1.111.15.3.1.1.33.1	oraEMNGEventSourceObjOwner
1.3.6.1.4.1.111.15.1.2.1.4.1	oraEM4JobAlertJobType	1.3.6.1.4.1.111.15.3.1.1.32.1	oraEMNGEventSourceObjSubType
1.3.6.1.4.1.111.15.1.2.1.5.1	oraEM4JobAlertJobStatus	1.3.6.1.4.1.111.15.3.1.1.62.1	oraEMNGEventTypeAttr2
1.3.6.1.4.1.111.15.1.2.1.6.1	oraEM4JobAlertTargets	1.3.6.1.4.1.111.15.3.1.1.23.1	oraEMNGEventTargetType and
		1.3.6.1.4.1.111.15.3.1.1.21.1	oraEMNGEventTargetName
1.3.6.1.4.1.111.15.1.2.1.7.1	oraEM4JobAlertTimeStamp	1.3.6.1.4.1.111.15.3.1.1.10.1	oraEMNGEventReportedTime
1.3.6.1.4.1.111.15.1.2.1.8.1	oraEM4JobAlertRuleName	1.3.6.1.4.1.111.15.3.1.1.39.1	oraEMNGEventRuleSetName
1.3.6.1.4.1.111.15.1.2.1.9.1	oraEM4JobAlertRuleOwner	1.3.6.1.4.1.111.15.3.1.1.41.1	oraEMNGEventRuleOwner
1.3.6.1.4.1.111.15.1.2.1.10.1	oraEM4JobAlertMetricName	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.11.1	oraEM4JobAlertMetricValue	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.12.1	oraEM4JobAlertContext	1.3.6.1.4.1.111.15.3.1.1.44.1	oraEMNGEventContextAttrs
1.3.6.1.4.1.111.15.1.2.1.13.1	oraEM4JobAlertKeyName	N/A	N/A

Table C-5 (Cont.) Job Status Change Mappings

Pre-12c OID Number	Pre-12c OID Name	12c OID Number	12c OID Name
1.3.6.1.4.1.111.15.1.2.1.14.1	oraEM4JobAlertKeyValue	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.15.1	oraEM4JobAlertSeverity	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.16.1	oraEM4JobAlertJobId	N/A	N/A
1.3.6.1.4.1.111.15.1.2.1.17.1	oraEM4JobAlertJobExecId	1.3.6.1.4.1.111.15.3.1.1.61.1	oraEMNGEventAttr1

*** Note:** When mapping 1.3.6.1.4.1.111.15.1.1.1.7.1 oraEM4AlertKeyValue to 12g metric_alert event to 1.3.6.1.4.1.111.15.1.1.1.7.1 oraEM4AlertKeyValue, you need to look at 1.3.6.1.4.1.111.15.3.1.1.84.1 oraEMNGEventAttr24.

```

if oraEMNGEventAttr24 is null
then
    oraEM4AlertKeyValue is null

if oraEMNGEventAttr24 value = "Number of keys=1"
    oraEM4AlertKeyValue --> oraEMNGEventAttr8

if oraEMNGEventAttr24 value = "Number of keys=x" where x is greater
than 1
=> check the values for the following pairs of attributes.
    <oraEMNGEventAttr10, oraEMNGEventAttr11>
    <oraEMNGEventAttr12, oraEMNGEventAttr13>
    <oraEMNGEventAttr14, oraEMNGEventAttr15>
    <oraEMNGEventAttr16, oraEMNGEventAttr17>
    <oraEMNGEventAttr18, oraEMNGEventAttr19>
    <oraEMNGEventAttr20, oraEMNGEventAttr21>
    <oraEMNGEventAttr22, oraEMNGEventAttr23>
    ...
    ...

```

As many pairs as the number of parts present in the key would be populated, the rest of will be set to null.

For each non-null pair of attributes, the first attribute provides the name for that part of the key and second attribute provides the value for that part of the key.

Important: OID 1.3.6.1.4.1.111.15.3.1.1.13.1 specifies the event type.

Examples:

For a metric_alert event type

oraEMNGEventAttr=Metric Alert

For a target_availability event type,

oraEMNGEventAttr=Target Availability

For a job_status_change event type

oraEMNGEventAttr=Job Status Change

Standby OMSs using Standby WebLogic Domain

The following standby OMS implementation is used for Enterprise Manager Release 12.1.0.2 and earlier installations. For Enterprise Manager Release 12.1.0.3 and later, see [Chapter 28, "Enterprise Manager Disaster Recovery."](#)

IMPORTANT: Standby OMSs using Standby WebLogic Domain are still supported but have been deprecated and may be desupported in a future release (see My Oracle Support Note 1563541.1 for details).

D.1 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.

D.1.1 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 till 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 plug-ins 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 plug-ins 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,orac
le.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.

