

Oracle® Enterprise Manager

Administration

11g Release 1 (11.1.0.1)

E16790-04

March 2013

Oracle Enterprise Manager Administration, 11g Release 1 (11.1.0.1)

E16790-04

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	xvii
Audience	xvii
Documentation Accessibility	xvii
Related Documents	xvii
Conventions	xviii
1 Monitoring	
1.1 Systems Monitoring: Breadth and Depth	1-1
1.2 Monitoring Basics	1-2
1.2.1 Out-of-Box Monitoring	1-2
1.2.1.1 Metric Thresholds	1-3
1.2.2 Alerts	1-4
1.2.3 Notifications	1-5
1.2.3.1 Customizing Notifications	1-5
1.2.4 Corrective Actions	1-6
1.2.5 Blackouts	1-6
1.3 Monitoring Templates	1-7
1.4 User-Defined Metrics	1-7
1.5 Accessing Monitoring Information	1-8
2 Enterprise Manager Security	
2.1 About Oracle Enterprise Manager Security	2-1
2.2 Enterprise Manager Authentication	2-2
2.2.1 Repository-Based Authentication	2-2
2.2.2 Single Sign-On Based Authentication	2-3
2.2.2.1 Registering Enterprise Manager as a Partner Application	2-4
2.2.2.2 Removing Single Sign-On Configuration	2-5
2.2.2.3 Registering Single Sign-On Users as Enterprise Manager Administrators	2-6
2.2.2.4 Grid Control as a Single Sign-On Partner Application	2-9
2.2.2.5 Bypassing the Single Sign-On Logon Page	2-9
2.2.3 Enterprise User Security Based Authentication	2-9
2.2.3.1 Registering Enterprise Users as Enterprise Manager Users	2-10
2.3 Enterprise Manager Authorization	2-11
2.3.1 Classes of Users	2-11
2.3.2 Privileges and Roles	2-12

2.3.2.1	Granting Privileges.....	2-13
2.4	Configuring Security for Grid Control	2-16
2.4.1	About Enterprise Manager Framework Security	2-16
2.4.2	Overview of the Steps Required to Enable Enterprise Manager Framework Security	2-17
2.4.3	Enabling Security for the Oracle Management Service.....	2-18
2.4.3.1	Checking the Security Status	2-20
2.4.3.2	Creating a New Certificate Authority	2-21
2.4.3.3	Viewing the Security Status and OMS Port Information	2-21
2.4.3.4	Configuring Transparent Layer Security	2-22
2.4.4	Enabling Security for the Oracle Management Agent	2-23
2.4.5	Enabling Security with Multiple Management Service Installations.....	2-25
2.4.6	Restricting HTTP Access to the Management Service	2-26
2.4.7	Managing Agent Registration Passwords.....	2-27
2.4.7.1	Using the Grid Control Console to Manage Agent Registration Passwords....	2-28
2.4.7.2	Using emctl to Add a New Agent Registration Password	2-28
2.4.8	Enabling Security with a Server Load Balancer	2-29
2.4.9	Enabling Security for the Management Repository Database	2-29
2.4.9.1	About Oracle Advanced Security and the sqlnet.ora Configuration File	2-30
2.4.9.2	Configuring the Management Service to Connect to a Secure Management Repository Database 2-30	
2.4.9.3	Enabling Oracle Advanced Security for the Management Repository.....	2-32
2.4.9.4	Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database 2-32	
2.4.10	Configuring Third Party Certificates	2-32
2.4.10.1	Configuring Third Party Certificate for HTTPS Upload Virtual Host	2-33
2.4.10.2	Configuring Third Party Certificate for HTTPS WebTier Virtual Host	2-33
2.5	Accessing Managed Targets	2-34
2.5.1	Credential Subsystem.....	2-34
2.5.1.1	Managing Credentials Using EMCLI	2-35
2.5.2	Pluggable Authentication Modules (PAM) Support for Hosts.....	2-36
2.5.2.1	Configuring PAM for RHEL4 Users	2-36
2.5.2.2	Configuring PAM for AIX Users.....	2-36
2.5.3	Sudo and PowerBroker Support.....	2-37
2.5.3.1	Creating a Privilege Delegation Setting	2-38
2.6	Cryptographic Support	2-39
2.6.1	Configuring the emkey	2-39
2.6.2	emctl Commands	2-40
2.6.2.1	emctl status emkey	2-40
2.6.2.2	emctl config emkey -copy_to_credstore.....	2-41
2.6.2.3	emctl config emkey -copy_to_repos	2-41
2.6.2.4	emctl config emkey -copy_to_file_from_credstore.....	2-41
2.6.2.5	emctl config emkey -copy_to_file_from_repos	2-42
2.6.2.6	emctl config emkey -copy_to_credstore_from_file.....	2-42
2.6.2.7	emctl config emkey -copy_to_repos_from_file	2-42
2.6.2.8	emctl config emkey -remove_from_repos	2-42
2.6.3	Install and Upgrade Scenarios	2-43
2.6.3.1	Installing the Management Repository	2-43

2.6.3.2	Installing the First Oracle Management Service	2-43
2.6.3.3	Upgrading from 10.2 to 11.1	2-43
2.6.3.4	Recreating the Management Repository	2-43
2.7	Setting Up the Auditing System for Enterprise Manager	2-43
2.7.1	Configuring the Enterprise Manager Audit System.....	2-44
2.7.2	Configuring the Audit Data Export Service	2-44
2.7.3	Updating the Audit Settings	2-44
2.7.4	Searching the Audit Data	2-45
2.7.5	List of Operations Audited.....	2-47
2.8	Additional Security Considerations.....	2-49
2.8.1	Changing the SYSMAN and MGMT_VIEW Passwords.....	2-49
2.8.1.1	Changing the SYSMAN User Password	2-49
2.8.1.2	Changing the MGMT_VIEW User Password.....	2-49
2.8.2	Responding to Browser-Specific Security Certificate Alerts	2-50
2.8.2.1	Responding to the Internet Explorer Security Alert Dialog Box	2-50
2.8.2.2	Responding to the Netscape Navigator New Site Certificate Dialog Box	2-51
2.8.2.3	Preventing the Display of the Internet Explorer Security Information Dialog Box....	2-52
2.8.3	Configuring Beacons to Monitor Web Applications Over HTTPS.....	2-53
2.8.4	Using ORACLE_HOME Credentials	2-54
2.8.5	Patching Oracle Homes When the User is Locked	2-56
2.8.6	Cloning Oracle Homes.....	2-57

3 Notifications

3.1	Setting Up Notifications.....	3-1
3.1.1	Setting Up a Mail Server for Notifications.....	3-1
3.1.1.1	Setting Up Repeat Notifications	3-4
3.1.2	Setting Up E-mail for Yourself.....	3-5
3.1.2.1	Defining E-mail Addresses	3-5
3.1.2.2	Setting Up a Notification Schedule.....	3-7
3.1.2.3	Subscribe to Receive E-mail for Notification Rules	3-7
3.1.3	Setting Up E-mail for Other Administrators	3-12
3.1.4	E-mail Customization.....	3-13
3.1.4.1	E-mail Customization Reference	3-14
3.2	Extending Notification Beyond E-mail.....	3-18
3.2.1	Custom Notification Methods Using Scripts and SNMP Traps	3-19
3.2.1.1	Adding a Notification Method based on an OS Command or Script.....	3-19
3.2.1.2	Adding a Notification Method Based on a PL/SQL Procedure	3-23
3.2.1.3	Adding a Notification Method Based on an SNMP Trap.....	3-26
3.3	Passing Corrective Action Status Change Information.....	3-28
3.3.1	Passing Corrective Action Execution Status to an OS Command or Script.....	3-29
3.3.2	Passing Corrective Action Execution Status to a PLSQL Procedure.....	3-30
3.4	Passing Job Execution Status Information.....	3-32
3.4.1	Passing Job Execution Status to a PLSQL Procedure	3-33
3.4.2	Passing Job Execution Status to an OS Command or Script.....	3-34
3.5	Passing User-Defined Target Properties to Notification Methods	3-35
3.6	Assigning Methods to Rules.....	3-35

3.7	Assigning Rules to Methods.....	3-36
3.8	Notification Coverage	3-37
3.9	Management Information Base (MIB).....	3-37
3.9.1	About MIBs.....	3-37
3.9.2	Reading the MIB Variable Descriptions	3-38
3.9.2.1	Variable Name	3-38
3.9.2.2	MIB Definition.....	3-39
3.10	Troubleshooting Notifications	3-45
3.10.1	General Setup	3-45
3.10.2	Notification System Errors	3-45
3.10.3	Notification System Trace Messages.....	3-45
3.10.4	E-mail Errors.....	3-47
3.10.5	OS Command Errors	3-47
3.10.6	SNMP Trap Errors	3-47
3.10.7	PL/SQL Errors	3-47

4 User-Defined Metrics

4.1	Extending Monitoring Capability.....	4-1
4.2	Creating OS-Based User-Defined Metrics	4-2
4.2.1	Create Your OS Monitoring Script.....	4-2
4.2.1.1	Code to check the status of monitored objects	4-2
4.2.1.2	Code to return script results to Enterprise Manager.....	4-2
4.2.1.3	Script Runtime Environment	4-4
4.2.2	Register the Script as a User-Defined Metric.....	4-5
4.2.3	OS-Based User-Defined Metric Example	4-8
4.3	Creating a SQL-Based User-Defined Metric	4-10
4.3.1	SQL-Based User-Defined Metric Examples	4-14
4.3.1.1	Example 1: Query Returning Tablespace Name and Percent Used	4-14
4.3.1.2	Example 2: Query Returning Segment Name/Type and Extent Count.....	4-15
4.3.1.3	Example 3: Embed a Long SQL statement in a PL/SQL Routine	4-15
4.4	Notifications, Corrective Actions, and Monitoring Templates	4-17
4.4.1	Getting Notifications for User-Defined Metrics	4-17
4.4.2	Setting Corrective Actions for User-Defined Metrics.....	4-20
4.4.3	Deploying User-Defined Metrics Across Many Targets Using Monitoring Templates....	4-20
4.4.4	Deleting User-Defined Metrics Across Many Targets Using Monitoring Templates	4-22
4.5	Changing User-Defined Metric Credentials	4-23

5 Group Management

5.1	Introduction to Groups	5-1
5.2	Managing Groups	5-2
5.2.1	Group Home Page	5-2
5.2.2	Group Charts Page	5-4
5.2.3	Group Administration Page.....	5-4
5.2.4	Group Members Page	5-5
5.2.5	System Dashboard.....	5-5

5.3	Out-of-Box Reports	5-6
5.4	Redundancy Groups.....	5-7
5.5	Privilege Propagating Groups.....	5-8
5.5.1	Creating Privilege Propagating Groups.....	5-9
5.5.2	Using the Group Administration Privilege	5-9
5.5.3	Adding Members to Privilege Propagating Groups	5-10
5.5.4	Converting Conventional Groups to Privilege Propagating Groups	5-10

6 Job System

6.1	Job System Purpose and Overview	6-1
6.1.1	What Are Job Executions and Job Runs?.....	6-2
6.1.2	Operations on Job Executions and Job Runs	6-2
6.2	Preliminary Considerations.....	6-3
6.2.1	Using Pre-defined Tasks.....	6-3
6.2.2	Creating Scripts.....	6-3
6.2.3	Sharing Job Responsibilities.....	6-4
6.2.4	Jobs and Groups.....	6-4
6.3	Creating Jobs.....	6-4
6.3.1	Selecting a Job Type.....	6-4
6.3.2	Creating an OS Command Job.....	6-5
6.3.2.1	Specifying a Single Operation.....	6-10
6.3.2.2	Specifying a Script	6-10
6.3.2.3	Access Level Rules.....	6-11
6.3.3	Creating a SQL Script Job	6-12
6.3.3.1	Specifying Targets	6-12
6.3.3.2	Options for the Parameters Page.....	6-12
6.3.3.3	Specifying Host and Database Credentials.....	6-13
6.3.3.4	Returning Error Codes from SQL Script Jobs.....	6-13
6.3.4	Creating a Multi-task Job.....	6-14
6.3.4.1	Job Capabilities	6-14
6.3.4.2	Specifying Targets for a Multi-task Job	6-15
6.3.4.3	Adding Tasks to the Job.....	6-15
6.4	Analyzing Job Activity.....	6-15

7 Starting and Stopping Enterprise Manager Components

7.1	Controlling the Oracle Management Agent.....	7-1
7.1.1	Starting, Stopping, and Checking the Status of the Management Agent on UNIX...	7-1
7.1.2	Starting and Stopping the Management Agent on Windows.....	7-2
7.1.3	Checking the Status of the Management Agent on Windows	7-4
7.2	Controlling the Oracle Management Service.....	7-4
7.2.1	Controlling the Management Service on UNIX	7-4
7.2.1.1	Using emctl to Start, Stop, and Check the Status of the Oracle Management Service	7-4
7.2.2	Controlling the Management Service on Windows.....	7-5
7.3	Controlling Fusion Middleware Control.....	7-5
7.4	Controlling the Database Control on UNIX.....	7-6

7.4.1	Starting the Database Control on UNIX.....	7-6
7.4.2	Stopping the Database Control on UNIX.....	7-6
7.4.3	Starting and Stopping the Database Control on Windows	7-6
7.5	Guidelines for Starting Multiple Enterprise Manager Components on a Single Host	7-7
7.6	Starting and Stopping Oracle Enterprise Manager 11g Grid Control	7-7
7.6.1	Starting Grid Control and All Its Components	7-7
7.6.2	Stopping Grid Control and All Its Components	7-8
7.7	Additional Management Agent Commands	7-9
7.7.1	Uploading and Reloading Data to the Management Repository	7-9
7.7.2	Specifying New Target Monitoring Credentials	7-10
7.7.2.1	Using the Grid Control Console to Modify the Monitoring Credentials	7-11
7.7.2.2	Using the Enterprise Manager Command Line to Modify the Monitoring Credentials	7-11
7.7.3	Listing the Targets on a Managed Host.....	7-11
7.7.4	Controlling Blackouts.....	7-14
7.7.5	Changing the Management Agent Time Zone.....	7-16
7.7.6	Reevaluating Metric Collections.....	7-17
7.8	emctl Commands	7-19
7.9	Using emctl.log File	7-32

8 Information Publisher

8.1	About Information Publisher	8-1
8.2	Out-of-Box Report Definitions	8-2
8.3	Custom Reports.....	8-4
8.3.1	Creating Custom Reports	8-4
8.3.2	Report Parameters	8-4
8.3.3	Report Elements	8-5
8.4	Scheduling Reports.....	8-6
8.4.1	Flexible Schedules.....	8-6
8.4.2	Storing and Purging Report Copies	8-6
8.4.3	E-mailing Reports	8-7
8.5	Sharing Reports.....	8-7

9 Sizing Your Enterprise Manager Deployment

9.1	Oracle Enterprise Manager Grid Control Architecture Overview	9-1
9.2	Enterprise Manager Grid Control Sizing and Performance Methodology	9-2
9.2.1	Step 1: Choosing a Starting Platform Grid Control Deployment	9-3
9.2.1.1	Network Topology Considerations	9-4
9.2.2	Step 2: Periodically Evaluate the Vital Signs of Your Site	9-5
9.2.3	Step 3: Use DBA and Enterprise Manager Tasks To Eliminate Bottlenecks Through Housekeeping	9-10
9.2.3.1	Online Weekly Tasks.....	9-10
9.2.3.2	Offline Monthly Tasks	9-10
9.2.4	Step 4: Eliminate Bottlenecks Through Tuning.....	9-11
9.2.4.1	High CPU Utilization.....	9-11
9.2.4.2	Loader Vital Signs	9-12
9.2.4.3	Rollup Vital Signs	9-14

9.2.4.4	Rollup Process.....	9-14
9.2.4.5	Job, Notification, and Alert Vital Signs	9-16
9.2.4.6	I/O Vital Signs	9-16
9.2.4.7	The Oracle Enterprise Manager Performance Page.....	9-17
9.2.5	Step 5: Extrapolating Linearly Into the Future for Sizing Requirements	9-18
9.3	Oracle Enterprise Manager Backup, Recovery, and Disaster Recovery Considerations	9-18
9.3.1	Best Practices for Backup	9-18
9.3.2	Best Practices for Recovery.....	9-19
9.3.2.1	Recovering the Management Repository	9-19
9.3.2.2	Recovering the Oracle Management Service	9-20
9.3.2.3	Recovering the Oracle Management Agent.....	9-20
9.3.2.4	Preventing Data Loss and Down Time While Switching From a Non-shared File System to a Shared File System	9-21
9.3.3	Best Practice for Disaster Recovery (DR)	9-21
9.3.3.1	Management Repository	9-21
9.3.3.2	Oracle Management Service	9-22
9.3.3.3	Management Agent.....	9-22

10 Maintaining and Troubleshooting the Management Repository

10.1	Management Repository Deployment Guidelines	10-1
10.2	Management Repository Data Retention Policies.....	10-2
10.2.1	Management Repository Default Aggregation and Purging Policies.....	10-2
10.2.2	Management Repository Default Aggregation and Purging Policies for Other Management Data	10-3
10.2.3	Modifying the Default Aggregation and Purging Policies.....	10-3
10.2.4	Modifying Data Retention Policies When Targets Are Deleted	10-4
10.2.5	How to Modify the Retention Period of Job History	10-5
10.3	Changing the SYSMAN Password	10-6
10.3.1	Overview of the MGMT_VIEW User.....	10-8
10.4	Dropping and Recreating the Management Repository	10-8
10.4.1	Dropping the Management Repository.....	10-8
10.4.2	Recreating the Management Repository	10-9
10.4.2.1	Using the RepManager Script to Create the Management Repository.....	10-9
10.4.2.2	Using a Connect Descriptor to Identify the Management Repository Database	10-10
10.5	Troubleshooting Management Repository Creation Errors	10-10
10.5.1	Package Body Does Not Exist Error While Creating the Management Repository.....	10-11
10.5.2	Server Connection Hung Error While Creating the Management Repository.....	10-11
10.5.3	General Troubleshooting Techniques for Creating the Management Repository	10-11
10.6	Cross Platform Enterprise Manager Repository Migration.....	10-12
10.6.1	Common Prerequisites.....	10-12
10.6.2	Methodologies.....	10-13
10.6.2.1	Cross Platform Transportable Tablespaces.....	10-13
10.6.2.2	Data Pump	10-16
10.6.2.3	Export/Import	10-18

10.6.3	Post Migration Verification	10-20
10.7	Improving the Login Performance of the Console Home Page	10-20

11 Setting Up the JVM Diagnostics Manager

11.1	Deploying the JVM Diagnostics Manager	11-1
11.1.1	Prerequisites	11-1
11.1.2	Running the DeployAD4Jmanager.sh Script	11-1
11.1.3	Deploying the JVM Diagnostics Agent	11-5
11.1.4	Creating a Special Repository User for Loading JVM Diagnostics Heaps	11-7
11.1.5	JVM Diagnostics Architecture	11-8

12 Setting Up Request Monitoring

12.1	Prerequisites	12-1
12.2	Request Monitoring Architecture	12-1
12.3	Setting Up Request Monitoring	12-2

13 Managing Compliance

13.1	Compliance Overview	13-1
13.2	Compliance Management	13-1
13.2.1	Accessing Compliance Management Pages in Grid Control	13-2
13.2.2	Investigating Policy Violations and Policy Group Evaluation Results	13-2
13.2.3	Assessing Security	13-3
13.2.4	Viewing Policy Violations Results	13-3
13.2.5	Policy Violations Reports	13-4
13.3	Setting Up Compliance Evaluations	13-4
13.3.1	Scheduling an Evaluation	13-4
13.3.2	Viewing Policy Group Evaluation Results	13-4
13.3.3	Out-of-Box Policies and Policy Groups	13-4
13.3.4	Customizing Policies	13-4
13.3.4.1	Defining Corrective Actions	13-5
13.3.4.2	Using Templates for Monitoring	13-5
13.4	Policy Groups Provided by Oracle	13-5
13.4.1	Secure Configuration for Oracle Database	13-5
13.4.2	Secure Configuration for Oracle Real Application Cluster	13-6
13.4.3	Secure Configuration for Oracle Listener	13-6

14 Configuring Services

14.1	Summary of Service Management Tasks	14-1
14.2	Setting up the System	14-3
14.3	Creating a Service	14-4
14.4	Configuring a Service	14-5
14.4.1	Availability Definition	14-6
14.4.2	Performance Metrics	14-7
14.4.3	Usage Metrics	14-8
14.4.4	Business Metrics	14-9
14.4.5	Service Tests and Beacons	14-9