D.1.2 Installing the First Standby Management Service (Oracle Installer, Software Only Mode)

1. Copy the emkey to the Management Repository by running the following command on the first OMS on the primary site:

```
emctl config emkey -copy_to_repos
```

2. Export the configuration from the first OMS on the primary site.

```
emctl exportconfig oms -dir <location for the export file>
```

After the configuration is exported, do not make any configuration changes to the primary site until the standby OMS is configured.

3. On the remote host, perform a software-only installation of the standby OMS as described in the Oracle® Enterprise Manager Cloud Control Advanced Installation and Configuration Guide chapter "Installing Enterprise Manager Software Now and Configuring Later" (Installing in Graphical Mode). Ensure that you install the software binaries in the same middleware location as that of the primary OMS.

Note: Make sure you select same set of plug-ins as the primary OMS.

If you have installed extra plug-ins which are not part of the installation, then complete the following steps.

1. Connect to repository and execute the following query to obtain the list of installed plug-ins:

```
SELECT epv.plugin_id "plugin id", epv.version "version", epv.rev_version  
"revision", epv.su_entity_id "update id"  
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;
```

Write down the returned list of plug-ins installed.

2. Extract the plug-in archives from the primary site.

```
emcli export_update -id=<update id> -host=<standby OMS host>  
-dir=<directory to export archives> <host credential options>
```

where <update id> is the value returned by the four column of query above.

Note that this command generates a file with a .zip extension. Rename the file from <file_name>.zip to <file_name>.opar .

Repeat above steps for all plugins. In standby OMS host, place all .opar files in one directory <plugin dir>.

3. Invoke runInstaller and choose software-only install. This will install only the middleware home and OMS Oracle homes.
4. To install the required plug-ins, run the plug-in install script.

```
<OMS_HOME>/sysman/install/PluginInstall.sh -pluginLocation <plugin dir>
```

where <plugin dir> is the directory where you placed extra plug-ins in step two above. Running this command displays a screen that will let you select the plug-ins to install. Verify that all the plug-ins listed in step 1 appear and that you select them. If there is a mismatch, OMS configuration may fail at a later step.

4. On the software-only installation, apply all the patches you applied on the primary OMS so that both the primary and standby OMSs are identical and are in sync.

5. Configure the OMS 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 Enterprise Manager 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 plug-ins by running the following command:

```
pluginca -action deploy -isFirstOMS true -plugins <plug-in list> -oracleHome
<oms oracle home> -middlewareHome <wls middleware home>
```

where *plug-in list* is the list of plug-ins 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;
```

The plug-in list must be 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,oracle.sysman.emas=
12.1.0.1.0,oracle.sysman.emfa=12.1.0.1.0,oracle.sysman.db=12.1.0.1.0,oracle.sys
man.emct=12.1.0.1.0,oracle.sysman.vt=12.1.0.1.0,oracle.sysman.ssa=12.1.0.1.0"
```

7. Copy over the configuration exported from the primary OMS in step 2 above to the standby OMS host. Import the exported configuration on the standby OMS:

```
emctl importconfig oms -file <full path of the export file>
```

Note that this command generates a warning about a failed export and prompts for confirmation in order to proceed. The warning can be ignored by entering "y" to proceed.

This command will automatically start the OMS.

8. Stop the Management Service but leave the Administration Server running using:

```
emctl stop oms
```

9. Configure the Management Agent on the second OMS host by running the following command from the OMS home:

```
$<AGENT_HOME>/sysman/install/agentDeploy.sh AGENT_BASE_DIR=<middleware_
home>/agent OMS_HOST=<oms_host_name> EM_UPLOAD_PORT=<oms_port> AGENT_
REGISTRATION_PASSWORD=<password> -configOnly
```

Note: If you have an SLB configured, then directly specify the host and port of the primary load balancer.

If no SLB is configured, then use the secure upload port of primary OMS.

Example:

```
./agentDeploy.sh AGENT_BASE_DIR=$MIDDLEWARE_BASE/agent OMS_HOST=prim_oms_
hhost.domain.com EM_UPLOAD_PORT=4900 AGENT_REGISTRATION_PASSWORD=password
-configOnly
```

10. Deploy the required plug-ins on the Management Agent.

For information about deploying plug-ins, refer to the section *Deploying and Updating Plug-ins* in the [Chapter 21, "Updating Cloud Control"](#).

11. 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" since the standby OMS is down at this stage.

12. If you have Single Sign-On configured on the primary site, follow the same steps to configure SSO on the standby OMS.
13. 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.

D.1.3 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".

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 plug-ins 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 plug-ins 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,orac
le.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.

D.1.4 Installing Additional Standby Management Services (Oracle Installer, Software Only Mode)

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 until the standby management service is configured.

3. On the remote host, perform a software-only installation of the standby OMS as described in the Oracle® Enterprise Manager Cloud Control Advanced Installation and Configuration Guide chapter "Installing Enterprise Manager Software Now and Configuring Later" (Installing in Graphical Mode). Ensure that you install the software binaries in the same middleware location as that of the primary OMS.

Note: Make sure you select same set of plug-ins as the primary OMS.

If you have installed extra plug-ins which are not part of the installation, then complete the following steps.

1. Connect to repository and execute the following query to obtain the list of installed plug-ins:

```
SELECT epv.plugin_id "plugin id", epv.version "version", epv.rev_version  
"revision", epv.su_entity_id "update id"  
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;
```

Write down the returned list of plug-ins installed.

2. Extract the plug-in archives from the primary site.

```
emcli export_update -id=<update id> -host=<standby OMS host>  
-dir=<directory to export archives> <host credential options>
```

where <update id> is the value returned by the fourth column of query above.

Note that this command generates a file with a .zip extension. Rename the file from <file_name>.zip to <file_name>.opar.

Repeat above steps for all plug-ins. In standby OMS host, place all .opar files in one directory <plugin dir>.

3. Invoke runInstaller and choose software-only install. This will install only the middleware home and OMS Oracle homes.
4. To install the required plug-ins, run the plug-in install script.

```
<OMS_HOME>/sysman/install/PluginInstall.sh -pluginLocation <plugin dir>
```

where <plugin dir> is the directory where you placed extra plug-ins in step two above. Running this command displays a screen that will let you select the plug-ins to install. Verify that all the plug-ins listed in step 1 appear and that you select them. If there is a mismatch, OMS configuration may fail at a later step.

4. On the software-only installation, apply all the patches you applied on the primary OMS so that both the primary and standby OMSs are identical and are in sync.

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.

To keep the ports consistent with your primary site, provide all the ports on the command line itself. For example:

```
omsca add -MSPORT 7202 -MS_HTTPS_PORT 7301 -EM_NODEMGR_PORT 7403 -EM_UPLOAD_
PORT 4889 -EM_UPLOAD_HTTPS_PORT 4900 -EM_CONSOLE_PORT 7788 -EM_CONSOLE_HTTPS_
PORT 7799 -nostart
```

6. Configure the required plug-ins by running the following command:

```
pluginca -action deploy -isFirstOMS false -plugins
<plugin-list> -oracleHome <oms oracle home> -middlewareHome
<wls middleware home>
```

where the plugin-list is the list of plug-ins returned by the SQL query above and is a comma separated 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,orac
le.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 automatically start the OMS.

8. Stop the OMS using:

```
emctl stop oms
```

9. Configure the Management Agent on the second OMS host by running the following command from the OMS home:

```
$<AGENT_HOME>/sysman/install/agentDeploy.sh AGENT_BASE_DIR=<middleware_
home>/agent OMS_HOST=<second_oms_host_name>
EM_UPLOAD_PORT=<second_oms_port> AGENT_REGISTRATION_PASSWORD=<password>
-configOnly
```

Note: If you have an SLB configured, then directly specify the host and port of the primary load balancer.

If no SLB is configured, then use secure upload port of primary OMS.

10. Deploy the required plug-ins on the Management Agent.

For information about deploying plug-ins, refer to the section *Deploying and Updating Plug-ins* in the [Chapter 21, "Updating Cloud Control"](#).

11. 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.
12. If you have Single Sign On configured on the primary site, follow the same steps to configure SSO on the standby OMS.
13. 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.

D.1.5 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.

D.1.5.1 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 Plug-ins

When new plug-ins are deployed on the primary site or existing plug-ins 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 assumes that the standby site was setup as per the documented process and the standby management services are currently down and point to the primary repository. The plug-in(s) deployment on the primary site has been completed successfully.

Deploying a New Plug-in or Upgrading a Plug-in on Standby Site

1. Extract the plug-in archives from the primary site

Go to the Self Update Home, click on Plug-ins, select the required plug-in 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. Run prerequisite checks for the plug-in and apply the required patch reported by after running the following command.

```
./emctl plugin deploy -action prereqcheck -isFirstOMS true -archives <exported_plugin_archive>
```

where `<exported_plugin_archive>` is the zip file that was exported in Step 1:
`<version>_OMS_<platform>_<revision>.zip`

4. Configure the plug-in on the first standby OMS Oracle Home.

```
./emctl plugin deploy -action deploy -isFirstOMS true -archives <exported_plugin_archive>
```

where `exported_plugin_archive` represents the zip file `<version>_OMS_<platform>_<revision>.zip`

5. Repeat steps 3 and 4 for each standby additional OMS/

```
./emctl plugin deploy -action prereqcheck -isFirstOMS true -archives <exported_plugin_archive>
```

```
./emctl plugin deploy -action deploy -isFirstOMS true -archives <exported_plugin_archive>
```

This completes the plug-in deployment on the standby site.

Sync up Sysman Credentials

When administrators modify sysman credentials on the primary site, the following procedure must be run on the standby site in order to keep the standby management service in sync with respect to sysman credentials.