14.4.5.1	Configuring the Beacons	14-10
14.4.5.2	Configuring Windows Beacons for Web Transaction (Browser) Playback	14-11
14.4.6	Root Cause Analysis Configuration.....	14-13
14.4.6.1	Getting the Most From Root Cause Analysis	14-14
14.5	Recording Web Transactions.....	14-14
14.6	Monitoring Settings	14-15
14.7	Configuring Aggregate Services.....	14-15
14.8	Configuring End-User Performance Monitoring	14-16
14.8.1	Configuring End-User Performance Monitoring Using Oracle HTTP Server Based on Apache 2.0 or Apache HTTP Server 2.0	14-17
14.8.1.1	Setting up the Third Party Apache Server.....	14-19
14.8.2	Configuring End-User Performance Monitoring Using Oracle Application Server Web Cache	14-19
14.8.2.1	Configuring Oracle Application Server Web Cache 10.1.2	14-20
14.8.2.2	Configuring Oracle Application Server Web Cache 9.0.4	14-21
14.8.2.3	Configuring End-User Performance Monitoring Using Earlier Versions of Oracle Application Server Web Cache	14-22
14.8.2.4	Configuring End-User Performance Monitoring Using Standalone Oracle Application Server Web Cache	14-25
14.8.3	Configuring End-User Performance Monitoring for Web Page Extensions.....	14-26
14.8.4	Configuring End-User Performance Monitoring for Web Pages Having the Same URI..	14-27
14.8.5	Starting and Stopping End-User Performance Monitoring.....	14-28
14.8.6	Verifying and Troubleshooting End-User Performance Monitoring.....	14-29
14.8.7	Enabling End-User Performance Monitoring for Third-Party Application Servers.....	14-30
14.9	Managing Forms Applications	14-31
14.9.1	Recording and Monitoring Forms Transactions	14-32
14.9.1.1	Setting the Permissions of the .java.policy File	14-32
14.9.1.2	Using a Trusted Enterprise Manager Certificate	14-33
14.9.1.3	Adding a Forms Certificate to the Enterprise Manager Agent.....	14-34
14.9.1.4	Configuring the Forms Server	14-34
14.9.1.5	Installing the Transaction Recorder to Record and Play Back Forms Transactions ...	14-35
14.9.2	Monitoring the End-User Performance of Forms Applications.....	14-35
14.9.2.1	Configuring the Forms Server for End-User Performance Monitoring	14-36
14.9.2.2	Configuring the OracleAS Web Cache	14-37
14.9.2.3	Configuring the Oracle HTTP Server / Apache HTTP Server	14-38
14.9.2.4	Starting and Stopping End-User Performance Monitoring.....	14-39
14.10	Configuring OC4J for Request Performance Diagnostics	14-40
14.10.1	Selecting OC4J Targets for Request Performance Diagnostics	14-40
14.10.2	Configuring Interactive Transaction Tracing	14-41
14.10.3	Configuring OC4J Tracing for Request Performance Data	14-42
14.10.4	Additional Configuration for Monitoring UIX Applications.....	14-43
14.11	Setting Up Monitoring Templates	14-43
14.11.1	Configuring Service Tests and Beacons.....	14-44
14.12	Configuring Service Levels.....	14-44
14.12.1	Defining Service Level Rules	14-45

14.12.2	Viewing Service Level Details	14-46
14.13	Configuring a Service Using the Command Line Interface.....	14-46
14.14	Troubleshooting Service Tests	14-46
14.14.1	Verifying and Troubleshooting Forms Transactions.....	14-47
14.14.1.1	Troubleshooting Forms Transaction Playback.....	14-47
14.14.1.2	Troubleshooting Forms Transaction Recording	14-48
14.14.1.3	Troubleshooting End-User Performance of Forms Transactions	14-49
14.14.2	Verifying and Troubleshooting Web Transactions.....	14-50
15	Extending Enterprise Manager	
15.1	Benefits of Extending Enterprise Manager	15-1
15.2	More Extensibility Information.....	15-1
16	High Availability Solutions	
16.1	Latest High Availability Information.....	16-1
16.2	Defining High Availability	16-2
16.2.1	Levels of High Availability	16-2
16.3	Determining Your High Availability Needs.....	16-3
16.4	RTO, RPO, and Availability Levels.....	16-3
17	High Availability: Single Resource Configurations	
17.1	About Single Resource Configurations	17-1
17.2	Deploying Grid Control Components on a Single Host	17-2
17.3	Backup and Recovery	17-4
17.3.1	Repository Backup and Recovery	17-4
17.3.1.1	Repository Backup	17-4
17.3.1.2	Repository Recovery	17-7
17.3.1.3	Recovery Scenarios.....	17-9
17.3.2	Oracle Management Service Backup and Recovery	17-11
17.3.2.1	Backing Up the OMS.....	17-11
17.3.2.2	Recovering the OMS	17-12
17.3.2.3	OMS Recovery Scenarios.....	17-13
17.3.3	Agent Backup and Recovery	17-18
17.3.3.1	Backing Up Agents.....	17-18
17.3.3.2	Recovering Agents	17-18
17.3.3.3	Agent Recovery Scenarios.....	17-18
17.3.4	Recovering from a Simultaneous OMS-Repository Failure	17-19
17.3.4.1	Collapsed Configuration: Incomplete Repository Recovery, Primary OMS on the Same Host	17-19
17.3.4.2	Distributed Configuration: Incomplete Repository Recovery, Primary OMS and additional OMS on Different Hosts, SLB Configured	17-20
17.3.5	EMCTL High Availability Commands.....	17-20
18	High Availability: Multiple Resource Configurations	
18.1	Managing Multiple Hosts and Deploying a Remote Management Repository	18-1
18.2	Using Multiple Management Service Installations.....	18-3

18.3	High Availability Configurations	18-5
18.3.1	Configuring the Management Repository	18-6
18.3.1.1	Post Management Service - Install Management Repository Configuration ...	18-6
18.3.2	Configuring the Management Services	18-7
18.3.2.1	Management Service Install Location.....	18-7
18.3.2.2	Configuring the First Management Service for High Availability.....	18-8
18.3.2.3	Configuring Management Service to Management Repository Communication	18-8
18.3.2.4	Configuring Shared File System Loader	18-8
18.3.2.5	Configuring Software Library	18-10
18.3.2.6	Configuring a Load Balancer	18-10
18.3.3	Configuring Additional Management Services	18-13
18.3.3.1	Installing a Fresh Additional Management Service According MAA Best Practices	18-13
18.3.3.2	Retrofitting MAA Best Practices on Existing Additional Management Service	18-13
18.3.4	Configuring the Management Agent.....	18-14
18.3.4.1	Load Balancing Connections Between the Management Agent and the Management Service	18-15
18.3.5	Disaster Recovery	18-17
18.3.5.1	Prerequisites	18-18
18.3.5.2	Setup Standby Database	18-18
18.3.5.3	Setup Standby Management Service	18-19
18.3.5.4	Switchover	18-21
18.3.5.5	Failover.....	18-22
18.3.5.6	Automatic Failover.....	18-24
18.4	Installation Best Practices for Enterprise Manager High Availability	18-26
18.4.1	Configuring the Management Agent to Automatically Start on Boot and Restart on Failure	18-26
18.4.2	Configuring Restart for the Management Agent	18-27
18.4.3	Installing the Management Agent Software on Redundant Storage	18-27
18.4.4	Install the Management Service Shared File Areas on Redundant Storage.....	18-27
18.5	Configuration With Grid Control.....	18-27
18.5.1	Console Warnings, Alerts, and Notifications	18-28
18.5.2	Configure Additional Error Reporting Mechanisms.....	18-28
18.5.3	Component Backup.....	18-28
18.5.4	Troubleshooting.....	18-28
18.5.4.1	Upload Delay for Monitoring Data.....	18-28
18.5.4.2	Notification Delay of Target State Change	18-29
18.6	Configuring Oracle Enterprise Manager for Active and Passive Environments	18-29
18.7	Using Virtual Host Names for Active and Passive High Availability Environments in Enterprise Manager Database Control	18-29
18.7.1	Set Up the Alias for the Virtual Host Name and Virtual IP Address	18-30
18.7.2	Set Up Shared Storage.....	18-30
18.7.3	Set Up the Environment.....	18-30
18.7.4	Ensure That the Oracle USERNAME, ID, and GROUP NAME Are Synchronized on All Cluster Members	18-30
18.7.5	Ensure That Inventory Files Are on the Shared Storage.....	18-31

18.7.6	Start the Installer	18-31
18.7.6.1	Windows Specific Configuration Steps.....	18-31
18.7.7	Start Services.....	18-31
18.8	Configuring Grid Control Repository in Active/Passive High Availability Environments... 18-32	
18.8.1	Installation and Configuration	18-32
18.8.2	Set Up the Virtual Host Name/Virtual IP Address.....	18-33
18.8.3	Set Up the Environment.....	18-33
18.8.4	Synchronize Operating System User IDs	18-34
18.8.5	Set Up Inventory	18-34
18.8.6	Install the Software	18-34
18.8.6.1	Windows Specific Configuration Steps.....	18-35
18.8.7	Startup of Services	18-35
18.8.8	Summary	18-35
18.9	How to Configure Grid Control OMS in Active/Passive Environment for High Availability Failover Using Virtual Host Names 18-35	
18.9.1	Overview and Requirements	18-35
18.9.2	Installation and Configuration	18-36
18.9.3	Setting Up the Virtual Host Name/Virtual IP Address	18-36
18.9.4	Setting Up Shared Storage.....	18-36
18.9.5	Setting Up the Environment	18-37
18.9.6	Synchronizing Operating System IDs.....	18-37
18.9.7	Setting Up Shared Inventory.....	18-37
18.9.8	Installing the Software	18-37
18.9.8.1	Windows Specific Configuration Steps.....	18-38
18.9.9	Starting Up Services	18-38
18.9.10	Summary	18-38
18.10	Configuring Targets for Failover in Active/Passive Environments	18-39
18.10.1	Target Relocation in Active/Passive Environments	18-39
18.10.2	Installation and Configuration	18-39
18.10.2.1	Prerequisites	18-39
18.10.2.2	Configuration Steps.....	18-40
18.10.3	Failover Procedure.....	18-40
18.10.4	Fallback Procedure	18-41
18.10.5	EM CLI Parameter Reference.....	18-41
18.10.6	Script Examples.....	18-41
18.10.6.1	Relocation Script	18-42
18.10.6.2	Start Listener Script	18-43
18.10.6.3	Stop Listener Script	18-43

19 Management Agent and Management Services

19.1	Reconfiguring the Oracle Management Agent.....	19-1
19.1.1	Configuring the Management Agent to Use a New Management Service.....	19-1
19.1.2	Securing the Management Agent	19-2
19.1.3	Changing the Management Agent Port.....	19-2
19.1.4	Controlling the Amount of Disk Space Used by the Management Agent	19-3
19.1.5	About the Management Agent Watchdog Process.....	19-4

19.1.6	Setting the Management Agent Time Zone	19-5
19.1.6.1	Understanding How the Management Agent Obtains Time Zone Information	19-5
19.1.6.2	Resetting the Time Zone of the Management Agent Due to ...Inconsistency of Time Zones	19-7
19.1.6.3	Troubleshooting Management Agent Time Zone Problems.....	19-7
19.1.7	Adding Trust Points to the Management Agent Configuration.....	19-9
19.2	Reconfiguring the Oracle Management Service	19-9
19.2.1	Configuring the Management Service to Use a New Management Repository	19-9
19.2.1.1	Changing the Repository Properties.....	19-9
19.2.1.2	About Changing the Repository Password	19-10
19.2.2	Configuring the Management Service to Prompt You When Using Execute Commands	19-11
19.2.3	Troubleshooting Management Service Time Zone Problems	19-11

Index

Preface

The Enterprise Manager Administrator's Guide is an administrative reference that teaches you how to perform day-to-day Enterprise Manager administrative tasks. The goal of this book is to help you understand the concepts behind Enterprise Manager. It teaches you how to perform all common administration tasks needed to effectively monitor and manage targets within your Enterprise Manager environment.

Audience

This document is intended for all levels of Enterprise Manager administrators, as well as database, network, and application administrators using Enterprise Manager to monitor and maintain their IT infrastructure.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at

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

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

For more information, see the following manuals in the Oracle Enterprise Manager 10g Release 2 documentation set:

- *Oracle Enterprise Manager Grid Control Quick Installation Guide*
- *Oracle Enterprise Manager Grid Control Basic Installation Guide*
- *Oracle Enterprise Manager Grid Control Advanced Installation and Configuration Guide*
- *Oracle Enterprise Manager Configuration for Oracle Collaboration Suite*
- *Oracle Enterprise Manager Policy Reference Manual*
- *Oracle Enterprise Manager Metric Reference Manual*
- *Extending Oracle Enterprise Manager*

- *Oracle Enterprise Manager Command Line Interface*
- *SNMP Support Reference Guide*

The latest versions of this and other Enterprise Manager books can be found at:

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

Oracle Enterprise Manager also provides extensive online help. Click **Help** on any Oracle Enterprise Manager page to display the online help system.

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

Basic Administration

This section of the guide covers using Enterprise Manager framework functionality to monitor your managed environment.

Part I contains the following chapters:

- [Chapter 1, "Monitoring"](#)
- [Chapter 2, "Enterprise Manager Security"](#)
- [Chapter 3, "Notifications"](#)
- [Chapter 4, "User-Defined Metrics"](#)
- [Chapter 5, "Group Management"](#)
- [Chapter 6, "Job System"](#)
- [Chapter 7, "Starting and Stopping Enterprise Manager Components"](#)
- [Chapter 8, "Information Publisher"](#)
- [Chapter 9, "Sizing Your Enterprise Manager Deployment"](#)
- [Chapter 10, "Maintaining and Troubleshooting the Management Repository"](#)

Monitoring

Because of the size, complexity, and criticality of today's enterprise IT operations, the challenge for IT professionals is to be able to maintain high levels of component availability and performance for both applications and all components that make up the application's technology stack. Monitoring the performance of these components and quickly correcting problems before they can impact business operations is crucial. Enterprise Manager provides comprehensive, flexible, easy-to-use monitoring functionality that supports the timely detection and notification of impending IT problems across your enterprise.

This chapter covers the following topics:

- [Systems Monitoring: Breadth and Depth](#)
- [Monitoring Basics](#)
- [Monitoring Templates](#)
- [User-Defined Metrics](#)
- [Accessing Monitoring Information](#)

1.1 Systems Monitoring: Breadth and Depth

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

Enterprise Manager provides the most comprehensive management features for all Oracle products. For example, Enterprise Manager's monitoring functionality is tightly integrated with Oracle Database 10g manageability features such as server-generated alerts. These alerts are generated by the database itself about problems it has self-detected. Server-generated alerts can be managed from the Enterprise Manager console and include recommendations on how problems can be resolved. Performance problems such as poorly performing SQL and corresponding recommendations that are generated by the database's self-diagnostic engine, called Automatic Database Diagnostic Monitor (ADDM), are also captured and exposed through the Enterprise Manager console. This allows Enterprise Manager administrators to implement ADDM recommendations with ease and convenience.

Enterprise Manager also makes it easy to expand the scope of system monitoring beyond individual components. Using Enterprise Manager's group management functionality, you can easily organize monitorable targets into groups, allowing you to monitor and manage many components as one.

1.2 Monitoring Basics

System 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 (hosts, operating systems, storage, and so on).

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

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

1.2.1 Out-of-Box Monitoring

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

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

- In-depth monitoring with Oracle-recommended metrics and thresholds.
- 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 notification for detected critical alerts.

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

Some examples of monitored metrics are:

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

Some metrics have associated predefined limiting parameters called thresholds that cause alerts to be triggered when collected metric values exceed these limits.

Enterprise Manager allows you to set metric threshold values for two levels of alert severity:

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

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

1.2.1.1 Metric Thresholds

As mentioned earlier, some metric thresholds come predefined out-of-box. While these values are acceptable 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 OPERATOR or greater privilege on a target in order to perform any metric threshold changes.

1.2.1.1.1 Metric Snapshots A metric snapshot is a named collection of a target's performance metrics that have been collected at a specific point in time. A metric snapshot can be used as an aid in calculating metric threshold values based on the target's past performance.

The key in defining a metric snapshot for a target is to select a date during which target performance was acceptable under typical workloads. Given this date, actual values of the performance metrics for the target are retrieved and these represent what is normal or expected performance behavior for the target. Using these values, you can then use Enterprise Manager to calculate warning and critical thresholds for the metrics that are a specified percentage 'worse' than the actual metric snapshot values. These represent values which, when crossed, could indicate performance problems. After thresholds are calculated, you can still edit the calculated values if needed.

You can define a metric snapshot for a target based on a date and (optionally) time. If you only specify a date, the metric snapshot is the set of average daily values of the

target's performance metrics for that date. If you also specify an hour within the date, then the metric snapshot is the set of Low and High metric values for the preceding hour.

Metric snapshots apply to all monitored targets except 10.2 or higher databases, Services, and Web applications. For these targets, the Metric Baseline feature is supported.

1.2.1.1.2 Metric Baselines Metric baselines are statistical characterizations of system performance over well-defined time periods. Metric baselines can be used to implement adaptive alert thresholds for certain performance metrics as well as provide normalized views of system performance. Adaptive alert thresholds are used to detect unusual performance events. Baseline normalized views of metric behavior help administrators explain and understand such events.

Metric baselines are well defined time intervals (baseline periods) over which Enterprise Manager has captured system performance metrics. The underlying assumption of metric baselines is that systems with relatively stable performance should exhibit similar metric observations (that is, values) over times of comparable workload. Two types of baseline periods are supported: moving window baseline periods and static 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 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).

Adaptive Thresholds

Once metric baselines are defined, they can be used to establish alert thresholds that are statistically significant and adapt to expected variations across time. For example, you can define alert thresholds to be generated based on significance level, such as the HIGH significance level thresholds are values that occur 5 in 100 times. Alternatively, you can generate thresholds based on a percentage of the maximum value observed within the baseline period. These can be used to generate alerts when performance metric values are observed to exceed normal peaks within that period.

Baseline Normalized Views

Enterprise Manager provides charts which graphically display the values of observed performance and workload metrics normalized against the baseline. Using these charts, statistically significant values are easily seen as 'blips' in the charts. These allow administrators to easily perform time-correlation of events. For example, performance events can be related to significantly increased demand or significantly unusual workload.

Metric baselines are supported for databases (version 10.2 or higher) and for Services and Web Application target types.

1.2.2 Alerts

When a metric threshold value is reached, an alert is generated. An alert indicates a potential problem; either a warning or critical threshold for a monitored metric has been crossed. An alert can also be generated for various target availability states, such as:

- Target is down.
- Oracle Management Agent monitoring the target is unreachable.

When an alert is generated, you can access details about the alert from the Enterprise Manager console. See "[Accessing Monitoring Information](#)" on page 1-8 for more information on viewing alert information.

Enterprise Manager provides various options to respond to alerts. Administrators can be automatically notified when an alert triggers and/or corrective actions can be set up to automatically resolve an alert condition.

1.2.3 Notifications

When a target becomes unavailable or if thresholds for performance are crossed, alerts are generated in the Enterprise Manager console and notifications are sent to the appropriate administrators. Enterprise Manager supports notifications via e-mail (including e-mail-to-page systems), SNMP traps, and/or by running custom scripts.

Enterprise Manager supports these various notification mechanisms via notification methods. A notification method is used to specify the particulars associated with a specific notification mechanism, for example, which SMTP gateway(s) to use for e-mail, which OS script to run to log trouble-tickets, and so on. Super Administrators perform a one-time setup of the various types of notification methods available for use. Once defined, other administrators can create notification rules that specify the set of criteria that determines when a notification should be sent and how it should be sent. The criteria defined in notification rules include the targets, metrics and severity states (clear, warning or critical) and the notification method that should be used when an alert occurs that matches the criteria. For example, you can define a notification rule that specifies e-mail should be sent to you when CPU Utilization on any host target is at critical severity, or another notification rule that creates a trouble-ticket when any database is down. Once a notification rule is defined, it can be made public for sharing across administrators. For example, administrators can subscribe to the same rule if they are interested in receiving alerts for the same criteria defined in the rule. Alternatively, an Enterprise Manager Super Administrator can assign notification rules to other administrators such that they receive notifications for alerts as defined in the rule.

Notifications are not limited to alerting administrators. Notification methods can be extended to execute any custom OS script or PL/SQL procedure, and thus can be used to automate any type of alert handling. For example, administrators can define notification methods that call into a trouble ticketing system, invoke third-party APIs to share alert information with other monitoring systems, or log a bug against a product.

1.2.3.1 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; e-mail addresses that are used to send pages can be associated with the 'short' format. The 'long' format contains full details about the alert; 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 alerts, they simply leave that part of the schedule blank. All alerts that are sent to an administrator automatically adhere to his specified schedule.

1.2.4 Corrective Actions

Corrective actions allow you to specify automated responses to alerts. Corrective actions ensure that routine responses to alerts are automatically executed, thereby saving administrator time and ensuring problems 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 thus 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 alert has triggered. 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. For example, if an Oracle Application Server's J2EE container (called an OC4J container) triggers a warning alert indicating it is approaching its limit on the number of requests it can handle, a corrective action can be defined to automatically start up another OC4J container on another host, thereby sharing application load among different containers. As shown by this example, corrective actions can be used to dynamically allocate resources as demand increases, thereby preventing performance bottlenecks before they impact overall application availability.