If the sysman credentials for the primary and standby sites are not kept in sync, they will fail to start when switchover or failover operations are attempted. In addition, the administrator will not be able to log into the standby OMS once the switchover operation is complete.

The procedure below assumes that the standby site was set up according to the documented process and the standby management services are currently down and point to the primary repository. The plug-in(s) deployment on the primary site has been completed successfully.

Perform following steps on standby management service

1. Start the standby Administration Server, if it is down

```
emctl start oms -admin_only
```

2. Change password for the sysman user. You must also provide the SYS password.

```
emctl config oms -change_repos_pwd -use_sys_pwd -sys_pwd  
<sys_pwd>
```

3. Stop the standby Administration Server,

```
emctl stop oms -all
```

Index

A

- accessing Software Library Administration
 - page, 11-7
- accessing Software Library console, 11-1
- Adapters, 8-18
- Add HTTP Location, 11-13
- Add NFS Location, 11-13
- Add OMS Agent file system, 11-12
- Add OMS Agent file system location, 11-13
- Add OMS Shared file system, 11-11
- adding targets, 18-6
- administration group
 - creating, 6-9
 - definition, 6-1
 - hierarchy, 6-11
 - home page, 6-25
- Administration Groups, 1-6
- Administrators, 12-1
- Agent patching directory structure, 22-1
- Agents, updating, 21-6
- aggregation and purging policies
 - See* data retention policies
- alerts
 - automated responses, 1-4
 - corrective actions, 1-4
 - notification methods, 1-7
 - notifications for, 1-7
- analyzing
 - job activity, 9-16
- archive logging
 - for Management Repository database, 20-1
- Auto Apply Templates, 7-2
- automated
 - responses to alerts, 1-4
- automated patching
 - Offline mode, 22-2
 - Online mode, 22-2
- automated patching advantages, 22-2
- Availability History Report, picture of, 14-2

B

- baseline normalized views, 1-3
- Beacons
 - introduction, 1-2

- benefits of Information Publisher, 14-1
- blackouts
 - command-line interface and, 1-5
 - controlling with emctl, 24-13
 - examples, 24-14
 - functionality of, 1-5
 - retroactive, 1-5

C

- catalog archives, 18-14
- Cell administration, 16-9
- changing Management Agent time zone, 24-15
- Checkpoint Firewall, Oracle ecosystem and, 1-2
- Cloud Control
 - starting, 24-9
 - starting all components of, 24-9
- Cloud Control Mobile, 3-21
 - acknowledge issue, 15-8
 - change views, 15-8
 - force quit, 15-9
 - iDevice, 15-1
 - Incident Manager, 15-4
 - incidents and problems, 15-6
 - iTunes App Store, 15-1
 - logging in, 15-3
 - manage issue workflow, 15-8
 - manage settings, 15-3
 - My Oracle Support, 15-6
 - requirements, 15-1
 - setup, 15-1
 - SR number, 15-6
 - touch-and-hold, 15-8
- Command Line Interface (EMCLI)
 - blackouts and, 1-5
- configuring
 - blackouts, functionality of, 1-5
 - monitoring templates, 1-6
- configuring Services
 - availability, 10-5, 10-10
 - beacons, 10-5
 - key beacons, 10-5
 - system-based, 10-5, 10-10
 - test-based service, 10-5, 10-10
- Command Line Interface, 10-34
 - creating, 10-4

- metrics
 - performance, 10-5
 - usage, 10-12
 - usage metrics, 10-5
- monitoring settings, 10-30
 - beacon overrides, 10-30
 - Collection Settings tab, 10-31
 - Data Granularity property, 10-31
 - frequency, 10-30
- monitoring templates, 10-31
 - beacons, 10-32
 - service tests, 10-32
 - service tests and beacons, 10-32
 - variables, 10-32
- performance, 10-5
- performance metrics, 10-11
 - aggregation function, 10-11
- recording transactions, 10-20
- Root Cause Analysis, 10-6, 10-11
 - Topology page, 10-6
- Service Level Rules, 10-32
 - actual service level, 10-33
 - availability, 10-33
 - business hours, 10-33
 - expected service level, 10-33
 - Information Publisher, 10-34
 - performance criteria, 10-33
 - Services Dashboard, 10-34
- service test-based availability
 - key service tests, 10-10
- service tests and beacons
 - configuring dedicated beacons, 10-10
 - configuring Web proxy, 10-10
 - selecting test type, 10-14
- system key components, 10-5
- system-based availability
 - key components, 10-11
- thresholds
 - critical, 10-5
 - warning, 10-5
- time zone, 10-5
- types of Services
 - Aggregate Service, 10-13
- configuring Software Librar
 - installation procedure
 - OMS Agent storage, 11-11
 - OMS shared file system, 11-11
 - referenced storage location, 11-13
- configuring Software Library, 11-1
 - administrators privileges, 11-4
 - installation procedure, 11-10
 - maintenance procedure, 11-29
 - deleting Software Library storage location, 11-30
 - periodic maintenance tasks, 11-30
 - re-importing Oracle owned entity files, 11-30
- overview, 11-1
- prerequisites, 11-10
- roles and Software Library privileges, 11-4
- storage, 11-6

- user roles and privileges, 11-3
- connect descriptor
 - using to identify the Management Repository database, 20-9, 20-10
- Connectors, 1-11
- controlling Oracle Management Service on UNIX, 24-5
- corrective actions
 - alerts and, 1-4
 - privileges required for, 1-4
- creating
 - administration groups, 6-9
 - custom reports, 14-2
 - report definitions, 14-2
- Critical URL Monitoring, as substitute for Management Agent, 1-2
- custom reports, 14-2
- customizing
 - notifications, 1-7
- customizing Cloud Control pages, 17-1

D

- dashboard
 - groups, 5-10
- data retention policies
 - for Application Performance Management data, 20-3
 - for other Management data, 20-3
 - modifying default, 20-3
 - of the Management Repository, 20-2
 - when targets are deleted, 20-4
- DBSNMP database user, 24-11
 - setting the password for, 24-11
- Default Templates, 7-2
- deleting targets
 - data retention policies when, 20-4
- deploying plug-ins, 18-16, 18-20
- deployment, 18-3
- deployment plug-ins, 18-18
- deployment status, 18-18, 18-20
- Disaster Recover, 28-1
- discovering targets, 18-5
- disk mirroring and stripping
 - Management Repository guideline, 20-1
- disk space management
 - controlling the size and number of log and trace files, 25-17
 - controlling the size of log and trace files, 25-18
- downloading logs, 25-7
- downloading plug-ins, 18-13, 18-14
- dropping the Management Repository, 20-8

E

- E-mail Customization, 4-11
- e-mail notifications, upper limits, 4-5
- e-mails, formats of, 1-8
- EMCLI, setting up, 21-3
- emctl

- controlling blackouts, 24-13
- listing targets on a managed host, 24-12
- starting, stopping, and checking the Management Service, 24-5
- emctl commands, 24-17
 - OMS, 24-17
- emctl config agent listtargets, 24-12
- emctl istop, 24-3
- emctl reload, 24-11
- emctl start agent, 24-2
- emctl start blackout, 24-14
- emctl start oms, 24-5
- emctl status agent, 24-2
- emctl status blackout, 24-14
- emctl status oms, 24-6
- emctl stop agent, 24-2
- emctl stop blackout, 24-14
- emctl stop oms, 24-6
- emctl upload, 24-11
- emctl.log, 25-10, 25-16
- emctl.log file, 24-24
- emctl.msg, 25-16
- emoms_pbs.log, 25-16
- emoms_pbs.trc, 25-16
- emoms.log, 25-16, 25-17
- emomslogging.properties
 - MaxBackupIndex, 25-18
 - MaxFileSize, 25-18
- emoms.trc, 25-16
- Enterprise Manager
 - blackouts, functionality of, 1-5
 - monitoring templates, 1-6
- Enterprise Manager, maintaining, 19-1
- Error Cases, OPatchauto, 23-28
- event attributes, 3-4
- Event connectors, 1-11
- Event Management, 3-1
- event severity, 3-7
- event types, 3-5
- Events, 1-4, 3-4
- Extended Network, as substitute for Management Agent, 1-2
- extensibility paradigm, 18-1

F

- Fetchlets, 12-6
- For, 20-3
- force quit
 - Cloud Control Mobile, 15-9

G

- gcagent_errors.log, 25-10
- gcagent_mdu.log, 25-10
- gcagent.log, 25-10
- generating HTML reports, 14-2
- Grid Control
 - stopping, 24-10
 - stopping all components of, 24-10

- Group Hierarchy, 6-3
- Group Members page, picture of, 5-10
- groups
 - central monitoring location, 5-8
 - dashboard, 5-10
 - description and purpose, 5-1
 - management features, 5-4
 - member targets, 5-10

H

- Health Overview, 19-2
- helpdesk connectors, 1-11
- home page
 - setting, 17-3

I

- IBM WebSphere
 - Oracle ecosystem and, 1-2
- iDevice
 - supported by Cloud Control Mobile, 15-1
- Incident Attributes, 3-8
- incident creation, 3-12
- Incident Management, 3-1
- Incident Manager, 1-9
 - Cloud Control Mobile, 15-4
- Incident Priority, 3-8
- Incidents, 3-7
- incidents and problems
 - Cloud Control Mobile, 15-6
- incidents, working with, 3-36
- Information Publisher
 - Create Like function, 14-2
 - generating HTML reports, 14-2
 - overview of, 14-1
 - predefined reports, 5-11
 - report
 - definitions, 14-2
 - elements, 14-3
 - reporting framework, 14-1
 - sharing reports, 14-4
 - viewing reports, 14-4
- informational updates, 21-5
- istop
 - emctl command, 24-3
- iTunes App Store
 - Cloud Control Mobile, 15-1