Corrective actions for a target can be defined by all Enterprise Manager administrators who have been granted OPERATOR 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.

1.2.5 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 (EMCLI). The EMCLI 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.

1.3 Monitoring Templates

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

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

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

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

Comparing Differences Between Targets and Monitoring Templates

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

1.4 User-Defined Metrics

User-defined metrics allow you to extend the reach of Enterprise Manager's monitoring to conditions specific to particular environments via custom scripts or SQL queries and function calls. Once a user-defined metric is defined, it will be monitored, aggregated in the repository, and can trigger alerts like any other metric in Enterprise Manager. There are two types of user-defined metrics: Operating System and SQL.

- *Operating System (OS) User-Defined Metrics*: Accessed from Host target home pages, these user-defined metrics allow you to implement custom monitoring functions via OS scripts.
- *SQL User-Defined Metrics*: Accessed from the Database target home pages, these user-defined metrics allow you to implement custom database monitoring using SQL queries or function calls.

Creating a User-Defined Metric

To monitor a particular condition (example: check successful completion of monthly system maintenance routines), you can write a custom OS script to monitor that condition, then register it as a user-defined metric in Enterprise Manager. Each time the metric is evaluated by Enterprise Manager, it uses this script to evaluate the condition. SQL user-defined metrics do not use external scripts: you enter SQL directly into the Enterprise Manager console at the time of metric creation. Once a user-defined metric is defined, all other monitoring features, such as alerts, notifications, historical collections, and corrective actions are automatically available to it.

If you already have your own library of custom monitoring scripts, you can leverage Enterprise Manager's monitoring features by integrating these scripts with Enterprise Manager as OS user-defined metrics. Likewise, existing SQL queries or function calls currently used to monitor database conditions can be easily integrated into Enterprise Manager's monitoring framework as SQL user-defined metrics. For more information about user-defined metrics, see *Oracle Enterprise Manager Advanced Configuration*.

1.5 Accessing Monitoring Information

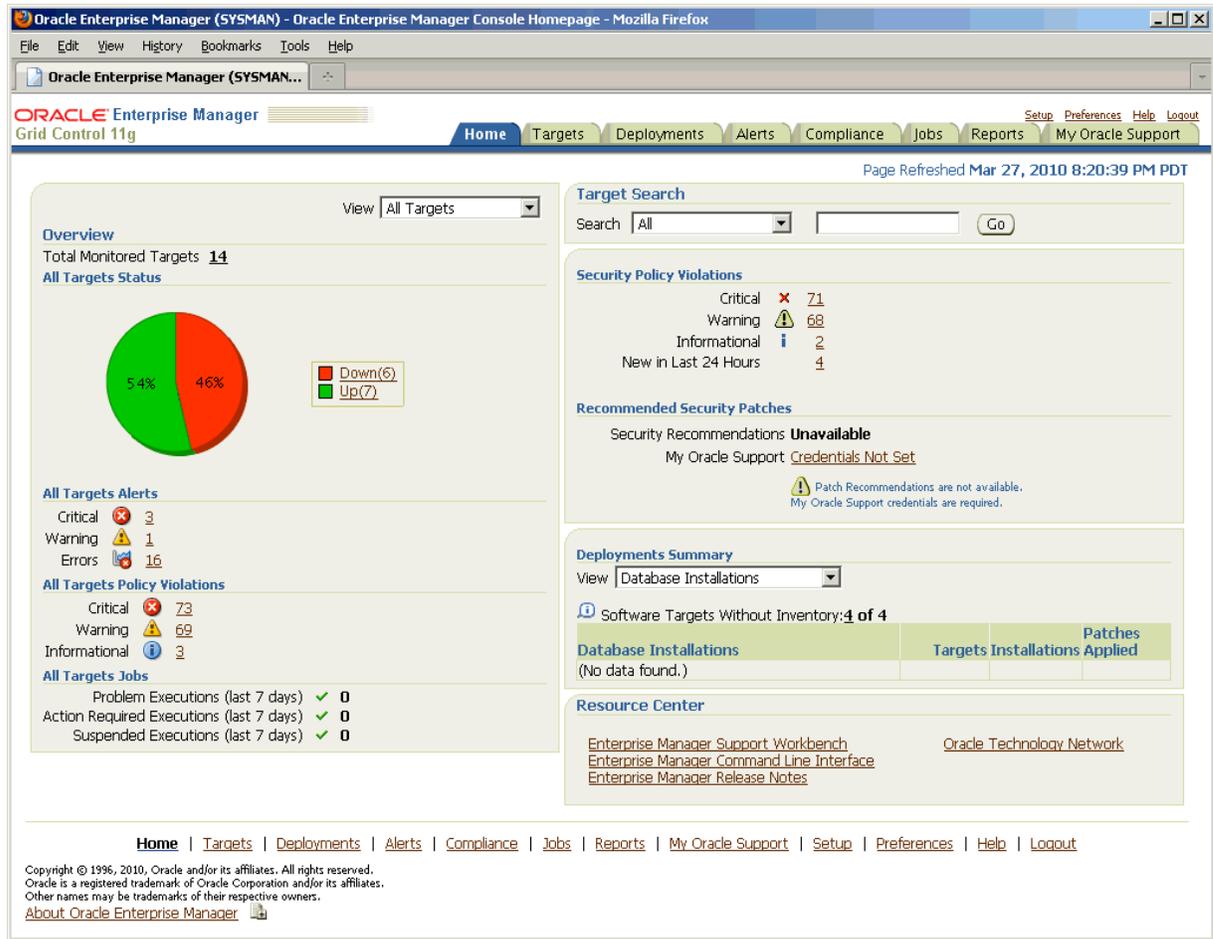
All monitoring information is accessed via the Enterprise Manager console, providing quick views into the health of your monitored environment.

Enterprise Manager Console Home Page

The Enterprise Manager console home page shown in [Figure 1–1](#) gives you an at-a-glance view of the overall status of your monitored environment. As shown in the following figure, the home page summarizes key monitoring areas such as availability across all managed targets, open alerts, policy violations, and recent problems with job executions. Links on this page allow you to drill down to detailed performance information.

The Resource Center is your central access point to Enterprise Manager documentation as well as the comprehensive technical resources of the Oracle Technology Network (OTN).

Figure 1-1 Enterprise Manager Console



From the home page, you can easily access alert information. For example, you can click on the Down link in the All Targets Status legend to determine which targets are currently down. Under All Target Alerts, you can click on the Warning alerts value to access a list of warning alerts for all monitored targets (Figure 1-2).

Figure 1–2 Warning Alerts Page

The screenshot shows the Oracle Enterprise Manager interface for 'Warning Alerts'. The page title is 'Warning Alerts' and it indicates the page was refreshed on Mar 28, 2010 at 3:43:24 AM UTC. A search bar is present with 'All Targets' selected for the target type and 'Day(s)' for the alert trigger window. Below the search bar is a table of alerts. The table has the following columns: Target, Type, Alert Triggered, Message, Acknowledged, Last Collected Value, and Last Collected Timestamp. The alerts are sorted by most recent first.

Target	Type	Alert Triggered	Message	Acknowledged	Last Collected Value	Last Collected Timestamp
hello	Cluster Database	Mar 27, 2010 5:26:06 AM PDT	db recovery file dest size of 6442450944 bytes is 85.18% used and has 954551296 remaining ...		0.37	Mar 22, 2010 9:05:48 AM PDT
adc2100510.us.oracle.com	Host	Mar 22, 2010 6:36:56 AM PDT	Memory Utilization is 83.09%, crossed warning (80) or critical (95) threshold.		82.66	Mar 27, 2010 8:36:56 PM PDT
hello	Cluster Database	Mar 22, 2010 6:33:26 AM PDT	Streams component PROPAGATION\$ 16 has 1 errors.		1	Mar 24, 2010 4:37:43 AM PDT
hello_hello1	Database Instance	Mar 22, 2010 6:30:34 AM PDT	User SYS logged on from stait09.		SYS	Mar 25, 2010 9:00:34 PM PDT
emrep.us.oracle.com	Database Instance	Mar 22, 2010 1:02:48 PM GMT	Metrics "Current Open Cursors Count" is at 3,405		3724	Mar 28, 2010 3:35:01 AM GMT

The most recent alerts are listed first. You can change the sorting methodology by clicking on the appropriate column header. By clicking on a specific alert message, you can drill down to explicit details about the metric in alert ().

Figure 1–3 Warning Alert: Metric Details

Oracle Enterprise Manager (SYSMAN) - Metric Alert - Mozilla Firefox

ORACLE Enterprise Manager
Grid Control 11g

Database Instance: Oemrep_Database > All Metrics > Current Open Cursors Count >

Metric Alert : Current Open Cursors Count

Page Refreshed Mar 27, 2010 8:26:38 PM PDT

Alert Details

Metric **Current Open Cursors Count**
Severity **Warning**
Alert Triggered **Mar 26, 2010 6:26:40 PM**
Last Updated **Mar 26, 2010 6:26:40 PM**
Acknowledged **No**
Acknowledged By **n/a**
Message **Metrics "Current Open Cursors Count" is at 1,790**

Recommendations

Run ADDM to get more performance analysis about your system.

Actions

[Edit Thresholds](#)

Links

[Additional Advice](#) [Initialization Parameters](#)

Metric Data

Last Known Value **2092**
Last Collection Timestamp **Mar 27, 2010 9:11:24 PM**

2,500
2,000
1,500
1,000
500
0

7:21 9 12 AM 3 6 9 12 PM 3 6 9
Mar 26, 2010 27

Oemrep_Database

Metric Settings

Warning Threshold **1200**
 Critical Threshold **Not Defined**
Occurrences Before Alert **3**

Updates

New Comment

Select	Timestamp	Type	Administrator	Message
<input checked="" type="radio"/>	Mar 26, 2010 6:26:40 PM	Warning	<SYSTEM>	Metrics "Current Open Cursors Count" is at 1,790

Home | **Targets** | Deployments | Alerts | Compliance | Jobs | Reports | My Oracle Support | Setup | Preferences | Help | Logout

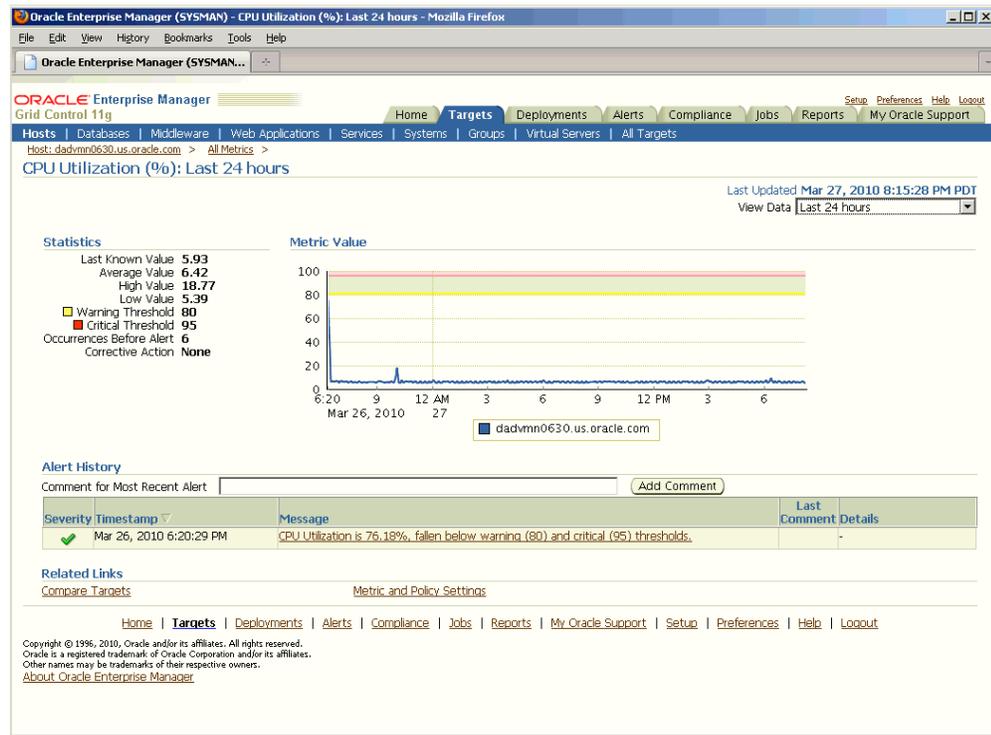
Copyright © 1996, 2010, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.
[About Oracle Enterprise Manager](#)

By default, metric values shown on this page reflect the last 24 hours of collected data. You can also select another time period or specify a custom time period with which to view metric data and easily assess if the problem occurred recently or across a long time period. Because Enterprise Manager collects and aggregates metric data in the Management Repository, you can click on the Compare Targets related link to display metric data for more than one target simultaneously, thus allowing you to compare performance across multiple targets (Figure 1–4).

If you do not wish to view metrics collected over time, you can choose one of several Real Time metric refresh periods:

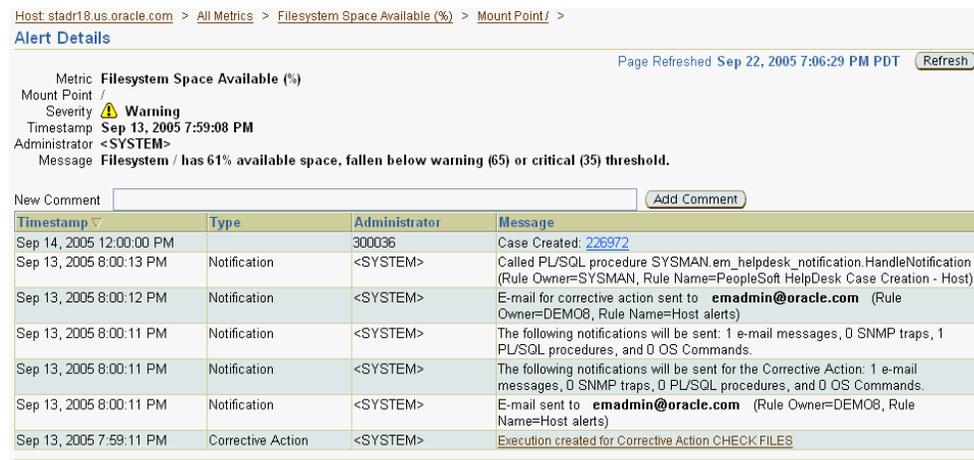
- Manual
- 30 Second
- 1 Minute
- 5 Minutes

Figure 1–4 Compare Targets



The Alert History table shows alerts generated over the selected time period. You can view explicit details about a specific alert in this table by clicking on the eyeglasses icon in the Details column. Figure 1–5 shows the Alert Details page.

Figure 1–5 Alert Details



The Alert Details page shows all notifications for an alert, any corrective actions that have been executed, and any custom notifications, for example, the opening of a case ticket for an alert. On this page, you also have the option of adding annotations or comments for other administrators to see.

Enterprise Manager Security

This chapter describes how to configure Oracle Enterprise Manager Security. Specifically, this chapter contains the following sections:

- [About Oracle Enterprise Manager Security](#)
- [Enterprise Manager Authentication](#)
- [Enterprise Manager Authorization](#)
- [Configuring Security for Grid Control](#)
- [Accessing Managed Targets](#)
- [Cryptographic Support](#)
- [Setting Up the Auditing System for Enterprise Manager](#)
- [Additional Security Considerations](#)

2.1 About Oracle Enterprise Manager Security

Oracle Enterprise Manager provides tools and procedures to help you ensure that you are managing your Oracle environment in a secure manner. The goals of Oracle Enterprise Manager security are:

- To be sure that only users with the proper privileges have access to critical monitoring and administrative data.

This goal is met by requiring username and password credentials before users can access the Enterprise Manager consoles and appropriate privileges for accessing the critical data.
- To be sure that all data transferred between Enterprise Manager components is transferred in a secure manner and that all data gathered by each Oracle Management Agent can be transferred only to the Oracle Management Service for which the Management Agent is configured.

This goal is met by enabling Enterprise Manager Framework Security. Enterprise Manager Framework Security automates the process of securing the Enterprise Manager components installed and configured on your network.
- To be sure that sensitive data such as credentials used to access target servers are protected.

This goal is met by Enterprise Manager's encryption support. The sensitive data is encrypted with an **emkey**. By following the best practice, even the repository owner and the SYSDBA will not be able to access the sensitive data.

- To be sure that access to managed targets is controlled through user authentication and privilege delegation.

This goal is met by configuring the Management Agent with PAM and LDAP for user authentication and using privilege delegation tools like Sudo and PowerBroker.

2.2 Enterprise Manager Authentication

Grid Control Authentication is the process of determining the validity of the user accessing Enterprise Manager Grid Control. The authentication feature is available across the different user interfaces such as Enterprise Manager Grid Control console and Enterprise Manager Command Line Interface.

The following authentication schemes are available:

- **Repository-Based Authentication:** This is the default authentication option. An Enterprise Manager administrator is also a repository (database) user. By using this option, you can take advantage of all the benefits that this authentication method provides like password control via password profile, enforced password complexity, password life time, number of failed attempts allowed and controls. During the password grace period, the administrator is prompted to change the password but when the password has expired, it must be changed. For more details, refer to [Repository-Based Authentication](#).
- **SSO-Based Authentication:** The single sign-on based authentication provides strengthened and centralized user identity management across the enterprise. After you have configured Enterprise Manager to use the Oracle Application Server Single Sign-On, you can register any single sign-on user as an Enterprise Manager administrator. You can then enter your single sign-on credentials to access the Oracle Enterprise Manager Grid Control console. For more details, refer to [Single Sign-On Based Authentication](#).
- **Enterprise User Security Based Authentication:** The Enterprise User Security (EUS) option enables you to create and store enterprise users and roles for the Oracle database in an LDAP-compliant directory server. Once the repository is configured with EUS, you can configure Enterprise Manager to use EUS as its authentication mechanism as described in [Enterprise User Security Based Authentication](#). You can register any EUS user as an Enterprise Manager administrator.

EUS helps centralize the administration of users and roles across multiple databases. If the managed databases are configured with EUS, the process of logging into these databases is simplified. When you drill down to a managed database, Enterprise Manager will attempt to connect to the database using Enterprise Manager Grid Control credentials. If successful, Enterprise Manager will directly connect you to the database without displaying a login page.

2.2.1 Repository-Based Authentication

Enterprise Manager Grid Control allows you to create and manage new administrator accounts. Each administrator account includes its own login credentials as well as a set of roles and privileges that are assigned to the account. You can also assign a password profile to the administrator. To create, edit, or view an administrator account:

1. From Enterprise Manager Grid Control, click **Setup**.
2. Click **Administrators** in the vertical navigation bar.

3. Click the appropriate task button on the Administrators page. The following screen is displayed:

Figure 2–1 Create / Edit Administrator

ORACLE Enterprise Manager 11g
Grid Control

Home Targets Deployments Alerts Compliance Jobs Reports My Oracle Support

Enterprise Manager Configuration | Management Services and Repository | Agents

Properties Roles System Privileges Target Privileges Review

Edit Administrator SYS: Properties

Cancel Step 1 of 5 Next Review

Password: [masked]
Confirm Password: [masked]
Password Profile: DEFAULT View
 Prevent password change
When checked, administrator is not allowed to change its own password.
 Expire password now
When selected, administrator account will be created with expired state. On next login, administrator will be forced to change password.
E-mail Address: [text area]
Description: [text area]
 Super Administrator

Cancel Step 1 of 5 Next Review

On this page, you can specify the type of administrator account being created, select the password profile, and the password expiry period. The password cannot be changed by the administrator if the **Prevent Password Change** checkbox is selected.

If you select the **Expire Password Now** checkbox, the password for administrator account will be set to an expired state. If the password has expired, when you login the next time, the following screen is displayed and you are prompted to change the password.

Figure 2–2 Password Expiry Page

ORACLE Enterprise Manager 11g
Grid Control

Home Targets Deployments Alerts Compliance Jobs Reports My Oracle Support

Setup Preferences Help Logout

Preferences

General Preferred Credentials Notification My Rules Public Rules Schedule Target Subtabs Custom Audit Attributes Accessibility

Information
Your current password has expired. Please change password first.

General Revert Apply

Password
To change your password, specify and confirm a new password.

Administrator TEST002
• Current Password [text input]
• New Password [text input]
• Confirm New Password [text input]

Revert Apply

Enter your current password and the new password and click **Apply**. You can now start using Enterprise Manager Grid Control.

2.2.2 Single Sign-On Based Authentication

If you are currently using Oracle Application Server Single Sign-On to control access and authorization for your enterprise, you can extend those capabilities to the Grid Control Console.

By default, when you navigate to the Grid Control Console, Enterprise Manager displays the Enterprise Manager login page. However, you can configure Enterprise

Manager so it uses Oracle Application Server Single Sign-On to authenticate your Grid Control Console users. Instead of seeing the Enterprise Manager login page, Grid Control Console users will see the standard Oracle Application Server Single Sign-On login page. From the login page, administrators can use their Oracle Application Server Single Sign-On credentials to access the Oracle Enterprise Manager 10g Grid Control Console.