J

- Job Activity page, 9-1
- jobs
 - analyzing job activity, 9-16
 - definition of, 9-1
 - Job Activity page, 9-1
 - job executions, 9-2
 - job runs, 9-2
 - multitask, 9-15
 - notification rules for e-mail, 9-9
 - operations on runs and executions, 9-2

- privileges for sharing job responsibilities, 9-4
- purpose of, 9-1

L

- load balancer switches
 - BIG-IP, Oracle ecosystem and, 1-2
- Load Balancers, 28-5
- local store, 21-3
- log files
 - controlling the size and number of, 25-17
 - locating and configuring, 25-1
 - locating Management Agent, 25-11
 - locating Management Service, 25-16
 - Management Agent, 25-9
 - Oracle Management Service, 25-15
 - searching, 25-5
- log4j.appender.emlogAppender.
 - MaxBackupIndex, 25-18
- log4j.appender.emlogAppender. MaxFileSize, 25-18
- log4j.appender.emtrcAppender.
 - MaxBackupIndex, 25-18
- log4j.appender.emtrcAppender. MaxFileSize, 25-18
- logging in to Cloud Control Mobile, 15-3
- LVM (Logical Volume Manager), 20-1

M

- Management Agent, 25-10
 - additional Management Agent commands, 24-11
 - checking the status on UNIX, 24-2
 - checking the status on Windows, 24-3
 - Critical URL Monitoring as substitute, 1-2
 - Extended Network as substitute, 1-2
 - purpose of, 1-1
 - starting and stopping on UNIX, 24-1
 - starting and stopping on Windows, 24-3
- Management Agent logs
 - setting log levels, 25-11, 25-12, 25-13
 - setting trace levels, 25-15
- Management Agent time zone, 24-15
- Management Information Base (MIB), 4-42
 - definition, 4-42
 - MIB variable descriptions, 4-43
- Management Repository
 - introduction of, 1-2
 - See* Oracle Management Repository
- Management Repository Deployment
 - Guideline, 20-1
- Management Service
 - starting and stopping on Windows systems, 24-6
- managing
 - groups, 5-4
- managing Cloud Control Mobile sites, 15-3
- managing logs, 25-1
- MaxBackupIndex
 - property in emomslogging.properties, 25-18
- MaxFileSize
 - property in emomslogging.properties, 25-18
- metric

- thresholds, 1-3
- Metric Baselines, 1-3
- Metric Columns, Delta, 8-12
- Metric Columns, Rate, 8-12
- Metric Extension, 1-4
- Metric Extension Lifecycle, 8-3
- Metric Extension, creating, 8-9
- Metric Extension, deleting, 8-16
- Metric Extension, editing, 8-14
- Metric Extension, exporting, 8-16
- Metric Extension, importing, 8-15
- Metric extensions, 8-1
- metric extensions, 12-6
- Metric Extensions, administrator privileges, 8-6
- Metric Extensions, deploying, 8-16
- Metric Extensions, updating older versions, 8-17
- metrics
 - threshold values, 1-3
 - thresholds, 1-3
- MGMT_ADMIN.DISABLE_METRIC_DELETION, 20-5
- MGMT_ADMIN.ENABLE_METRIC_DELETION, 20-5
- MGMT_METRICS_1DAY table, 20-4
- MGMT_METRICS_1HOUR table, 20-4
- MGMT_METRICS_RAW table, 20-4
- MIB

- <italic>See Management Information Base (MIB)

- Migrating, 11-30
- modes of patching, 22-2
- monitoring
 - alerts as they occur, 5-10
 - basics of, 1-1
 - templates
 - function of, 1-6
- monitoring credentials
 - defined, 24-11
 - setting, 24-11
- Monitoring Overview, 1-1
- Monitoring Template, creating, 7-3
- Monitoring Template, definition of, 7-2
- Monitoring Template, editing, 7-4
- Monitoring Template, retention period, 7-10
- Monitoring Templates, 7-1
- Monitoring Templates, applying to targets, 7-4
- Monitoring Templates, compare with targets, 7-7
- Monitoring Templates, exporting/importing, 7-9
- multitask jobs, 9-15
- My Oracle Support, OMS Patches, 23-8
- My Oracle Support, OPatchauto, 23-9

N

- NetApp Filers
 - Oracle ecosystem and, 1-2
- new product announcements, 21-5
- Notification, 12-4
- notification methods
 - based on a PL/SQL Procedure, 4-24
 - based on an SNMP trap, 4-38