Note:

- You can configure Enterprise Manager Grid Control to either use Oracle Application Server Single Sign-On or the Enterprise User Security features. You cannot use both options at the same time.
 - When Enterprise Manager is configured to use Single Sign-On with Server Load Balancer, make sure that the correct monitoring settings have been defined. For details, refer to the chapter on *Grid Control Common Configurations*.
-
-

The following sections describe how to configure Enterprise Manager as an OracleAS Single Sign-On Partner Application:

- [Registering Enterprise Manager as a Partner Application](#)
- [Removing Single Sign-On Configuration](#)
- [Registering Single Sign-On Users as Enterprise Manager Administrators](#)
- [Grid Control as a Single Sign-On Partner Application](#)
- [Bypassing the Single Sign-On Logon Page](#)

2.2.2.1 Registering Enterprise Manager as a Partner Application

To register Enterprise Manager as a partner application manually, follow these steps:

1. Enter the following URL to navigate to the SSO Administration page.

```
http://sso_host:sso_port/pls/orasso
```
2. Login as `orcladmin` user and click **SSO Administration**.
3. Click **Administer Partner Applications** and then click **Add Partner Application**.
4. Enter the following information on the Add Partner Application page.

```
Name: <EMPartnerName>
Home URL: protocol://em_host:em_port
Success URL: protocol://em_host:em_port/osso_login_success
Logout URL: protocol://em_host:em_port/osso_logout_success
Administrator Email: user@host.com
```

where `host`, `port`, and `protocol` refer to the EM Host, port and the protocol (`http` or `https`) used.

5. After entering these details, click **Edit <EMPartnerName>** and enter the following parameters to generate the `osso.txt`. Sample values for these parameters are shown below:

```
sso_server_version: v1.2
cipher_key: <EncryptionKeyValue>
site_id: <IDValue>
site_token: <TokenValue>
```

```
login_url=protocol://sso_host:sso_port/pls/orasso/orasso.wvssso_app_admin.ls_
login
logout_url=protocol://sso_host:sso_port/pls/orasso/orasso.wvssso_app_admin.ls_
logout
cancel_url=protocol://em_host:em_port
sso_timeout_cookie_name=SSO_ID_TIMEOUT
sso_timeout_cookie_key=9E231B3C1A3A808A
```

6. Enter the following command to generate the `osso.conf` file:

```
WEBTIER_HOME/ohs/bin/iasobf osso.txt osso.conf root
```

7. Use the `osso.conf` file and configure it as necessary using the `emctl` command as follows:

```
emctl config oms sso -ossoconf <ossoconf file> -dasurl <dasurl> [-unsecure]
[-sysman_pwd <pwd>] [-domain <domain>]
```

where:

- `-ossoconf` is the path to the `osso.conf` file
- `-dasurl` is the URL specifying the host and port for the Delegated Administration Service (DAS). Generally, the DAS host name and port are the same as the host name and port of the Oracle Application Server Single Sign-On server. For example:

```
http://mgmthost1.acme.com:7777
```
- `-unsecure` is used to register the http port with the single sign-on server.
- `-sysman_pwd` is the sysman user password. If this parameter is not specified, you will be prompted to enter it.
- `-domain` is the name of the host domain. This parameter needs to be specified if the fully qualified name of the host is not available.

The sample output for this command is shown below:

```
emctl config oms sso -ossoconf /tmp/osso.conf -dasurl
http://somehost.domain.com:7777
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
Enter SYSMAN user password :
SSO Configuration done successfully, Please restart OMS.
```

8. Restart WebTier and OMS as follows:

```
emctl stop oms
emctl start oms
```

2.2.2.2 Removing Single Sign-On Configuration

To remove the single sign-on configuration, run the following command:

```
emctl config oms sso -remove [-sysman_pwd <pwd>]
```

where `-sysman_pwd` is the sysman repository password.

Example 2-1 Sample Output of the `emctl config oms -remove` command

```
emctl config oms sso -remove
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
```

```
Enter SYSMAN user password :  
SSO Configuration removed successfully, Please restart OMS.  
Restart OMS using  
emctl stop oms  
emctl start oms
```

2.2.2.3 Registering Single Sign-On Users as Enterprise Manager Administrators

After you have configured Enterprise Manager to use the Single Sign-On logon page, you can register any Single Sign-On user as an Enterprise Manager administrator. You can register single sign-on users using:

- Enterprise Manager Grid Control Graphical User Interface
- Enterprise Manager Grid Control Command Line Interface

2.2.2.3.1 Registering Single Sign-On Users Using the Graphical User Interface

You can use the graphical user interface to register single sign-on users by following these steps:

1. Go the Enterprise Manager Grid Control Console URL.

For example:

```
http://mgmthost1.acme.com:7777/em
```

The browser is redirected to the standard Single Sign-On Logon page.

2. Enter the credentials for a valid Single Sign-On user.

If the Single Sign-On user is not an Enterprise Manager administrator, the browser is redirected to a modified version of the Enterprise Manager logon page (Figure 2-3).

3. Log in to Enterprise Manager as a Super Administrator.
4. Click **Setup** and then click **Administrators** to display the Administrators page.

See Also: "Creating, Editing, and Viewing Administrators" in the Enterprise Manager online Help

Because Enterprise Manager has been configured to use Single Sign-On, the first page in the Create Administrator wizard now offers you the option of creating an administrator based on a registered Oracle Internet Directory user.

5. Select **Oracle Internet Directory** and advance to the next page in the wizard.
6. Enter the name and e-mail address of the Oracle Internet Directory user, or click the flashlight icon to search for a user name in the Oracle Internet Directory.
7. Use the rest of the wizard pages to define the roles, system privileges, and other characteristics of the Enterprise Manager administrator and then click **Finish**.

Enterprise Manager displays a summary page that lists the characteristics of the administrator account.

8. Click **Finish** to create the new Enterprise Manager administrator.

The OID user is now included in the list of Enterprise Manager administrators. You can now verify the account by logging out of the Grid Control Console and logging back in using the OID user credentials on the Single Sign-On logon page.

Figure 2–3 Modified Enterprise Manager Logon Page When Configuring SSO


ORACLE Enterprise Manager 10g [Help](#)
Grid Control

Login

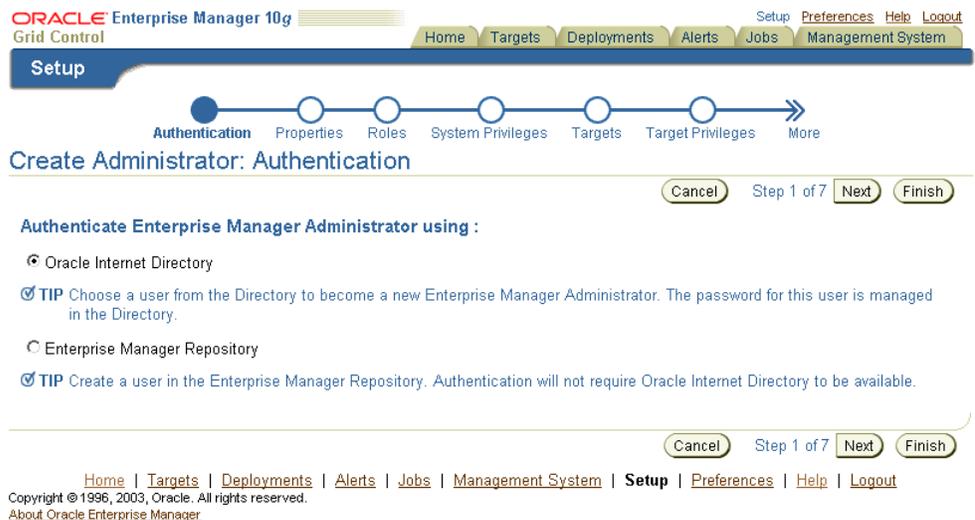
Log on as a different Single Sign On user.
Log On to Enterprise Manager using Enterprise Manager repository authentication below.

Login to Oracle Enterprise Manager

* User Name
* Password

Login

Copyright © 1996, 2003, Oracle. All rights reserved.

Figure 2–4 Create Administrator Page When SSO Support Is Enabled


ORACLE Enterprise Manager 10g [Setup](#) [Preferences](#) [Help](#) [Logout](#)
Grid Control [Home](#) [Targets](#) [Deployments](#) [Alerts](#) [Jobs](#) [Management System](#)

Setup

Authentication Properties Roles System Privileges Targets Target Privileges More

Create Administrator: Authentication

Cancel Step 1 of 7 Next Finish

Authenticate Enterprise Manager Administrator using :

Oracle Internet Directory

TIP Choose a user from the Directory to become a new Enterprise Manager Administrator. The password for this user is managed in the Directory.

Enterprise Manager Repository

TIP Create a user in the Enterprise Manager Repository. Authentication will not require Oracle Internet Directory to be available.

Cancel Step 1 of 7 Next Finish

[Home](#) | [Targets](#) | [Deployments](#) | [Alerts](#) | [Jobs](#) | [Management System](#) | **Setup** | [Preferences](#) | [Help](#) | [Logout](#)
Copyright © 1996, 2003, Oracle. All rights reserved.
[About Oracle Enterprise Manager](#)

2.2.2.3.2 Registering Single Sign-On Users Using EMCLI s

You can use the following EMCLI command to create Single Sign-On users:

```
emcli create_user -name=ssouser -type=EXTERNAL_USER
```

This command creates a user with the name **ssouser** who is authenticated against the single sign-on user.

Argument	Description
-name	Name of the administrator.
-type	The type of user. The default value for this parameter is EM_USER. The other possible values are: <ul style="list-style-type: none"> EXTERNAL_USER: Used for single-sign-on based authentication. DB_EXTERNAL_USER: Used for enterprise user based security authentication.
-password	The password for the administrator.

Argument	Description
-roles	The list of roles that can be granted to this administrator.
-email	The list of email addresses for this administrator.
-privilege	The system privileges that can be granted to the administrator. This option can be specified more than once.
-profile	The name of the database profile. This is an optional parameter. The default profile used is DEFAULT.
-desc	The description of the user being added.
-expired	This parameter is used to set the password to "expired" status. This is an optional parameter and is set to False by default.
-prevent_change_password	When this parameter is set to True, the user cannot change the password. This is an optional parameter and is set to False by default.
input_file	This parameter allows the administrator to provide the values for any of these arguments in an input file. The format of value is name_of_argument:file_path_with_file_name.

Example 1

```
emcli create_user
  -name="new_admin"
  -password="oracle"
  -email="first.last@oracle.com;joe.shmoe@shmoeshop.com"
  -roles="public"
  -privilege="view_job;923470234ABCDFE23018494753091111"
  -privilege="view_target;<host>.com:host"
```

This example creates an Enterprise Manager administrator named `new_admin`. This administrator has two privileges: the ability to view the job with ID `923470234ABCDFE23018494753091111` and the ability to view the target `<host>.com:host`. The administrator `new_admin` is granted the PUBLIC role.

Example 2

```
emcli create_user
  -name="User1"
  -type="EXTERNAL_USER"
  -input_file="privilege:/home/user1/priv_file"
```

Contents of `priv_file` are:
`view_target;<host>.com:host`

This example makes `user1` which has been created externally as an Enterprise Manager user. `user1` will have view privileges on `<host>.com:host`.

Example 3

```
emcli create_user
  -name="User1"
  -desc="This is temp hire."
  -prevent_change_password="true"
  -profile="MGMT_ADMIN_USER_PROFILE"
```

This example sets `user1` as an Enterprise Manager user with some description. The `prevent_change_password` is set to `true` to indicate that the password cannot be changed by `user1` and the `profile` is set to `MGMT_ADMIN_USER_PROFILE`.

Example 4

```
emcli create_user
      -name="User1"
      -desc="This is temp hire."
      -expire="true"
```

This example sets `user1` as an Enterprise Manager with some description. Since the password is set to expire immediately, when the user logs in for the first time, he is prompted to change the password.

2.2.2.4 Grid Control as a Single Sign-On Partner Application

The `emctl config oms sso` command adds the Oracle Enterprise Manager Grid Control Console as an Oracle Application Server Single Sign-On partner application. Partner applications are those applications that have delegated authentication to the Oracle Application Server Single Sign-On Server.

To see the list of partner applications, navigate to the following URL:

```
http://hostname:port/pls/orasso/orasso.home
```

For example:

```
http://ssohost1.acme.com:7777/pls/orasso/orasso.home
```

2.2.2.5 Bypassing the Single Sign-On Logon Page

After you configure Enterprise Manager to use the Single Sign-On logon page, you can bypass the Single Sign-On page at any time and go directly to the Enterprise Manager logon page by entering the following URL:

```
http://hostname.domain:port/em/console/logon/logon
```

For example:

```
http://mgmthost1.acme.com:7777/em/console/logon/logon
```

2.2.3 Enterprise User Security Based Authentication

Enterprise User Security enables you to create and store Oracle database information as directory objects in an LDAP-compliant directory server. For example, an administrator can create and store enterprise users and roles for the Oracle database in the directory, which helps centralize the administration of users and roles across multiple databases.

See Also: Enterprise User Security Configuration Tasks and Troubleshooting in the *Oracle Database Advanced Security Administrator's Guide*

If you currently use Enterprise User Security for all your Oracle databases, you can extend this feature to Enterprise Manager. Configuring Enterprise Manager for use with Enterprise User Security simplifies the process of logging in to database targets you are managing with the Oracle Enterprise Manager Grid Control Console.

To configure Enterprise Manager for use with Enterprise User Security:

1. Ensure that you have enabled Enterprise User Security for your Oracle Management Repository database, as well as the database targets you will be managing with the Grid Control Console. Refer to *Oracle Database Advanced Security Administrator's Guide* for details.

2. Using the `emctl set property` command, set the following properties:

```
oracle.sysman.emSDK.sec.DirectoryAuthenticationType=EnterpriseUser
oracle.sysman.emSDK.sec.eus.Domain=<ClientDomainName> (For
example:mydomain.com)
oracle.sysman.emSDK.sec.eus.DASHostUrl=<das_url> (For example:
oracle.sysman.emSDK.sec.eus.DASHostUrl=http://my.dashost.com:7777 )
```

For example:

```
emctl set property -name oracle.sysman.emSDK.sec.DirectoryAuthenticationType
-value EnterpriseUser
```

3. Change directory to the `ORACLE_HOME/sysman/config` directory and open the `emoms.properties` file with your favorite text editor.

4. Add the following entries in the `emoms.properties` file:

```
oracle.sysman.emSDK.sec.DirectoryAuthenticationType=EnterpriseUser
oracle.sysman.emSDK.sec.eus.Domain=<ClientDomainName> (For example:
mydomain.com)
oracle.sysman.emSDK.sec.eus.DASHostUrl=<das_url> (For example:
oracle.sysman.emSDK.sec.eus.DASHostUrl=http://my.dashost.com:7777 )
```

5. Save and close the `emoms.properties` file.

6. Stop the Oracle Management Service.

See Also: [Controlling the Oracle Management Service](#) on page 7-4

7. Start the Management Service.

The next time you use the Oracle Enterprise Manager Grid Control Console to drill down to a managed database, Enterprise Manager will attempt to connect to the database using Enterprise User Security. If successful, Enterprise Manager will connect you to the database without displaying a login page. If the attempt to use Enterprise User Security fails, Enterprise Manager will prompt you for the database credentials.

2.2.3.1 Registering Enterprise Users as Enterprise Manager Users

After you have configured Enterprise Manager to use Enterprise Users, you can register existing enterprise users as Enterprise Manager Users and grant them the necessary privileges so that they can manage Enterprise Manager effectively.

You can register existing enterprise users by using:

- Enterprise Manager Grid Control Graphic User Interface
- Enterprise Manager Command Line Interface

2.2.3.1.1 Registering Enterprise Users Using the Graphical User Interface

You can use the graphical user interface to register enterprise users by following these steps:

1. Log into Enterprise Manager as a Super Administrator.

2. Click **Setup** and then click **Administrators** to display the Administrators page. Since Enterprise Manager has been configured to use Enterprise Users, the first page of the Create Administrator wizard will provide the option to create an administrator based on a registered Oracle Internet Directory user or a normal database user.
3. Select Oracle Internet Directory and click **Continue** to go to the next page in the wizard.
4. Enter the name and e-mail address of the Oracle Internet Directory user or click the flashlight icon to search for a user name in the Oracle Internet Directory.
5. Use the rest of the wizard pages to define the roles, system privileges, and other characteristics of the Enterprise Manager administrator and then click **Finish**. Enterprise Manager displays a summary page that lists the characteristics of the administrator account.
6. Click **Finish** to create the new Enterprise Manager administrator.
The OID user is now included in the list of Enterprise Manager administrators. You can now verify the account by logging out of the Grid Control Console and logging back in using the OID user credentials on the Single Sign-On logon page.

2.2.3.1.2 Registering Enterprise Users Using the Command Line Interface

To register Enterprise Users as Enterprise Manager users using EMCLI, enter the following command:

```
emcli create_user -name=eususer -type=DB_EXTERNAL_USER
```

This command registers the `eususer` as an Enterprise Manager user where `eususer` is an existing Enterprise User. For more details, refer to [Registering Single Sign-On Users Using EMCLI](#).

2.3 Enterprise Manager Authorization

System security is a major concern of any corporation. Giving the same level of access to all systems to all administrators is dangerous, but individually granting access to tens, hundreds, or even thousands of targets to every new member of the group is time consuming. With Enterprise Manager's administrator privileges and roles feature, this task can be performed within seconds, instead of hours. Authorization controls the access to the secure resources managed by Enterprise Manager via system, target, and object level privileges and roles.

This section describes Enterprise Manager's Authorization model including user classes, roles, and privileges assigned to each user class. The following topics are described:

- Classes of Users
- Privileges and Roles

2.3.1 Classes of Users

Oracle Enterprise Manager supports different classes of Oracle users, depending upon the environment you are managing and the context in which you are using Oracle Enterprise Manager Grid Control.

The Enterprise Manager administrators you create and manage in the Grid Control Console are granted privileges and roles to log in to the Grid Control Console and to manage specific target types and to perform specific management tasks. The default

super administrator for the Grid Control Console is the SYSMAN user, which is a database user associated with the Oracle Management Repository. You define the password for the SYSMAN account during the Enterprise Manager installation procedure.

By restricting access to privileged users and providing tools to secure communications between Oracle Enterprise Manager 11g components, Enterprise Manager protects critical information in the Oracle Management Repository.

The Management Repository contains management data that Enterprise Manager Grid Control uses to help you monitor the performance and availability of your entire enterprise. This data provides you with information about the types of hardware and software you have deployed, as well as the historical performance and specific characteristics of the applications, databases, applications servers, and other targets that you manage. The Management Repository also contains information about the Enterprise Manager administrators who have the privileges to access the management data.

You can create and manage Enterprise Manager administrator accounts. Each administrator account includes its own login credentials, as well as a set of roles and privileges that are assigned to the account. There are three administrator access categories:

- **Super Administrator:** Powerful Enterprise Manager administrator with full access privileges to all targets and administrator accounts within the Enterprise Manager environment. The Super Administrator, SYSMAN is created by default when Enterprise Manager is installed. The Super Administrator can create other administrator accounts.
- **Administrator:** Regular Enterprise Manager administrator.
- **Repository Owner:** Database administrator for the Management Repository. This account cannot be modified, duplicated, or deleted.

The types of management tasks that the administrator can perform and targets that he can access depends on the roles, system privileges, and target privileges that he is granted. The Super Administrator can choose to let certain administrators perform only certain management tasks, or access only certain targets, or perform certain management tasks on certain targets. In this way, the Super Administrator can divide the workload among his administrators.

2.3.2 Privileges and Roles

User privileges provide a basic level of security in Enterprise Manager. They are designed to control user access to data and to limit the kinds of SQL statements that users can execute. When creating a user, you grant privileges to enable the user to connect to the database, to run queries and make updates, to create schema objects, and more.

When Enterprise Manager is installed, the SYSMAN user (super administrator) is created by default. The SYSMAN Super Administrator then creates other administrator accounts for daily administration work. The SYSMAN account should only be used to perform infrequent system wide, global configuration tasks.

The Super Administrator divides workload among his administrators by filtering target access, or filtering access to management task, or both through the roles, System Privileges, and Target Privileges he grants them. For example, he can allow some administrators to view any target and to add any target in the enterprise and other administrators to only perform specific operations such as maintaining and cloning on a target for which they are responsible.

A role is a collection of Enterprise Manager system privileges, or target privileges, or both, which you can grant to administrators or to other roles. These roles can be based upon geographic location (for example, a role for Canadian administrators to manage Canadian systems), line of business (for example, a role for administrators of the human resource systems or the sales systems), or any other model. Administrators do not want to perform the task of individually granting access to tens, hundreds, or even thousands of targets to every new member of their group.

By creating roles, an administrator needs only to assign the role that includes all the appropriate privileges to his team members instead of having to grant many individual privileges. He can divide workload among his administrators by filtering target access, or filtering access to management task, or both.

Public Role: Enterprise Manager creates one role by default called **Public**. This role is unique in that it is automatically assigned to all new non-super administrators when they are created. By default it has no privileges assigned to it. The Public role should be used to define default privileges you expect to assign to a majority of non-super administrators you create. Privileges need not be assigned to Public initially - they can be added at any time. The role may be deleted if your enterprise does not wish to use it. If deleted, it can be added back in later if you later decide to implement it.

2.3.2.1 Granting Privileges

A privilege is a right to perform management actions within Enterprise Manager. Privileges can be divided into three categories:

- System Privileges
- Target Privileges
- Object Privileges

System Privileges: These privileges allow a user to perform system wide operations. To set the System Privileges, click the **Setup** link to navigate to the Setup Overview page and click on the **Administrators** option in the left panel. Select an administrator from the list and click **Edit**. The Edit Administrator wizard is displayed. Click **Next** to navigate through the wizard to see the System Privileges page:

Figure 2–5 System Privileges

ORACLE Enterprise Manager
Grid Control 11g

Enterprise Manager Configuration | Management Services and Repository | Agents

Home Targets Deployments Alerts Compliance Jobs Reports My Oracle Support

Authentication Properties Roles **System Privileges** Target Privileges Review

Create Administrator test: System Privileges

Cancel Back Step 4 of 6 Next Review

Select the System Privileges that you want to grant to this Enterprise Manager Administrator. System Privileges give the administrator the right to perform particular management actions.

Select	Name	Description
<input type="checkbox"/>	JVM Diagnostics Administrator	Able to manage JVM Diagnostics admin operations
<input type="checkbox"/>	JVM Diagnostics User	View JVM Diagnostics Data
<input type="checkbox"/>	Request Monitoring Administrator	Able to manage Request Monitoring admin operations
<input type="checkbox"/>	Request Monitoring User	View Request Monitoring data
<input type="checkbox"/>	Publish Report	Ability to publish reports for public viewing
<input type="checkbox"/>	Use any Beacon	Use any Beacon on any monitored host to monitor transactions, URLs, and network components. Beacon is installed with the Oracle Agent.
<input type="checkbox"/>	Add any Target	Add any target in Enterprise Manager
<input type="checkbox"/>	View any Target	View all discovered targets in the environment (including Agents, and Management Service and Repository targets). This system privilege automatically includes the MONITOR ENTERPRISE MANAGER system privilege.
<input type="checkbox"/>	Create Privilege Propagating Group	Ability to create privilege propagating groups. Privileges granted on a privilege propagating group will be automatically granted on the members of the group
<input type="checkbox"/>	Monitor Enterprise Manager	Monitor Enterprise Manager performance. Performance is determined by how efficiently the Management Services and their components are running and processing, and how well the DBMS jobs handling the maintenance and monitoring of Enterprise Manager are running.

Cancel Back Step 4 of 6 Next Review

Table 2–1 System Privileges

System Privilege	Description
USE ANY BEACON	Allows the administrator to use any Beacon on any monitored host to monitor transactions, URLs, and network components.
ADD ANY TARGET	Allows the administrator to add any target to Enterprise Manager for monitoring, administration and management.
VIEW ANY TARGET	Allows the administrator to view any target on the system, including Oracle Management Agents and Management Services. Whenever the <code>VIEW ANY TARGET</code> privilege is granted, the <code>MONITOR ENTERPRISE MANAGER</code> privilege is also granted by default.
CREATE PRIVILEGE PROPAGATING GROUP	Allows the administrator to create privilege propagating groups. Privileges granted to such groups will be automatically granted to all members of the group.
MONITOR ENTERPRISE MANAGER	Allows the administrator to monitor the availability and performance of Enterprise Manager itself, and grants the administrator access to the following targets: the database used for the Management Repository, the Management Service and Management Repository, and all Oracle Management Agents in the global enterprise.
PUBLISH REPORT	Allows the administrator to publish reports for public use.
JVM Diagnostics Administrator	Allows the administrator to manage JVM Diagnostics operations.
JVM Diagnostics User	Allows the user to view JVM Diagnostics data.
Request Monitoring Administrator	Allows the user to manage E2E administrator operations.
Request Monitoring User	Allows the user to view E2E data.

Select the check box to select the system privilege to be granted to the administrator and click **Next**. The Target Privileges page is displayed.

Target Privileges: These privileges allow an administrator to perform operations on a target. The Target Privileges page shows a list of targets for which privileges can be granted. Select a target from the list and click the pencil icon in the Privilege column. The following screen is displayed.

Figure 2–6 Target Privileges

ORACLE Enterprise Manager
Grid Control 11g

Home Targets Deployments Alerts Compliance Jobs Reports My Oracle Support

Properties Roles System Privileges Target Privileges Review

Create Administrator test: Target Privileges

Cancel Continue

Select the Target Privileges that you want to grant to this Enterprise Manager administrator. Target Privileges give the administrator the right to perform particular actions on targets selected on the previous page.

Full
Ability to do all operations on the target, including delete the target

Operator
Ability to do normal administrative operations on the target, such as configure a blackout and edit the target properties

Blackout Target
Ability to create, edit, schedule and stop a blackout on the target

Manage Target Metrics
Ability to edit threshold for metric and policy setting, apply monitoring templates, and manage User Defined Metrics

Configure Target
Ability to edit target properties and modify monitoring configuration

Manage Target Alerts
Ability to clear stateless alerts, manually re-evaluate alerts and acknowledge alerts for the target

View
Ability to view properties, inventory and monitor information about a target

Cancel Continue

Select the check box to specify the privileges that are to be granted and click **Continue**. For more details on setting these privileges, see the Enterprise Manager Online Help.

Note: If a target has certain privileges, *** need info ***

Table 2–2 Target Privileges

Target Privilege	Description
FULL	Implicitly grants all the target privileges and allows the administrator to delete the target from the Enterprise Manager system.
OPERATOR	Allows the administrator to perform normal administrative operations on a target such as configure a blackout, or edit the properties.
BLACKOUT TARGET	Allows the administrator to create, edit, schedule, and stop blackout on a target.
MANAGE TARGET ALERTS	Allows the administrator to clear stateless alerts, manually re-evaluate alerts and acknowledge alerts for the target.
CONFIGURE TARGET	Allows the administrator to edit target properties and modify monitoring configurations.
MANAGE TARGET METRICS	Allows the administrator to edit thresholds for metric and policy settings, apply monitoring templates and manage user defined metrics.
VIEW	Allows the administrator to view properties, inventory and monitor information about a target.

Object Privileges: These privileges allow an administrator to perform a particular action on a specific schema object. Different object privileges are available for different types of schema objects.

Table 2–3 Object Privileges

Target Privilege	Description
VIEW JOB	Provides the administrator with the ability to view the job and its definition.
FULL JOB	Provides the administrator with the ability to view, edit, submit, and delete the job.
VIEW REPORT	Provides the administrator with the ability to view the report.
VIEW TEMPLATE	The ability to view the template definition.
FULL TEMPLATE	The ability to edit the template definition.

2.4 Configuring Security for Grid Control

This section contains the following topics:

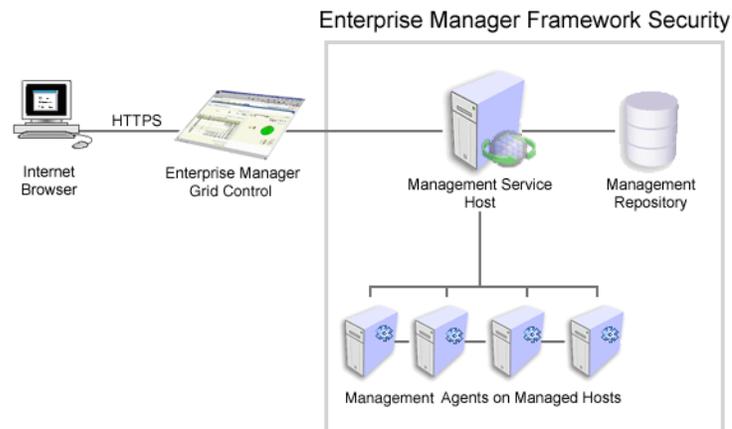
- [About Enterprise Manager Framework Security](#)
- [Overview of the Steps Required to Enable Enterprise Manager Framework Security](#)
- [Enabling Security for the Oracle Management Service](#)
- [Enabling Security for the Oracle Management Agent](#)
- [Enabling Security with Multiple Management Service Installations](#)
- [Restricting HTTP Access to the Management Service](#)
- [Managing Agent Registration Passwords](#)
- [Enabling Security with a Server Load Balancer](#)
- [Enabling Security for the Management Repository Database](#)

2.4.1 About Enterprise Manager Framework Security

Enterprise Manager Framework Security provides safe and secure communication channels between the components of Enterprise Manager. For example, Framework Security provides secure connections between your Oracle Management Service and its Management Agents.

See Also: *Oracle Enterprise Manager Concepts* for an overview of Enterprise Manager components

The following figure shows how Enterprise Manager Framework Security provides security for the connections between the Enterprise Manager components.

Figure 2–7 Enterprise Manager Framework Security

Enterprise Manager Framework Security implements the following types of secure connections between the Enterprise Manager components:

- HTTPS and Public Key Infrastructure (PKI) components, including signed digital certificates, for communications between the Management Service and the Management Agents.

See Also: *Oracle Security Overview* for an overview of Public Key Infrastructure features, such as digital certificates and public keys

- Oracle Advanced Security for communications between the Management Service and the Management Repository.

See Also: *Oracle Database Advanced Security Administrator's Guide*

2.4.2 Overview of the Steps Required to Enable Enterprise Manager Framework Security

To enable Enterprise Manager Framework Security, you must configure each of the Enterprise Manager components in a specific order. The following list outlines the process for securing the Management Service and the Management Agents that upload data to the Management Service:

Note: The Enterprise Manager components are configured during installation. You can use the following commands if you want to reconfigure any of the components.

1. Use the `emctl stop oms` command to stop the OMS and the WebTier.
2. Use `emctl secure oms` to enable security for the Management Service.
3. Restart the OMS and the WebTier using the `emctl start oms` command.
4. For each Management Agent, stop the Management Agent, use the `emctl secure agent` command to enable security for the Management Agent, and restart the Management Agent.

5. After security is enabled for all the Management Agents, use the `emctl secure lock` command to restrict HTTP Access to the Management Service. This will ensure that Management Agents for which security has not been enabled will not be able upload data to the Management Service.

The following sections describe how to perform each of these steps in more detail.

Note: To resolve errors from `emctl secure` operations, refer to `EM_INSTANCE_HOME/sysman/log/secure.log` for more details.

2.4.3 Enabling Security for the Oracle Management Service

To enable Enterprise Manager Framework Security for the Management Service, you use the `emctl secure oms` utility, which is located in the following subdirectory of the Management Service home directory:

```
ORACLE_HOME/bin
```

The `emctl secure oms` utility performs the following actions:

- Generates a Root Key within your Management Repository. The Root Key is used during distribution of Oracle Wallets containing unique digital certificates for your Management Agents.
- Modifies your WebTier to enable an HTTPS channel between your Management Service and Management Agents, independent from any existing HTTPS configuration that may be present in your WebTier.
- Enables your Management Service to accept requests from Management Agents using Enterprise Manager Framework Security.

To run the `emctl secure oms` utility you must first choose an Agent Registration Password. The Agent Registration password is used to validate that future installation sessions of Oracle Management Agents and Oracle Management Services are authorized to load their data into this Enterprise Manager installation.

To enable Enterprise Manager Framework Security for the Oracle Management Service:

1. Stop the Management Service, the WebTier, and the other application server components using the following command:

```
OMS_ORACLE_HOME/bin/emctl stop oms
```

2. Enter the following command:

```
OMS_ORACLE_HOME/bin/emctl secure oms
```

3. You will be prompted for the Enterprise Manager Root Password. Enter the `SYSMAN` password.
4. You will be prompted for the Agent Registration Password, which is the password required for any Management Agent attempting to secure with the Management Service. Specify an Agent Registration Password for the Management Service.
5. When the operation is complete, restart the WebLogic Server and the deployed Enterprise Manager application.

```
OMS_ORACLE_HOME/bin/emctl start oms
```

6. After the Management Service restarts, test the secure connection to the Management Service by browsing to the following secure URL using the HTTPS protocol:

```
https://hostname.domain:https_upload_port/em
```

For example:

```
https://mgmthost1.acme.com:1159/em
```

If the Management Service security has been enabled, your browser displays the Enterprise Manager Login page.

Note: The 1159 port number is the default secure port used by the Management Agents to upload data to the Management Service. This port number may vary if the default port is unavailable.

Example 2-2 Sample Output of the `emctl secure oms` Command

```
emctl secure oms
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
Securing OMS... Started.
Securing OMS... Successful
```

Alternatively, you can enter the `emctl secure oms` command all on one line, but if you enter the command on one line, the passwords you enter will be displayed on the screen as you type the command.

Example 2-3 Usage of the `emctl secure oms` Command (II)

```
emctl secure oms [-sysman_pwd <sysman password>] [-reg_pwd <registration
password>] [-host <hostname>] [-slb_port <slb port>] [-slb_console_port <slb
console port>] [-reset] [-console] [-lock] [-lock_console] [-secure_port <secure_
port>] [-upload_http_port <upload_http_port>] [-root_dc <root_dc>] [-root_country
<root_country>] [-root_email <root_email>] [-root_state <root_state>] [-root_loc
<root_loc>] [-root_org <root_org>] [-root_unit <root_unit>] [-wallet <wallet_loc>
-trust_certs_loc <certs_loc>] [-wallet_pwd <pwd>] [-key_strength <strength>]
[-cert_validity <validity>] [-protocol <protocol>]
Valid values for <protocol> are the allowed values for Apache's SSLProtocol
directive
```

The parameters are explained below:

- `sysman_pwd` - Oracle Management Repository user password.
- `reg_pwd` - The Management Agent registration password.
- `host` - The host name to be used in the certificate used by the Oracle Management Service. You may need to use the SLB host name if there is an SLB before the Management Service.
- `reset` - A new certificate authority will be created. All the Agents and Oracle Management Services need to be resecured.
- `secure_port` - The port to be used for secure communication.
- `upload_http_port` - The port used for unsecure upload communications.
- `slb_port` - This parameter is required when Server Load Balancer is used. It specifies the secure upload port configured in the Server Load Balancer.

- `slb_console_port` - This parameter is required when Server Load Balancer is used. It specifies the secure console port configured in the Server Load Balancer.
- `root_dc` - The domain component used in the root certificate. The default value is `com`.
- `root_country` - The country to be used in the root certificate. The default value is `US`.
- `root_state` - The state to be used in the root certificate. The default value is `CA`.
- `root_loc` - The location to be used in the root certificate. The default value is **EnterpriseManager on <hostname>**.
- `root_org` - The organization name to be used in the root certificate. The default value is EnterpriseManager on <hostname>.
- `root_unit` - The organizational unit to be used in the root certificate. The default value is EnterpriseManager on <hostname>.
- `root_email` - The email address to be used in the root certificate. The default value is EnterpriseManager@<hostname>.
- `wallet`: This is the location of the wallet containing the third party certificate. This parameter should be specified while configuring third party certificates.
- `trust_certs_loc` - The location of the `trusted_certs.txt` (required when third party certificates are used).
- `key_strength`: The strength of the key to be used. Valid values are 512, 1024, 2048, and 4096.
- `cert_validity`: The number of days for which the self-signed certificate is valid. The valid range is between 1 to 3650.
- `protocol`: This parameter is used to configure Oracle Management Service in TLSv1-only or SSLv3-only or mixed mode (default). Valid values are the allowed values as per **Apache's SSLProtocol** directive.

Note: The `key_strength` and `cert_validity` parameters are applicable only when the `-wallet` option is not used.

2.4.3.1 Checking the Security Status

You can check whether security has been enabled for the Management Service by entering the `emctl status oms -secure` command.

Example 2-4 Sample Output of the `emctl status oms -details` Command

```
emctl status oms -details [-sysman_pwd <pwd>]
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
Enter Enterprise Manager Root (SYSMAN) Password :
Console Server Host : omshost.mydomain.com
HTTP Console Port   : 7788
HTTPS Console Port  : 7799
HTTP Upload Port    : 4889
HTTPS Upload Port   : 1159
OMS is not configured with SLB or virtual hostname
Agent Upload is locked.
OMS Console is locked.
Active CA ID: 1
```

2.4.3.2 Creating a New Certificate Authority

You may need to create a new Certificate Authority (CA) if the current CA is expiring or if you want to change the key strength. A unique identifier is assigned to each CA. For instance, the CA created during installation may have an identifier as ID 1, subsequent CAs will have the IDs 2,3, and so on. At any given time, the last created CA is active and issues certificates for OMSs and Agents.

Example 2-5 Creating a New Certificate Authority

```
emctl secure createca [-sysman_pwd <pwd>] [-host <hostname>] [-key_
strength<strength>] [-cert_validity <validity>] [-root_dc <root_dc>] [-root_
country <root_country>] [-root_email <root_email>] [-root_state <root_state>]
[-root_loc <root_loc>] [-root_org <root_org>] [-root_unit <root_unit>]
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
Creating CA... Started.
Successfully created CA with ID 2
```

Example 2-6 Viewing Information about a Certificate Authority

```
emcli get_ca_info -ca_id="1;2" -details
Info about CA with ID: 1
CA is not configured
DN: CN=myhost.mydomain.com, C=US
Serial# : 3423643907115516586
Valid From: Tue Mar 16 11:06:20 PDT 2010
Valid Till: Sat Mar 14 11:06:20 PDT 2020
Number of Agents registered with CA ID 1 is 1
myhost.mydomain.com:3872

Info about CA with ID: 2
CA is configured
DN: CN=myhost.mydomain.com, C=US, ST=CA
Serial# : 1182646629511862286
Valid From: Fri Mar 19 05:17:15 PDT 2010
Valid Till: Tue Mar 17 05:17:15 PDT 2020
There are no Agents registered with CA ID 2
```

Note: The WebLogic Administrator and Node Manager passwords are stored in the Administration Credentials Wallet. This is present in the `EM_INSTANCE_HOME/sysman/config/adminCredsWallet` directory. To recreate Administrator Credentials wallet, run the following command on each machine on which the Management Service is running:

```
emctl secure create_admin_creds_wallet [-admin_pwd
<pwd>] [-nodemgr_pwd <pwd>]
```

2.4.3.3 Viewing the Security Status and OMS Port Information

To view the security status and OMS port information, use the following command

Example 2-7 emctl status oms -details

```
$ emctl status oms -details [-sysman_pwd welcome1]
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
Console Server Host : myhost.mydomain.com
HTTP Console Port : 7788
HTTPS Console Port : 7799
```

```

HTTP Upload Port    : 4889
HTTPS Upload Port   : 1159
OMS is not configured with SLB or virtual hostname
Agent Upload is locked.
OMS Console is locked.
Active CA ID: 1
    
```

2.4.3.4 Configuring Transparent Layer Security

The Oracle Management Service can be configured in the following modes:

- **TLSv1-only mode:** To configure the OMS to use only TLSv1 connections, do the following:

1. Stop the OMS by entering the following command:

```
OMS_ORACLE_HOME/bin/emctl stop oms
```

2. Enter the following command:

```
emctl secure oms -protocol TLSv1
```

3. Append `-Dweblogic.security.SSL.protocolVersion=TLS1` to `JAVA_OPTIONS` in `Domain_Home/bin/startEMServer.sh`. If this property already exists, update the value to `TLS1`.

4. Restart the OMS with the following command:

```
OMS_ORACLE_HOME/bin/emctl start oms
```

Note: If the OMS is configured in the TLSv1 only mode, the 10.2.x Agents cannot communicate with the OMS since those Agents do not support the TLS mode.

- **SSLv3 Only Mode:** To configure the OMS to use SSLv3 connections only, do the following:

1. Stop the OMS by entering the following command:

```
OMS_ORACLE_HOME/bin/emctl stop oms
```

2. Enter the following command:

```
emctl secure oms -protocol SSLv3
```

3. Append `-Dweblogic.security.SSL.protocolVersion=SSL3` to `JAVA_OPTIONS` in `Domain_Home/bin/startEMServer.sh` or `startEMServer.cmd` on Windows. If this property already exists, update the value to `SSL3`.

4. Restart the OMS with the following command:

```
OMS_ORACLE_HOME/bin/emctl start oms
```

- **Mixed Mode:** To configure the OMS to use both SSLv3 and TLSv1 connections, do the following:

1. Stop the OMS by entering the following command:

```
OMS_ORACLE_HOME/bin/emctl stop oms
```

2. Enter the following command:

```
emctl secure oms
```

3. Append `-Dweblogic.security.SSL.protocolVersion=ALL` to `JAVA_OPTIONS` in `Domain_Home/bin/startEMServer.sh`. If this property already exists, update the value to `ALL`.
4. Restart the OMS with the following command:

```
OMS_ORACLE_HOME/bin/emctl start oms
```

Note: By default, the OMS is configured to use the Mixed Mode. To configure the Management Agent in TLSv1 only mode, set `allowTLSOnly=true` in the `emd.properties` file and restart the Agent.