- based on operating system commands, 4-16
 - definition, 4-15
- notification rules
 - custom, 4-8
 - definition, 4-7
 - out-of-box, 4-8
 - out-of-the-box notification rules, 4-6
 - subscribing to, 4-7
- notification schedules, 4-6
- notification system
 - e-mail errors, 4-75
 - errors, 4-73
 - trace messages, 4-73
- notifications
 - alerts, 1-7
 - customizing, 1-7
 - defining multiple mail servers, 4-3
 - for jobs, 9-9
 - long e-mail notifications, 4-5
 - mail server settings, 4-2
 - management information base (MIB), 4-42
 - methods, 1-7
 - notification method, 1-7
 - notification schedules, 4-6
 - sample Operating System command script, 4-18
 - setting up, 4-2
 - short email notifications, 4-6

O

- Offline mode, 22-2
- OMS Configurations, OPatchauto, 23-1
- OMS Configurations, OPatchauto, 23-1
- OMS Core Patches, 23-9
- OMS patching, 23-1
- OMS Plug-in Patches, 23-11
- Online mode, 22-2
- OPatch, latest version, 23-7
- OPatchauto, 23-1, 23-3
- opatchauto apply, 23-12, 23-16
- OPatchauto Commands, 23-15
- OPatchauto help, 23-15
- OPatchauto Log, 23-23
- opatchauto lspatches, 23-13, 23-20
- OPatchauto Parameters, 23-3
- OPatchauto Property File, 23-4
- opatchauto rollback, 23-13, 23-18
- OPatchauto Troubleshooting, 23-23
- opatchauto version, 23-14, 23-21
- OPatchauto, MOS, 23-8
- OPatchauto, prerequisites, 23-4
- Operating System command
 - sample notification method for, 4-17
 - sample script, 4-18
- Operating System scripts, 4-16
- Oracle
 - ecosystem, 1-2
- Oracle Enterprise Manager
 - log files, 25-1
- Oracle Enterprise Manager Cloud Control, 24-9

- Oracle HTTP Server logs, 25-19
- Oracle Management Agent
 - about log and trace files, 25-10
 - location of log and trace files, 25-11
 - log and trace files, 25-9
 - starting and stopping, 24-1
- Oracle Management Repository
 - data retention policies, 20-2
 - dropping, 20-8
 - identifying with a connect descriptor, 20-9, 20-10
 - recreating, 20-8, 20-9
 - reloading data, 24-11
 - starting the Management Repository
 - database, 24-9
 - troubleshooting, 20-11
 - uploading data, 24-11
- Oracle Management Service, 24-5
 - about the log and trace files, 25-16
 - location the log and trace files, 25-16
 - log and trace files, 25-15
 - modifying monitoring credentials, 24-11
 - starting, stopping, and checking, 24-5
- Oracle Management Service logs, 25-16, 25-17
- Oracle Management Service trace files, 25-18
- Oracle Process Management and Notification (OPMN)
 - using to start and stop the Management Service, 24-6
- Oracle WebLogic Server logs, 25-19
- oraEM4JobAlertTable, A-32
- oraEMNGEvent, A-1
- orAgentTraps, A-22
- OS scripts
 - See* Operating System scripts
- OUI Inventory Configurations, 23-2
- Out-of-Box Monitoring, 1-1
- out-of-box reports, 14-2

P

- Patch Format, 23-3
- Patch Search region, 22-3, 22-4
- Patches and Updates, 22-2
 - Agent patching
 - Add All To Plan, 22-7
 - Create Plan, 22-5, 22-6
 - Patch Plans, 22-7
 - Deployment Options, 22-7
 - Patches page, 22-7
 - Plan Information, 22-7
 - Review and Deploy page, 22-9
 - Validation page, 22-9
 - View Plan, 22-6
 - warnings
 - Null Platform, 22-7
- Patch Search region, 22-3
 - Advanced Search, 22-4
 - Basic Search, 22-4
- patching Enterprise Manager

- Management Agents
 - automated patching, 22-2
 - accessing Patches and Updates, 22-2
 - applying Agent patches, 22-5
 - searching Patches, 22-3
 - validating applied agent patches, 22-10
 - verifying the applied agent patches, 22-10
 - viewing Patch recommendations, 22-3
 - manual patching, 22-13
 - overview, 22-1
 - patching Enterprise Manager core components, 22-1
 - patching Management Agents, 22-1
 - patching OMS, 22-1, 23-1
 - patching Repository, 22-1
- Performance Metrics
 - Beacon Aggregation Function
 - Average, 10-11, 10-12
 - Minimum, 10-11
 - Sum, 10-11
- performance metrics
 - Beacon Aggregation Function
 - maximum value, 10-11
 - minimum value, 10-12
 - sum of values, 10-12
 - System Aggregation Function
 - maximum value, 10-12
- personalize, 17-1
- personalizing Cloud Control pages, 17-1
- Planning Host Names, 28-4
- planning outage periods, blackouts, 1-5
- PL/SQL procedures, 4-16
 - while creating notification methods, 4-24
- plug-in archives, 18-15
- plug-in homes, 18-26
- plug-in id, 18-11
- plug-in manager, 18-1, 18-6, 18-7, 18-8
- plug-ins, 18-1, 18-2, 18-3, 18-10, 18-11, 18-13, 18-14, 18-27
- privileges
 - for corrective actions, 1-4
 - for sharing job responsibilities, 9-4
- Problem Management, 3-2
- Problems, 3-12
- ProcessManager
 - service used to control the Management Service on Windows systems, 24-6
- purging policies
 - See data retention policies