2.4.4 Enabling Security for the Oracle Management Agent

When you install the Management Agent on a host, you must identify the Management Service that will be used by the Management Agent. If the Management Service you specify has been configured to take advantage of Enterprise Manager Framework Security, you will be prompted for the Agent Registration Password and Enterprise Manager Framework Security will be enabled for the Management Agent during the installation.

Otherwise, if the Management Service has not been configured for Enterprise Manager Framework Security or if the Registration Password was not specified during installation, then security will not be enabled for the Management Agent. In those cases, you can later enable Enterprise Manager Framework Security for the Management Agent.

To enable Enterprise Manager Framework Security for the Management Agent, use the `emctl secure agent` utility, which is located in the following directory of the Management Agent home directory:

```
AGENT_HOME/bin (UNIX)
AGENT_HOME\bin (Windows)
```

The `emctl secure agent` utility performs the following actions:

- Obtains an Oracle Wallet from the Management Service that contains a unique digital certificate for the Management Agent. This certificate is required in order for the Management Agent to conduct SSL communication with the secure Management Service.
- Obtains an Agent Key for the Management Agent that is registered with the Management Service.
- Configures the Management Agent so it is available on your network over HTTPS and so it uses the Management Service HTTPS upload URL for all its communication with the Management Service.

To enable Enterprise Manager Framework Security for the Management Agent:

1. Ensure that your Management Service and the Management Repository are up and running.
2. Change directory to the following directory:

```
AGENT_HOME/bin (UNIX)
AGENT_HOME\bin (Windows)
```

3. Stop the Management Agent:

```
emctl stop agent
```

4. Enter the following command:

```
emctl secure agent (UNIX)
emctl secure agent (Windows)
```

The `emctl secure agent` utility prompts you for the Agent Registration Password, authenticates the password against the Management Service, and reconfigures the Management Agent to use Enterprise Manager Framework Security.

Note: Alternatively, you can enter the command all on one line, but if you enter the command on one line, the password you enter will be displayed on the screen as you type:

```
emctl secure agent agent_registration_pwd (UNIX)
emctl secure agent agent_registration_pwd (Windows)
```

[Example 2-8](#) shows sample output of the `emctl secure agent` utility.

5. Restart the Management Agent:

```
emctl start agent
```

6. Confirm that the Management Agent is secure by checking the Management Agent home page.

Note: You can also check if the Agent Management is secure by running the `emctl status agent -secure` command, or by checking the Agent and Repository URLs in the output of the `emctl status agent` command.

In the General section of the Management Agent home page ([Figure 2-8](#)), the **Secure Upload** field indicates whether or not Enterprise Manager Framework Security has been enabled for the Management Agent.

See Also: "Checking the Status of an Oracle Management Agent" in the Enterprise Manager online Help

Example 2-8 Sample Output of the emctl secure agent Utility

```
emctl secure agent
Oracle Enterprise Manager 11g Release 1 Grid Control.
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
Securing agent... Started
Securing agent... Successful.
```

Example 2-9 Sample Output of the emctl status agent secure Command

```
emctl status agent -secure
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
Checking the security status of the Agent at location set in
/private/home/oracle/product/102/em/agent10g/sysman/config/emd.properties...
```

```

Done.
Agent is secure at HTTPS Port 3872.
Checking the security status of the OMS at
http://gridcontrol.oraclecorp.com:4889/em/upload/... Done.
OMS is secure on HTTPS Port 4888
    
```

Figure 2–8 Secure Upload Field on the Management Agent Home Page

The screenshot shows the Oracle Enterprise Manager Grid Control 11g interface. The breadcrumb trail is: Alerts > Agent: adc2100768.us.oracle.com:3872. The page title is "Agent: adc2100768.us.oracle.com:3872". The "Latest Data Collected From Target" is "Mar 2, 2010 6:00:57 AM GMT".

General

Status	Up
Host	adc2100768.us.oracle.com
Availability (%)	100.00
Number of Restarts (last 24 hours)	1
Management Service	smpstst-svc5.us.oracle.com:1159
Agent to Management Service Response Time (ms)	0
Version	11.1.0.1.0
Operating System Owner	oracle
Oracle Home	/scratch/oracle/agent100226/agent11g
Agent State Directory	/scratch/oracle/agent100226/agent11g
Agent Heartbeat Interval (seconds)	60

Resource Utilization

CPU Usage (%)	0.00
Virtual Memory Usage (MB)	147
Regular Files Open	9
Threads Created	6

Upload

Secure Upload	Yes
Last Successful Upload	Mar 2, 2010 5:53:40 AM
Data Pending Upload	0 MB in 0 Files
Uploaded data (KB past hour)	0.00

Monitored Targets

If you want to modify target properties, press the Configure button. Depending on the target type, you may be able to update credentials, or choose monitoring levels if monitoring has been defined for that target type.

Add Database Instance [v] Go

Select	Name	Type
<input checked="" type="checkbox"/>	adc2100768.us.oracle.com	Host

Alerts

Metric Collection Errors 1

Metric Name	Severity	Message	Alert Triggered Value	Last Checked
No current alerts or warnings				

2.4.5 Enabling Security with Multiple Management Service Installations

If you already have a secure Management Service running and you install an additional Management Service that uses the same Management Repository, you will need to enable Enterprise Manager Framework Security for the new Management Service. This task is executed using the same procedure that you used to secure the first Management Service, by running the `emctl secure oms` utility.

Because you have already established at least one Agent Registration Password and a Root Key in your Management Repository, they must be used for your new Management Service. Your secure Management Agents can then operate against either Management Service.

All the registration passwords assigned to the current Management Repository are listed on the Registration Passwords page in the Oracle Enterprise Manager 11g Grid Control Console.

If you install a new Management Service that uses a new Management Repository, the new Management Service is considered to be a distinct enterprise. There is no way for the new Management Service to partake in the same security trust relationship as another Management Service that uses a different Management Repository. Secure Management Agents of one Management Service will not be able to operate against the other Management Service.

2.4.6 Restricting HTTP Access to the Management Service

By default, when you enable Enterprise Manager Framework Security on your Oracle Management Service there are no default restrictions on HTTP access. The Grid Control Console can also be accessed over HTTP and the Oracle Management Agents will be able to upload over HTTP as well as HTTPS.

However, it is important that only secure Management Agent installations that use the Management Service HTTPS channel are able to upload data to your Management Repository and Grid Control console is accessible via HTTPS only.

To restrict access so Management Agents can upload data to the Management Service only over HTTPS:

1. Stop the Management Service, the WebTier, and the other application server components:

```
cd ORACLE_HOME/opmn/bin
emctl stop oms
```

2. Change directory to the following location in the Management Service home:

```
ORACLE_HOME/bin
```

3. Enter the following command to prevent Management Agents from uploading data to the Management Service over HTTP:

```
emctl secure lock -upload
```

Note:

- To lock the console and prevent HTTP access to the console, enter the following command:

```
emctl secure lock -console
```

- To lock both, enter either of the following commands:

```
emctl secure lock or
emctl secure lock -upload -console
```

- To lock both the console access and uploads from Agents while enabling security on the Management Service, enter the following command:

```
emctl secure oms -lock [other options]
```

4. Restart the Management Service, the WebTier, and the other application server components:

```
cd ORACLE_HOME/opmn/bin
emctl start oms
```

5. Verify that you cannot access the Management Agent upload URL using the HTTP protocol:

For example, navigate to the following URL:

```
http://hostname.domain:4889/em/upload
```

You should receive an error message similar to the following:

```
Forbidden
```

You are not authorised to access this resource on the server.

6. Verify that you can access the Management Agent Upload URL using the HTTPS protocol:

For example, navigate to the following URL:

```
https://hostname.domain:4888/em/upload
```

You should receive the following message, which confirms the secure upload port is available to secure Management Agents:

```
Http XML File receiver
Http Receiver Servlet active!
```

To allow the Management Service to accept uploads from unsecure Management Agents, use the following command:

```
emctl secure unlock -upload
```

Note:

- To unlock the console and allow HTTP access to the console, enter the following command:

```
emctl secure unlock -console
```

- To unlock both, enter either of the following command:

```
emctl secure unlock
emctl secur unlock -console -upload
```

Example 2–10 Sample Output of the emctl secure lock Command

```
emctl secure lock
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
OMS Console is locked. Access the console over HTTPS ports.
Agent Upload is locked. Agents must be secure and upload over HTTPS port.
```

Example 2–11 Sample Output of the emctl secure unlock Command

```
emctl secure unlock
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
OMS Console is unlocked. HTTP ports too can be used to access console.
Agent Upload is unlocked. Unsecure Agents may upload over HTTP.
```

To restrict HTTP access to the Oracle Enterprise Manager 11g Grid Control Console, use the `emctl secure lock -console` command.

2.4.7 Managing Agent Registration Passwords

Enterprise Manager uses the Agent Registration password to validate that installations of Oracle Management Agents are authorized to load their data into the Oracle Management Service.

The Agent Registration password is created during installation when security is enabled for the Oracle Management Service.

Note: To avoid new Agents from being registered with the Oracle Management Service, you must delete all registration passwords.

2.4.7.1 Using the Grid Control Console to Manage Agent Registration Passwords

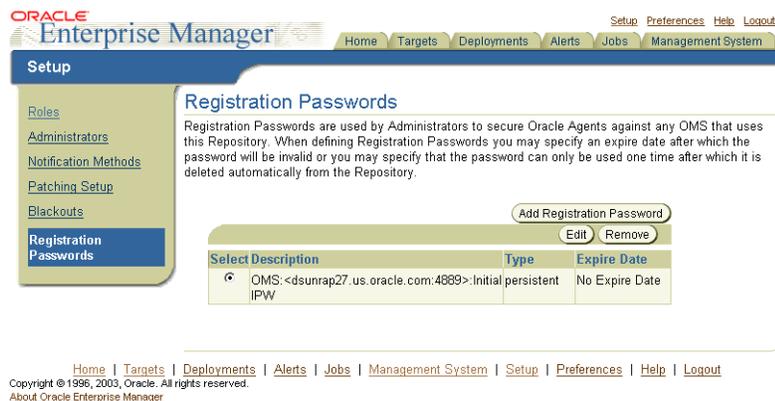
You can use the Grid Control Console to manage your existing registration passwords or create additional registration passwords:

1. Click **Setup** at the top of any Grid Control Console page.
2. Click **Registration Passwords**.

Enterprise Manager displays the Registration Passwords page (Figure 2–9). The registration password you created when you ran the `emctl secure oms` command appears in the Registration Passwords table.

3. Use the Registration Passwords page to change your registration password, create additional registration passwords, or remove registration passwords associated with the current Management Repository.

Figure 2–9 Managing Registration Passwords in the Grid Control Console



When you create or edit an Agent Registration Password on the Registration Passwords page, you can determine whether the password is persistent and available for multiple Management Agents or to be used only once or for a predefined period of time.

For example, if an administrator requests to install a Management Agent on a particular host, you can create a one-time-only password that the administrator can use to install and configure one Management Agent.

On the other hand, you can create a persistent password that an administrator can use for the next two weeks before it expires and the administrator must ask for a new password.

2.4.7.2 Using emctl to Add a New Agent Registration Password

To add a new Agent Registration Password, use the following `emctl` command on the machine on which the Management Service has been installed:

```
emctl secure setpwd [-sysman_pwd] [new registration pwd]
```

The `emctl secure setpwd` command requires that you provide the password of the Enterprise Manager super administrator user, `sysman`, to authorize the addition of the Agent Registration Password.

If you change the Agent Registration Password, you must communicate the new password to other Enterprise Manager administrators who need to install new Management Agents, enable Enterprise Manager Framework Security for existing Management Agents, or install additional Management Services.

As with other security passwords, you should change the Agent Registration Password on a regular and frequent basis to prevent it from becoming too widespread.

2.4.8 Enabling Security with a Server Load Balancer

When you deploy a Management Service that is available behind a Server Load Balancer (SLB), special attention must be given to the DNS host name over which the Management Service will be available. Although the Management Service may run on a particular local host, for example `myhost.mycompany.com`, your Management Agents will access the Management Service using the host name that has been assigned to the Server Load Balancer. For example, `oracleoms.mycompany.com`.

As a result, when you enable Enterprise Manager Framework Security for the Management Service, it is important to ensure that the Server Load Balancer host name is embedded into the Certificate that the Management Service uses for SSL communications. To do so, enter the following commands:

This may be done by using `emctl secure oms` and specifying the host name in the with an extra `-host` parameter as follows:

- Enable security on the Management Service by entering the following command:


```
emctl secure oms -host <slb_hostname> [-slb_console_port <slb UI port>] [-slb_port <slb upload port>] [other params]
```
- Create virtual servers and pools on the Server Load Balancer.
- Verify that the console can be accessed using the following URL:


```
https://slbhost:slb_console_port/em
```
- Re-secure the Agents with Server Load Balancer by using the following command:


```
emctl secure agent -emdWalletSrcUrl <SLB Upload or UI URL>
```

For example:

```
Agent_Home/bin/emctl secure agent -emdWalletSrcUrl
https://slbost:slb_upload_port/em https://slbost:slb_upload_
port/em
```

2.4.9 Enabling Security for the Management Repository Database

This section describes how to enable Security for the Oracle Management Repository. This section includes the following topics:

- [About Oracle Advanced Security and the `sqlnet.ora` Configuration File](#)
- [Configuring the Management Service to Connect to a Secure Management Repository Database](#)
- [Enabling Oracle Advanced Security for the Management Repository](#)
- [Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database](#)

2.4.9.1 About Oracle Advanced Security and the sqlnet.ora Configuration File

You enable security for the Management Repository by using Oracle Advanced Security. Oracle Advanced Security ensures the security of data transferred to and from an Oracle database.

See Also: *Oracle Database Advanced Security Administrator's Guide*

To enable Oracle Advanced Security for the Management Repository database, you must make modifications to the `sqlnet.ora` configuration file. The `sqlnet.ora` configuration file is used to define various database connection properties, including Oracle Advanced Security parameters.

The `sqlnet.ora` file is located in the following subdirectory of the Database home:

```
ORACLE_HOME/network/admin
```

After you have enabled Security for the Management Repository and the Management Services that communicate with the Management Repository, you must also configure Oracle Advanced Security for the Management Agent by modifying the `sqlnet.ora` configuration file in the Management Agent home directory.

See Also: ["Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database"](#)

It is important that both the Management Service and the Management Repository are configured to use Oracle Advanced Security. Otherwise, errors will occur when the Management Service attempts to connect to the Management Repository. For example, the Management Service might receive the following error:

```
ORA-12645: Parameter does not exist
```

To correct this problem, be sure both the Management Service and the Management Repository are configured as described in the following sections.

Note: The procedures in this section describe how to manually modify the `sqlnet.ora` configuration file to enable Oracle Advanced Security. Alternatively, you can make these modifications using the administration tools described in the *Oracle Database Advanced Security Administrator's Guide*.

2.4.9.2 Configuring the Management Service to Connect to a Secure Management Repository Database

If you have enabled Oracle Advanced Security for the Management Service database—or if you plan to enable Oracle Advanced Security for the Management Repository database—use the following procedure to enable Oracle Advanced Security for the Management Service:

1. Stop the Management Service:

```
ORACLE_HOME/bin/emctl stop oms
```

2. Set the `emoms.properties` by using the `emctl set property` command
3. Restart the Management Service.

```
ORACLE_HOME/bin/emctl start oms
```

Table 2–4 Oracle Advanced Security Properties in the Enterprise Manager Properties File

Property	Description
oracle.sysman.emRep.dbConn.enableEncryption	<p>Defines whether or not Enterprise Manager will use encryption between Management Service and Management Repository.</p> <p>Possible values are TRUE and FALSE. The default value is TRUE.</p> <p>For example:</p> <pre>oracle.sysman.emRep.dbConn.enableEncryption=true</pre>
oracle.net.encryption_client	<p>Defines the Management Service encryption requirement.</p> <p>Possible values are REJECTED, ACCEPTED, REQUESTED, and REQUIRED.</p> <p>The default value is REQUESTED. In other words, if the database supports secure connections, then the Management Service uses secure connections, otherwise the Management Service uses insecure connections.</p> <p>For example:</p> <pre>oracle.net.encryption_client=REQUESTED</pre>
oracle.net.encryption_types_client	<p>Defines the different types of encryption algorithms the client supports.</p> <p>Possible values should be listed within parenthesis. The default value is (DES40C).</p> <p>For example:</p> <pre>oracle.net.encryption_types_client=(DES40C)</pre>
oracle.net.crypto_checksum_client	<p>Defines the Client's checksum requirements.</p> <p>Possible values are REJECTED, ACCEPTED, REQUESTED, and REQUIRED.</p> <p>The default value is REQUESTED. In other words, if the server supports checksum enabled connections, then the Management Service uses them, otherwise it uses normal connections.</p> <p>For example:</p> <pre>oracle.net.crypto_checksum_client=REQUESTED</pre>
oracle.net.crypto_checksum_types_client	<p>This property defines the different types of checksums algorithms the client supports.</p> <p>Possible values should be listed within parentheses. The default value is (MD5).</p> <p>For example:</p> <pre>oracle.net.crypto_checksum_types_client=(MD5)</pre>

2.4.9.3 Enabling Oracle Advanced Security for the Management Repository

To be sure your database is secure and that only encrypted data is transferred between your database server and other sources, review the security documentation available in the Oracle Database documentation library.

See Also: *Oracle Database Advanced Security Administrator's Guide*

The following instructions provide an example of how you can confirm that Oracle Advanced Security is enabled for your Management Repository database and its connections with the Management Service:

1. Locate the `sqlnet.ora` configuration file in the following directory of the database Oracle Home:

```
ORACLE_HOME/network/admin
```

2. Using a text editor, look for the following entries (or similar entries) in the `sqlnet.ora` file:

```
SQLNET.ENCRYPTION_SERVER = REQUESTED  
SQLNET.CRYPTO_SEED = "abcdefg123456789"
```

See Also: "Configuring Network Data Encryption and Integrity for Oracle Servers and Clients in the Oracle Application Server 10g Administrator's Guide"

3. Save your changes and exit the text editor.

2.4.9.4 Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database

After you have enabled Oracle Advanced Security for the Management Repository, you must also enable Advanced Security for the Management Agent that is monitoring the Management Repository:

1. Locate the `sqlnet.ora` configuration file in the following directory inside the home directory for the Management Agent that is monitoring the Management Repository:

```
AGENT_HOME/network/admin (UNIX)  
AGENT_HOME\network\admin (Windows)
```

2. Using a text editor, add the following entry to the `sqlnet.ora` configuration file:

```
SQLNET.CRYPTO_SEED = "abcdefg123456789"
```

The `SQLNET.CRYPTO_SEED` can be any string between 10 to 70 characters.

See Also: "Configuring Network Data Encryption and Integrity for Oracle Servers and Clients in the Oracle Application Server Administrator's Guide"

3. Save your changes and exit the text editor.
4. Restart the Management Agent.

2.4.10 Configuring Third Party Certificates

You can configure third party certificates for:

- HTTPS Upload Virtual Host
- HTTPS Console Virtual Host

Note: Only Single Sign-On wallets are supported.

2.4.10.1 Configuring Third Party Certificate for HTTPS Upload Virtual Host

You can configure the third party certificate for the HTTPS Upload Virtual Host in two ways:

Method 1

1. Create a wallet for each OMS in the grid.
2. While creating the wallet, specify the host name of the machine where the OMS is installed or the Load Balancer Name if the OMS is behind the Load Balancer for Common Name.
3. Write the certificates of all the Certificate Authorities in the certificate chain (like the Root Certificate Authority, Intermediate Certificate Authority) into a file named `trusted_certs.txt`.
4. Download or copy the `trusted_certs.txt` file to the host machines on which each Agent that is communicating with the OMS is running.
5. Restart the Agent after running the `add_trust_cert` command.

```
emctl secure add_trust_cert -trust_certs_loc <location of the trusted_certs.txt file>
```

6. Secure the OMS and restart it.

```
emctl secure oms -wallet <location of wallet> -trust_certs_loc <loc of trusted_certs.txt> [any other options]
```

Method 2

1. Create a wallet for each OMS in the grid.
2. Specify the host name of the machine where the OMS is installed or the Load Balancer Name if the OMS is behind the Server Load Balancer for Common Name (CN).
3. Write the certificates of all the Certificate Authorities in the certificate chain (like the Root Certificate Authority, Intermediate Certificate Authority) into a file named `trusted_certs.txt`.
4. Restart the OMS after it has been secured.

```
emctl secure oms -wallet <location of wallet> -trust_certs_loc <loc of trusted_certs.txt> [any other options]
```

5. Either re-secure the Agent by running the `emctl secure agent` command or import the trust points by running the `emctl secure add_trust_cert -trust_certs_loc <location of the trusted_certs.txt file>` command. The `-trust_certs_loc` parameter must contain the path and the filename of the `trusted_certs.txt` file. This file must only contain certificates in base64 format and no special characters or empty lines.

2.4.10.2 Configuring Third Party Certificate for HTTPS WebTier Virtual Host

To configure the third party certificate for HTTPS WebTier Virtual Host:

1. Create a wallet for each OMS in the grid. Specify the host name of the machine where the OMS is installed or the Load Balancer Name if the OMS is behind the Server Load Balancer for Common Name.
2. Run the following command on each OMS:

```
emctl secure console -wallet <location of wallet>
```

Note: Only single-sign-on wallets are supported.

2.5 Accessing Managed Targets

The following topics are discussed in this section:

- Credential Subsystem
- Pluggable Authentication Modules (PAM) Support
- Sudo and Powerbroker Support

2.5.1 Credential Subsystem

Credentials like user names and passwords are typically required to access targets such as databases, application servers, and hosts. Credentials are encrypted and stored in Enterprise Manager. By using appropriate credentials, you can:

- Collect metrics in the background as well as real-time
- Perform jobs like backup, patching, cloning etc.
- Perform real-time target administration like start, stop etc.
- Connect to My Oracle Support

Based on their usage, credentials can be classified into the following categories:

- **Job Credentials:** The job system uses the credential subsystem to retrieve the credentials required to submit a job on the targets. The administrator can define their preferred and default credentials in the preference section of EM console. The user can override the default credentials by specifying different credentials while submitting the job.

Note: If the user chooses to use preferred credentials, these credentials will be used when the user submits the job. If the preferred credentials are not available, the default credentials will be used. If default credentials are not present, the job cannot be submitted.

- **Monitoring Credentials:** These credentials are used by the Management Agent to monitor certain types of targets. For example, most database monitoring involves connecting to the database, which requires a username, password, and optionally, a role. Monitoring credentials, if stored in the repository, can also be potentially used by management applications to connect directly to the target from the OMS.
- **Collection Credentials:** These credentials are associated with user-defined metrics.

To simplify the usage and management of credentials, the following features are available in Enterprise Manager:

- **Preferred Credentials:** Preferred credentials are used to simplify access to managed targets by storing target login credentials in the Management Repository. With preferred credentials set, users can access an Enterprise Manager target that recognizes those credentials without being prompted to log into the target. Preferred credentials are set on a per user basis, thus ensuring the security of the managed enterprise environment.
 - **Default Credentials:** Default credentials can be set for a particular target type and will be available for all the targets of the target type. It will be overridden by target preferred credentials.
 - **Target Credentials:** Target credentials are preferred credentials set for a particular target. They could be used by applications such as the job system, patch, etc. For example, if the user chooses to use preferred credentials while submitting a job, then the preferred credentials set for the target (target credentials) will be used. If the target credentials are not present, the default credentials (for the target type) will be used. If the default credentials are not present, the job will fail. If not specified, by default, preferred credentials refer to preferred target credentials"

For example, to set the host preferred credentials, click Preferences to navigate to the Preferences page. Click the **Preferred Credentials** link in the right panel. In the Preferred Credentials page, click the **Set Credentials** icon for the host. The Host and Cluster Preferred Credentials is displayed.

Figure 2–10 Host and Cluster Preferred Credentials

The screenshot shows the Oracle Enterprise Manager interface for setting preferred credentials. At the top, there's a navigation bar with 'Home', 'Targets', 'Deployments', 'Alerts', 'Compliance', 'Jobs', 'Reports', and 'My Oracle Support'. Below that is the 'Preferences' section, with 'Preferred Credentials' selected. The main heading is 'Host and Cluster Preferred Credentials'. There are 'Revert' and 'Apply' buttons. A tip explains that normal credentials are used by Enterprise Manager functions that need operating system access but do not require administrator privileges. Below this are two sections: 'Default Credentials' and 'Target Credentials'. The 'Default Credentials' section has a table with columns: Normal Username, Normal Password, Run as, Profile, Privileged Username, Privileged Password, Run as, Profile, and Run Privilege. The 'Target Credentials' section has a search box and a table with columns: Name, Scope Username, Normal Password, Run as, Profile, Privileged Username, Privileged Password, Run as, Profile, and Test. The table lists two hosts: 'stasb21.us.oracle.com' and 'dadymb0083.us.oracle.com'.

On this page, you can set both default and explicit preferred credentials for the host and cluster target types. For more details on setting preferred credentials, see the Enterprise Manager Online Help.

2.5.1.1 Managing Credentials Using EMCLI

You can manage passwords using EMCLI verbs. Using EMCLI, you can:

- Change the database user password in both the target database and Enterprise Manager.

```
emcli update_db_password -change_at_target=Yes|No -change_all_reference=Yes|No
```

- Update a password which has already been changed at the host target.

```
emcli update_host_password -change_all_reference=Yes|No
```

- Set preferred credentials for given users.

```
emcli set_credential
  -target_type="ttype"
  [-target_name="tname"]
  -credential_set="cred_set"
  [-user="user"]
  -columns="col1:newval1;col2:newval2;PDP:SUDO/POWERBROKER;RUNAS:oracle;
PROFILE:user1..."
  [-input_file="tag1:file_path1;tag2:file_path2;..."]
  [-oracle_homes="home1;home2"]
  [-monitoring]
```

For detailed descriptions of these verbs, refer to *Enterprise Manager Command Line Interface* guide.

2.5.2 Pluggable Authentication Modules (PAM) Support for Hosts

Pluggable authentication modules, or PAM, is a mechanism to integrate multiple low-level authentication schemes into a high-level application programming interface (API). It allows programs that rely on authentication to be written independently of the underlying authentication scheme. By using PAM, instead of using the local password file to authenticate the user accessing the host, you can take advantage of other authentication mechanisms such as LDAP, RADIUS and Kerberos. If your host authentication is configured over PAM, the Management Agent needs to be configured accordingly to enable PAM Authentication. Refer to note 422073.1 for deployment details.

Note: The local password file (usually `/etc/passwd`) will be checked and used first. This should be synchronized with the LDAP password if it is being used. If this fails, the Management Agent will switch to the external authentication module.

2.5.2.1 Configuring PAM for RHEL4 Users

For users on RHEL4, the PAM file configuration is as follows:

```
##PAM-1.0
auth required pam_ldap.so
account required pam_ldap.so
password required pam_ldap.so
session required pam_ldap.so
```

For more details, see

<https://www.redhat.com/docs/manuals/enterprise/RHEL-4-Manual/ref-guide/s1-pam-format.html>

2.5.2.2 Configuring PAM for AIX Users

For AIX users, use the `edit/etc/pam.conf` file and add the following lines:

emagent auth	required	/usr/lib/security/pam_aix
emagent account	required	/usr/lib/security/pam_aix
emagent password	required	/usr/lib/security/pam_aix
emagent session	required	/usr/lib/security/pam_aix

After editing the file, apply patch **5527130** and run `root.sh`

2.5.3 Sudo and PowerBroker Support

Privilege delegation allows a logged-in user to perform an activity with the privileges of another user. Sudo and PowerBroker are privilege delegation tools that allow a logged-in user to be assigned these privileges. Typically, the privileges that are granted to a specific user are administered centrally. For example, the sudo command can be used to run a script that requires root access:

```
sudo root root.sh
```

In the invocation of sudo in the example above, an administrator can use the sudo command to run a script as root provided he has been granted the appropriate privileges by the system administrator.

Enterprise Manager preferred credentials allow you to use two types of privilege delegation tools: Sudo and PowerBroker. You can use EMCLI or the Manage Privilege Delegation Settings page to set/edit privilege delegation settings for a host. See the *Enterprise Manager Command Line Interface* guide for more information on using the command line.

Sudo: sudo allows a permitted user to execute a command as the super user or another user, as specified in the sudoers file. If the invoking user is root or if the target user is the same as the invoking user, no password is required. Otherwise, sudo requires that users authenticate themselves with a password by default. Once a user has been authenticated, a timestamp is updated and the user may then use sudo without a password for a short period of time (5 minutes unless overridden in sudoers). sudo determines who is an authorized user by consulting the file `/etc/sudoers` file. For more information, see the manual page on sudo (`man sudo`) on Unix. Enterprise Manager authenticates the user using sudo, and executes the script as sudo. For example, if the command to be executed is `foo -arg1 -arg2`, it will be executed as `sudo -S foo -arg1 -arg2`.

PowerBroker: Symark PowerBroker enables UNIX system administrators to specify the circumstances under which other people may run certain programs such as root (or other important accounts). The result is that responsibility for such actions as adding user accounts, fixing line printer queues, and so on, can be safely assigned to the appropriate people, without disclosing the root password. The full power of root is thus protected from potential misuse or abuse—for example, modifying databases or file permissions, erasing disks, or more subtle damage.

Symark PowerBroker can access existing programs as well as its own set of utilities that execute common system administration tasks. Utilities being developed to run on top of Symark PowerBroker can manage passwords, accounts, backups, line printers, file ownership or removal, rebooting, logging people out, killing their programs, deciding who can log in to where from where, and so on. They can also provide TCP/IP, Load Balancer, cron, NIS, NFS, FTP, rlogin, and accounting subsystem management. Users can work from within a restricted shell or editor to access certain programs or files as root.

See your Sudo or PowerBroker documentation for detailed setup and configuration information.

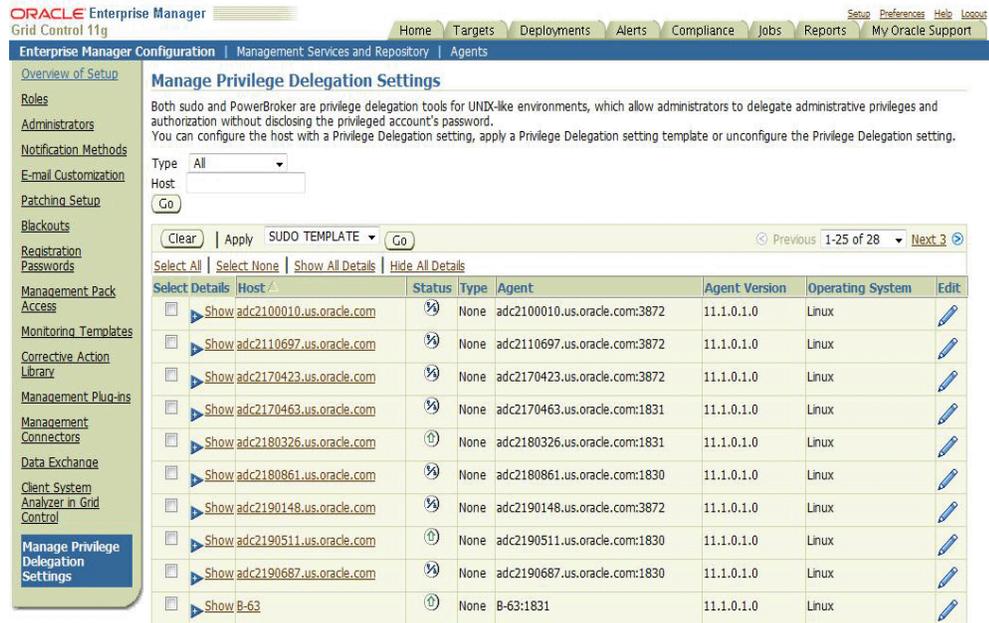
2.5.3.1 Creating a Privilege Delegation Setting

Enterprise Manager allows you to create privilege delegation settings either by creating the setting directly on a host target, or by creating a PDP setting template that you can apply to multiple hosts.

To create a privilege delegation setting directly on a host:

1. Login to Enterprise Manager and navigate to the Setup page. Click **Manage Privilege Delegation Settings** on the left panel. The following screen is displayed:

Figure 2–11 Manage Privilege Delegation Settings



2. For any host target appearing in the table, click **Edit**. Enterprise Manager takes you to the Host Privilege Delegation Setting page.
3. Select a privilege delegation type (Sudo or PowerBroker).
4. Enter the privilege delegation command to be used and, in the case of PowerBroker, the optional Password Prompt.
5. Click **Update** to apply the settings to the host. The following figure shows the Host Privilege Delegation Setting window that you can use to create a PowerBroker setting.

Figure 2–12 Host Privilege Delegation Setting - PowerBroker

The screenshot shows the Oracle Enterprise Manager Grid Control interface. The page title is "Host Privilege Delegation Setting : adc2100010.us.oracle.com". Below the title, there are "Cancel" and "Update" buttons. The main content area has three radio buttons: "None", "Sudo", and "PowerBroker", with "PowerBroker" selected. Under the "Settings" section, there are two input fields: "PowerBroker Password Prompt" with a placeholder "Password prompt" and "PowerBroker command." with a placeholder "For eg. /usr/bin/pbrun -i %RUNAS% %COMMAND%". To the right, under "Parameters", there is a table with the following content:

Name	Description
%COMMAND%	PowerBroker command.
%PROFILE%	Use this profile to run the command.
%RUNAS%	Run the command as this user.
%USERNAME%	Name of the user running the command.

At the bottom of the page, there are "Cancel" and "Update" buttons, a navigation menu with "Home", "Targets", "Deployments", "Alerts", "Compliance", "Jobs", "Reports", "My Oracle Support", "Setup", "Preferences", "Help", and "Logout", and a copyright notice: "Copyright © 1996, 2010, Oracle and/or its affiliates. All rights reserved. Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners. About Oracle Enterprise Manager".

Once you have created a privilege delegation setting, you must apply this setting to selected targets. This setting can be applied to one more hosts or to a composite (Group) target (the group must contain at least one host target). You can apply a Privilege Delegation setting using the Grid Control console by clicking **Setup** on the Enterprise Manager Home page and then choosing **Manage Privilege Delegation Settings** from the left menu panel.

2.6 Cryptographic Support

To protect the integrity of sensitive data in Enterprise Manager, a signing and verification method known as the `emkey` is used. Encryption key is the master key that is used to encrypt/decrypt sensitive data, such as passwords and preferred credentials that are stored in the Repository. The key is originally stored in the repository. It is removed from the repository and copied to the Credential Store during installation of the first OMS. (the `emkey` is secured out-of-the-box). A backup is created in `OMS_ORACLE_HOME/sysman/config/emkey.ora`. It is recommended to create a backup of this file on some other machine. When starting up, OMS reads the `emkey` from the Credential Store and the repository. If the `emkey` is not found or is corrupted, it fails to start. By storing the key separately from the Enterprise Manager schema, we ensure that sensitive data such as Preferred Credentials in the Repository remain inaccessible to the schema owner and other SYSDBA users (Privileged users who can perform maintenance tasks on the database) in the Repository. Moreover, keeping the key from the schema will ensure that sensitive data remain inaccessible while Repository backups are accessed. Further, the schema owner should not have access to the OMS/Repository Oracle homes.

2.6.1 Configuring the emkey

The `emkey` is an encryption key that is used to encrypt and decrypt sensitive data in Enterprise Manager such as host passwords, database passwords and others. By default, the `emkey` is stored in the `ORACLE_HOME/sysman/config/emkey.ora` file. The location of this file can be changed.

WARNING: If the `emkey.ora` file is lost or corrupted, all the encrypted data in the Management Repository becomes unusable. Maintain a backup copy of this file on another system.

During startup, the Oracle Management Service checks the status of the `emkey`. If the `emkey` has been properly configured, it uses it encrypting and decrypting data. If the `emkey` has not been configured properly, the following error message is displayed.

Example 2–12 `emctl start oms` Command

```
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
emctl start oms
Starting HTTP Server ...
Starting Oracle Management Server ...
Checking Oracle Management Server Status ...
Oracle Management Server is not functioning because of the following reason:
The Em Key is not configured properly. Run "emctl status emkey" for more details.
```

2.6.2 `emctl` Commands

The `emctl` commands related to `emkey` are given below:

- `emctl status emkey [-sysman_pwd <pwd>]`
- `emctl config emkey -copy_to_credstore [-sysman_pwd <pwd>]`
- `emctl config emkey -copy_to_repos [-sysman_pwd <pwd>]`
- `emctl config emkey -remove_from_repos [-sysman_pwd <pwd>]`
- `emctl config emkey -copy_to_file_from_credstore -admin_host <host> -admin_port <port> -admin_user <username> [-admin_pwd <pwd>] [-repos_pwd <pwd>] -emkey_file <emkey file>`
- `emctl config emkey -copy_to_file_from_repos (-repos_host <host> -repos_port <port> -repos_sid <sid> | -repos_conn_desc <conn desc>) -repos_user <username> [-repos_pwd <pwd>] [-admin_pwd <pwd>] -emkey_file <emkey file>`
- `emctl config emkey -copy_to_credstore_from_file -admin_host <host> -admin_port <port> -admin_user <username> [-admin_pwd <pwd>] [-repos_pwd <pwd>] -emkey_file <emkey file>`
- `emctl config emkey -copy_to_repos_from_file (-repos_host <host> -repos_port <port> -repos_sid <sid> | -repos_conn_desc <conn desc>) -repos_user <username> [-repos_pwd <pwd>] [-admin_pwd <pwd>] -emkey_file <emkey file>`

2.6.2.1 `emctl status emkey`

This command shows the health or status of the `emkey`. Depending on the status of the `emkey`, the following messages are displayed:

- When the `emkey` has been correctly configured in the Credential Store, the following message is displayed.

Example 2–13 `emctl status emkey` - Example 1

```
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
The EmKey is configured properly, but is not secure. Secure the EmKey by running
```

```
"emctl config emkey -remove_from_repos"
```

- When the emkey has been correctly configured in the Credential Store and has been removed from the Management Repository, the following message is displayed.

Example 2–14 emctl status emkey - Example 2

```
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
The EMKey exists in the Management Repository, but is not configured properly or
is corrupted in the credential store.
Configure the EMKey by running "emctl config emkey -copy_to_credstore".
```

- When the emkey is corrupted in the Credential Store and removed from the Management Repository, the following message is displayed.

Example 2–15 emctl status emkey - Example 3

```
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
The EMKey is not configured properly or is corrupted in the credential store and
does not exist in the Management Repository. To correct the problem:
1) Get the backed up emkey.ora file.
2) Configure the emkey by running "emctl config emkey -copy_to_credstore_from_
file"
```

2.6.2.2 emctl config emkey -copy_to_credstore

This command copies the emkey from the Management Repository to the Credential Store.

Example 2–16 Sample Output of the emctl config emkey -copy_to_credstore Command

```
emctl config emkey -copy_to_credstore [-sysman_pwd <pwd>]
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
The EMKey has been copied to the Credential Store.
```

2.6.2.3 emctl config emkey -copy_to_repos

This command copies the emkey from the Credential Store to Management Repository.

Example 2–17 Sample Output of the emctl config emkey -copy_to_repos Command

```
emctl config emkey -copy_to_repos [-sysman_pwd <pwd>]
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
The EMKey has been copied to the Management Repository. This operation will cause
the EMKey to become unsecure.
After the required operation has been completed, secure the EMKey by running
"emctl config emkey -remove_from_repos".
```

2.6.2.4 emctl config emkey -copy_to_file_from_credstore

This command copies the emkey from the Credential Store to a specified file.

Example 2–18 Sample Output of the emctl config emkey -copy_to_file_from_credstore Command

```
emctl config emkey -copy_to_file_from_credstore -admin_host <host> -admin_port
<port> -admin_user <username> [-admin_pwd <pwd>] [-repos_pwd <pwd>] -emkey_file
<emkey file>
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
The EMKey has been copied to file.
```

2.6.2.5 emctl config emkey -copy_to_file_from_repos

This command copies the emkey from the Management Repository to a specified file.

Example 2–19 Sample Output of the emctl config emkey -copy_to_file_from_repos Command

```
emctl config emkey -copy_to_file_from_repos (-repos_host <host> -repos_port <port>
-repos_sid <sid> | -repos_conndesc <conn desc>) -repos_user <username> [-repos_pwd
<pwd>] [-admin_pwd <pwd>] -emkey_file <emkey file>
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
The EMKey has been copied to file.
```

2.6.2.6 emctl config emkey -copy_to_credstore_from_file

This command copies the emkey from a specified file to the Credential Store.

Example 2–20 Sample Output of the emctl config emkey -copy_to_credstore_from_file Command

```
emctl config emkey -copy_to_credstore_from_file -admin_host <host> -admin_port
<port> -admin_user <username> [-admin_pwd <pwd>] [-repos_pwd <pwd>] -emkey_file
<emkey file>
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
The EMKey has been copied to the Credential Store.
```

2.6.2.7 emctl config emkey -copy_to_repos_from_file

This command copies the emkey from a specified file to the repository.