R

- RAID-capable disk
 - Management Repository guideline, 20-1
- Receivelets, 12-5

- reevaluating metric collections, 24-15
- Repeat Notifications, 4-14
- Repeat Notifications for Rules, 4-14
- RepManager script, 20-9
- reports
 - creating custom reports, 14-2
 - custom, 14-2
 - definitions, Information Publisher, 14-2
 - e-mailing, 14-4
 - generating HTML report, 14-2
 - Information Publisher, 14-1
 - out-of-box, Information Publisher, 14-2
 - predefined, 5-11
 - predefined report definitions, 14-2
 - report elements, 14-3
 - scheduling, 14-3
 - sharing, 14-4
 - storing and purging, 14-4
 - viewing, 14-4
- requirements
 - Cloud Control Mobile, 15-1
- retroactive blackouts, 1-5
- Root Cause Analysis
 - mode
 - automatic, 10-7
 - manual, 10-7
- Rule Actions, 3-19
- Rule Criteria, 3-17
- Rule Set Types, 3-15
- Rule Set, developing, 3-23
- rule sets, 3-13
- rule sets, out-of-box, 3-14
- Rule Sets, setting up, 3-30
- Rules, 1-7

S

- scheduled maintenance with blackouts, 1-5
- scheduling
 - reports, 14-3
 - reports, flexibility, 14-3
- searching logs, 25-5
- Self Update feature
 - setting up, 21-1
 - using, 21-1
- self-update, 18-3
- Server Connection Hung
 - error while creating the repository, 20-11
- Service Tests and Beacons
 - Tests
 - DNS, 10-14
 - FTP, 10-14
 - SOAP, 10-14
 - Web Transaction, 10-14
- Services control panel
 - using to start the Management Service, 24-6
- setting
 - metric threshold values, 1-3
- setting your home page, 17-3
- sharing reports, 14-4

- SNMP, 12-1
- SNMP Trap, 12-4
- SNMP Traps, 12-5
- SNMP traps, 4-16, 4-38
 - about, A-1
 - sample, 4-40
- Software Library, 21-3
 - designers, 11-4
 - Operators, 11-4
 - Super Administrators, 11-4
 - users, 11-4
- Software Library Administration, 11-7
 - referenced file locations, 11-9
 - Agent storage, 11-10
 - http storage, 11-9
 - NFS storage, 11-9
 - upload file locations, 11-8
- Software Library console, 11-2
- Software Library referenced locations, 11-6
- Software Library storage, 11-1
- Software library upload locations, 11-6
- Standby OMS System, patching, 23-14
- Standby OMS, Patching, 23-22
- status codes, corrective actions, 4-48
- system errors, notification, 4-73
- System Patch, 23-3
- system patch, 23-3

T

- target
 - definition of, 1-1
- target monitoring credentials
 - defined, 24-11
 - setting, 24-11
- target properties, 1-7
- targets
 - listing targets on a managed host, 24-12
- Template Collections, 1-6, 6-9
 - with administration groups, 6-18
- template collections, privileges, 6-17
- thresholds
 - definition of, 1-3
 - for metrics, 1-3
- time zone, 24-15
- touch-and-hold
 - Cloud Control Mobile, 15-8
- trace files
 - controlling the contents of Management Service, 25-18
 - controlling the size and number of, 25-17
 - locating Management Agent, 25-11
 - locating Management Service, 25-16
 - Management Agent, 25-9
 - Oracle Management Service, 25-15
- troubleshooting
 - general techniques while creating the Management Repository, 20-11
 - Management Agent startup errors, 24-4
 - Management Service, 24-6

- Management Service startup errors, 24-6
- notifications, 4-72
 - while creating the Management Repository, 20-10
- troubleshooting Management Agent, 24-4
- Troubleshooting Service Tests, 10-38
 - Forms Transactions, 10-38

U

- undeploying plug-ins, 18-25
- updates
 - applying in offline mode, 21-5
 - applying in online mode, 21-4
- updating Cloud Control, 21-1
- upgrading plug-ins, 18-19, 18-21
- Usage metrics
 - Aggregation Function
 - average value, 10-12
 - maximum value, 10-12
 - minimum value, 10-12
 - sum of values, 10-13

V

- viewing
 - reports, 14-4
- viewing logs, 25-3
- VPN
 - Cloud Control Mobile requirement, 15-1

W

- Web Application
 - Source
 - Step, 10-11
 - Step Group, 10-11
 - Transaction, 10-11