Example 2–21 Sample Output of the emctl config emkey -copy_to_repos_from_file Command

```
emctl config emkey -copy_to_repos_from_file (-repos_host <host> -repos_port <port>
-repos_sid <sid> | -repos_conndesc <conn desc>) -repos_user <username> [-repos_pwd
<pwd>] [-admin_pwd <pwd>] -emkey_file <emkey file>
Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
The EMKey has been copied to the Management Repository. This operation will cause
the EMKey to become unsecure.
After the required operation has been completed, secure the EMKey by running
"emctl config emkey -remove_from_repos".
```

2.6.2.8 emctl config emkey -remove_from_repos

This command removes the emkey from the repository.

Example 2–22 Sample Output of emctl config emkey -remove_from_repos Command

```
emctl config emkey -remove_from_repos [-sysman_pwd <pwd>]
```

Oracle Enterprise Manager 11g Release 1 Grid Control
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
The EMKey has been removed from the Management Repository.

Note: If the emkey is corrupted, you cannot remove it from the Management Repository.

2.6.3 Install and Upgrade Scenarios

This section explains the install and upgrade scenarios for emkey.

2.6.3.1 Installing the Management Repository

A new emkey is generated as a strong random number when the Management Repository is installed.

2.6.3.2 Installing the First Oracle Management Service

When the Oracle Management Service is installed, the Installer copies the emkey to Credential Store and removes it from repository (emkey is secured out-of-box).

2.6.3.3 Upgrading from 10.2 to 11.1

The Management Repository is upgraded as usual. When upgrading the OMS, the omsca (OMS Configuration Assistant) copies the emkey to Credential Store and removes from repository. If the emkey is already secured before upgrade or has been removed from repository, then omsca reads the emkey from emkey.ora file present in old OMS Oracle Home and copies it to Credential Store.

Note: After all the Oracle Management Service have been upgraded, you can secure the emkey, that is, remove it from the Management Repository by running the following command:

```
emctl config emkey -remove_from_repos
```

2.6.3.4 Recreating the Management Repository

When the Management Repository is recreated, a new emkey is generated. This new key will not be in synchronization with the existing emkey Credential Store.

1. Copy the new emkey to Credential Store by using the `emctl config emkey -copy_to_credstore` command.
2. Take a backup by entering the `emctl config emkey -copy_to_file_from_repos` command or the `emctl config emkey -copy_to_file_from_credstore` command.
3. Secure the emkey by using the `emctl config emkey -remove_from_repos` command.

2.7 Setting Up the Auditing System for Enterprise Manager

All operations performed by Enterprise Manager users such as creating users, granting privileges, starting a remote job like patching or cloning, need to be audited to ensure compliance with the Sarbanes-Oxley Act of 2002 (SAS 70). This act defines standards an auditor must use to assess the contracted internal controls of a service organization.

Auditing an operation enables an administrator to monitor, detect, and investigate problems and enforce enterprise wide security policies.

Irrespective of how the user has logged into Enterprise Manager, if auditing is enabled, each user action is audited and the audit details are stored in a record.

2.7.1 Configuring the Enterprise Manager Audit System

You can configure the Enterprise Manager Audit System by using the following emcli commands:

- `enable_audit`: Enables auditing for all user operations.
- `disable_audit`: Disables auditing for all user operations.
- `show_operations_list`: Shows a list of the user operations being audited.
- `show_audit_settings`: Shows the audit status, operation list, externalization service details, and purge period details.

2.7.2 Configuring the Audit Data Export Service

Audit data needs to be protected and maintained for several years. The volume of audit data may become very large and impact the performance of the system. To limit the amount of data stored in the repository, the audit data must be externalized or archived at regular intervals. The archived audit data is stored in an XML file complying with the ODL format. To externalize the audit data, the `EM_AUDIT_EXTERNALIZATION` API is used. Records of the format `<file-prefix>.NNNNN.xml`, where `NNNN` is a number are generated. The numbers start with `00001` and continue to `99999`.

You can set up the audit externalization service for exporting audit data into the file system by using the `update_audit_setting -externalization_switch` command.

2.7.3 Updating the Audit Settings

The `update_audit_settings` command updates the current audit settings in the repository and restarts the Management Service.

Example 2-23 Usage of the update_audit_setting command

```
emcli update_audit_settings
  -audit_switch="ENABLE/DISABLE"
  -operations_to_enable="name of the operations to enable, for all oprtations
  use ALL"
  -operations_to_disable="name of the operations to disable, for all
  oprtations use ALL"
  -externalization_switch="ENABLE/DISABLE"
  -directory_name="directory_name (DB Directory)"
  -file_prefix="file_prefix"
  -file_size="file_size (Bytes)"
  -data_retention_period="data_retention_period (Days)"
```

- `-audit_switch`: Enables auditing across Enterprise Manager. The possible values are `ENABLE/DISABLE`. Default value is `DISABLE`.
- `-operations_to_disable`: Enables auditing for specified operations. Enter **All** to enable all operations.

- `-operations_to_disable`: Disables auditing for specified operations. Enter **All** to disable all operations.
- `-externalization_switch`: Enables the audit data export service. The possible values are `ENABLE/DISABLE`. Default value is `DISABLE`.
- `-directory`: The database directory that is mapped to the OS directory where the export service archives the audit data files.
- `-file_prefix`: The file prefix to be used by the export service to create the file in which audit data is to be stored.
- `-file_size`: The size of the file on which the audit data is to be stored. The default value is 5000000 bytes.
- `data_retention_period`: The period for which the audit data is to be retained inside the repository. The default value is 365 days.

2.7.4 Searching the Audit Data

You can search for audit data that has been generated over a specified period. You can also search for the following:

- Audit details of a specific user operation or all user operations.
- Audit details of operations with a Success or Failure status or All operations.

To view the audit data, click the **Setup** option. On the Setup page, click the **Management Services and Repository** tab. The Overview page is displayed. Click the **Audit Data** link under the Audit section. The Audit Data page is displayed. Specify the search criteria in the fields and click **Go**. The results are displayed in the Summary table.

Figure 2–13 Audit Data Search Page

The screenshot displays the Oracle Enterprise Manager interface for searching audit data. The search criteria are as follows:

- Start Date:** Mar 9, 2010, Hour 00, Min 00
- End Date:** Mar 9, 2010, Hour 23, Min 59
- Operation:** All
- Status:** All
- Administrator:** (empty field)
- Rows Displayed:** 200

The search results are shown in a table with the following columns: Timestamp, Administrator, Operation, Status, and Message.

Timestamp	Administrator	Operation	Status	Message
Mar 9, 2010 12:05:49 AM	OIDS01	STOP JOB	Success	Successful Stop Job on Job DB TEST operation by OIDS01
Mar 9, 2010 12:05:49 AM	OIDS01	Delete Job	Success	Successful DELETE_JOB on Job DB TEST operation by OIDS01
Mar 9, 2010 12:05:35 AM	OIDS01	Delete Job	Success	Successful DELETE_JOB on Job TES 2 operation by OIDS01
Mar 9, 2010 12:05:34 AM	OIDS01	STOP JOB	Success	Successful Stop Job on Job TES 2 operation by OIDS01
Mar 9, 2010 12:04:03 AM	SYSMAN	Job Execution	Success	Job Output. Obtained successfully
Mar 9, 2010 12:03:48 AM	OIDS01	EM Logout	Success	OIDS01 Logged out successfully
Mar 9, 2010 12:02:52 AM	OIDS01	Job Execution	Success	Job Output. Obtained successfully
Mar 9, 2010 12:01:11 AM	OIDS01	Remote Operation Job	Success	Job TES 2 successfully dispatched
Mar 9, 2010 12:01:10 AM	OIDS01	Submit Job	Success	Job TES 2 submitted successfully by OIDS01
Mar 9, 2010 12:00:01 AM	SYSMAN	Job Execution	Success	Job Output. Obtained successfully
Mar 9, 2010 12:00:01 AM	SYSMAN	Job Execution	Success	Job Output. Obtained successfully
Mar 9, 2010 12:00:01 AM	SYSMAN	Job Execution	Success	Job Output. Obtained successfully

To view the details of each record that meets the search criteria, select **Detailed** in the View drop-down list. To drill down to the full record details, click on the **Timestamp**. The Audit Record page is displayed.

Figure 2–14 Audit Record Details Page

General

Operation Timestamp: Mar 9, 2010 12:09:03 AM (Timezone +00:00)
 Administrator: SYSMAN
 Operation: Job Execution
 Status: Success
 Message: Job Output Obtained successfully
 Normalized Timestamp: Mar 9, 2010 12:09:03 AM (Timezone +00:00)

Client Information

Session: 9
 IP Address: Not Applicable
 Hostname: stjajs01
 Upstream Component Type: Others
 Authentication Type: Not Applicable
 Upstream Component Name: Others

OMS Information

Hostname: stjajs01
 IP Address: 140.87.25.101
 Instance ID: 1

Operation Specific Information

Object Type: EventCollection
 Object Name: EVENTJOB_3BEC3979E25C4DF98529D4E313DFAD5B
 Object Owner: SYSMAN
 Job Name: EVENTJOB_3BEC3979E25C4DF98529D4E313DFAD5B
 Job Type: EventCollection
 Step_Status: COMPLETED

Field Name	Description
General	
Operation Timestamp	The date and time on which the operation took place.
Administrator	The id of the administrator who has logged into Enterprise Manager.
Operation	The type of operation being audited.
Status	The status of the operation which can be success or failure.
Message	A descriptive message indicating the status of the operation.
Normalized Timestamp	This is the UTC timestamp.
Client Information	
Session	This can either be the HTTP Session ID or the DBMS Session ID.
IP Address	The IP address of the client’s host machine.
Hostname	The name of the client’s host machine.
Upstream Component Type	The type of client, Console, Web Service, EMCLI, being used.
Authentication Type	The nature of the session (HTTP Session, DB Session).
Upstream Component Name	The name of the client being used.
OMS Information	
Hostname	The host name of the Oracle Management Service.
IP Address	The IP address of the Oracle Management Service.

Field Name	Description
Instance ID	The Instance ID of the Oracle Management Service.
Operation Specific Information	
Object Name	The operation being performed on an object

2.7.5 List of Operations Audited

The following table lists the names of operation and their description.

Table 2-5 List of Operations Audited

Operation Name	Description
change_password	Change Password
create_user	Create User
delete_user	Delete User
logon	Login
logoff	Logout
grant_role	Grant Role
grant_target_priv	Grant Target Privilege
revoke_role	Revoke Role
revoke_target_priv	Revoke Target Privilege
submit_job	Submit Job
edit_job	Edit Job
delete_job	Delete Job
change_pref_cred	Change Preferred Credential
modify_user	Modify User
grant_system_priv	Grant System Privilege
grant_job_priv	Grant Job Privilege
revoke_system_priv	Revoke System Privilege
revoke_job_priv	Revoke Job Privilege
remote_op	Remote Operation Job
get_file	Get File
put_file	Put File
file_transfer	File Transfer
create_role	Create Role
delete_role	Delete Role
modify_role	Modify Role
job_output	Job Output
suspend_job	Suspend Job
agent_resync	Agent Re synchronization Operation
repository_resync	Repository Re synchronization Operation

Table 2–5 (Cont.) List of Operations Audited

Operation Name	Description
remove_privilege_delegation_setting	Remove Privilege Delegation Setting
set_privilege_delegation_setting	Set Privilege Delegation Setting
add_agent_registration_password	Add Registration Password
edit_agent_registration_password	Edit Registration Password
delete_agent_registration_password	Delete Registration Password
agent_registration_password_usage	Registration Password Usage
audit_settings	Enable or Disable Auditing
audit_export_settings	Externalize Audit Data Settings
create_template	Create Template
edit_template	Edit Template
delete_template	Delete Template
apply_template	Apply Template
save_monitoring_settings	Save Monitoring Settings
modify_metric_settings	Modify Metric Settings
modify_policy_settings	Modify Policy Settings
create_udp	Create User Defined Policy
edit_udp	Edit User Defined Policy
delete_udp	Delete User Defined Policy
evaluate_udp	Evaluate User Defined Policy
import_udp	Import User Defined Policy
create_udpg	Create User Defined Policy Group
edit_udpg	Edit User Defined Policy Group
delete_udpg	Delete User Defined Policy Group
delete_pg_eval	Delete Policy Group Evaluation Results
create_pg_sched	Create Policy Group Schedule
edit_pg_sched	Edit Policy Group Schedule
delete_pg_sched	Delete Policy Group Schedule
db_login	Audit Database User Login
db_logout	Audit Database User Logout

Table 2–5 (Cont.) List of Operations Audited

Operation Name	Description
db_start	Audit Database Startup
db_shutdown	Audit Database Shutdown
db_restart	Audit Database Restart

2.8 Additional Security Considerations

After you enable security for the Enterprise Manager components and framework, there are additional security considerations. This section provides the following topics:

- [Changing the SYSMAN and MGMT_VIEW Passwords](#)
- [Responding to Browser-Specific Security Certificate Alerts](#)
- [Configuring Beacons to Monitor Web Applications Over HTTPS](#)
- [Using ORACLE_HOME Credentials](#)
- [Patching Oracle Homes When the User is Locked](#)
- [Cloning Oracle Homes](#)

2.8.1 Changing the SYSMAN and MGMT_VIEW Passwords

This section describes the commands used to change the SYSMAN and MGMT_VIEW passwords.

2.8.1.1 Changing the SYSMAN User Password

To change the password of the SYSMAN user, enter the following command:

```
emctl config oms -change_repos_pwd [-change_in_db] [-old_pwd <old_pwd>] [-new_pwd <new_pwd>] [-use_sys_pwd [-sys_pwd <sys_pwd>]]
```

You must run this command on each Management Service in your environment.

Parameter	Description
-change_in_db	This parameter is optional and is used to change the SYSMAN password in the repository. If there are multiple Management Services running, this parameter must be set to true for at least one Management service. If this parameter is not specified, the emoms.properties file will be updated with the new SYSMAN password.
-old_pwd	This is the current SYSMAN password.
-new_pwd	This is the new password.
-use_sys_pwd	This parameter is optional and is used to connect to the database as a SYS user.
-sys_pwd	This is the password for the SYS user.

2.8.1.2 Changing the MGMT_VIEW User Password

To change the password of the MGMT_VIEW user, enter the following command:

```
emctl config oms -change_view_user_pwd [-sysman_pwd <sysman_pwd>] [-user_pwd
```

```
<user_pwd>] [-auto_generate]
```

Parameter	Description
-sysman_pwd	The password for the SYSMAN user.
-user_pwd	The new password for the MGMT_VIEW user. This is an optional parameter and if it is not specified, the password is auto generated.
-auto_generate	If this option is specified, the password is auto generated.

2.8.2 Responding to Browser-Specific Security Certificate Alerts

This section describes how to respond to browser-specific security alert dialog boxes when you are using Enterprise Manager in a secure environment.

The security alert dialog boxes described in this section should appear only if you have enabled Enterprise Manager Framework Security, but you have not completed the more extensive procedures to secure your WebTier properly.

This section contains the following topics:

- [Responding to the Internet Explorer Security Alert Dialog Box](#)
- [Responding to the Netscape Navigator New Site Certificate Dialog Box](#)
- [Preventing the Display of the Internet Explorer Security Information Dialog Box](#)

2.8.2.1 Responding to the Internet Explorer Security Alert Dialog Box

If you enable security for the Management Service, but do not enable the more extensive security features of your WebTier, you will likely receive a Security Alert dialog box similar to the one shown in [Figure 2-15](#) when you first attempt to display the Grid Control Console using the HTTPS URL in Internet Explorer.

Note: The instructions in this section apply to Internet Explorer 5.5. The instructions may vary for other supported browsers.

Figure 2-15 Internet Explorer Security Alert Dialog Box



When Internet Explorer displays the Security Alert dialog box, use the following instructions to install the certificate and avoid viewing this dialog box again in future Enterprise Manager sessions:

1. In the Security Alert dialog box, click **View Certificate**.

Internet Explorer displays the Certificate dialog box.

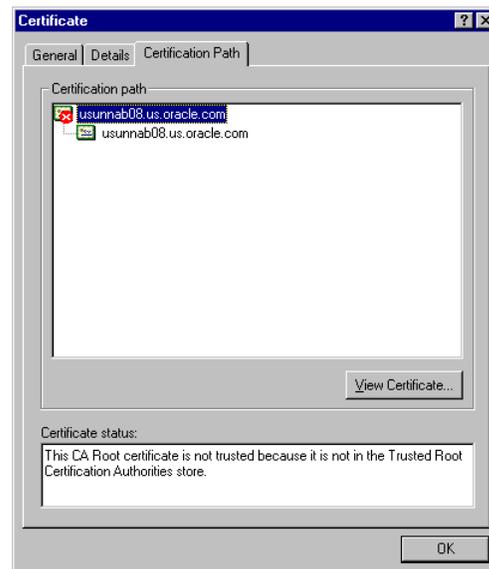
2. Click the **Certificate Path** tab and select the first entry in the list of certificates as shown in [Figure 2–16](#).
3. Click **View Certificate** to display a second Certificate dialog box.
4. Click **Install Certificate** to display the Certificate Import wizard.
5. Accept the default settings in the wizard, click **Finish** when you are done, and then click **Yes** in the Root Certificate Store dialog box.

Internet Explorer displays a message box indicating that the Certificate was imported successfully.

6. Click **OK** to close each of the security dialog boxes and click **Yes** on the Security Alert dialog box to continue with your browser session.

You should no longer receive the Security Alert dialog box in any future connections to Enterprise Manager when you use this browser.

Figure 2–16 Certificate Path Tab on the Internet Explorer Certificate Dialog Box



2.8.2.2 Responding to the Netscape Navigator New Site Certificate Dialog Box

If you enable security for the Management Service, but you do not enable the more extensive security features of your WebTier, you will likely receive a New Site Certificate dialog box similar to the one shown in [Figure 2–17](#) when you first attempt to display the Grid Control Console using the HTTPS URL in Netscape Navigator.

Note: The instructions in this section apply to Netscape Navigator 4.79. The instructions may vary for other supported browsers.

When Netscape Navigator displays the New Site Certificate dialog box, use the following instructions to install the certificate and avoid viewing this dialog box again in future Enterprise Manager sessions:

1. Review the instructions and information on each wizard page; click **Next** until you are prompted to accept the certificate.

2. Select **Accept this certificate forever (until it expires)** from the list of options.
3. On the last screen of the wizard, click **Finish** to close the wizard and continue with your browser session.

You should no longer receive the New Site Certificate dialog box when using the current browser.

Figure 2–17 Netscape Navigator New Site Certificate Dialog Box

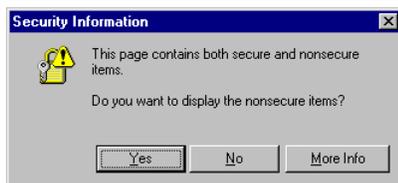


2.8.2.3 Preventing the Display of the Internet Explorer Security Information Dialog Box

After you enable Security for the Management Service, you may receive a dialog box similar to the one shown in [Figure 2–18](#) whenever you access certain Enterprise Manager pages.

Note: The instructions in this section apply to Internet Explorer 6.0. The instructions may vary for other supported browsers.

Figure 2–18 Internet Explorer Security Information Dialog Box



To stop this dialog box from displaying:

1. Select **Internet Options** from the Internet Explorer **Tools** menu.
2. Click the **Security** tab.
3. Select **Internet** and then click **Custom Level**.

Internet Explorer displays the Security Settings dialog box.

4. Scroll down to **Miscellaneous** settings and enable the **Display Mixed Content** option.

2.8.3 Configuring Beacons to Monitor Web Applications Over HTTPS

Oracle Beacons provide application performance availability and performance monitoring. They are part of the Application Service Level Management features of Enterprise Manager.

See Also: "About Application Service Level Management" in the Enterprise Manager Online Help

When a Beacon is used to monitor a URL over Secure Sockets Layer (SSL) using an HTTPS URL, the Beacon must be configured to recognize the Certificate Authority that has been used by the Web site where that URL resides.

See Also: "The Public Key Infrastructure Approach to Security" in the *Oracle Security Overview* for an overview of Public Key Infrastructure features, such as Certificate Authorities

The Beacon software is preconfigured to recognize most commercial Certificate Authorities that are likely to be used by a secure Internet Web Site. However, you may encounter Web Sites that, although available over HTTPS, do not have a Certificate that has been signed by a commercial Certificate Authority recognized by the Beacon. The following are out-of-box certificates recognized by Beacons:

- Class 1 Public Primary Certification Authority by VeriSign, Inc.
- Class 2 Public Primary Certification Authority by VeriSign, Inc.
- Class 3 Public Primary Certification Authority by VeriSign, Inc.
- Secure Server Certification Authority by RSA Data Security, Inc.
- GTE CyberTrust Root by GTE Corporation
- GTE CyberTrust Global Root by GTE CyberTrust Solutions, Inc.
- Entrust.net Secure Server Certification Authority by Entrust.net ((c) 1999
- Entrust.net Limited, www.entrust.net/CPS incorp. by ref. (limits liab.))
- Entrust.net Certification Authority (2048) by Entrust.net ((c) 1999
- Entrust.net Limited, www.entrust.net/CPS_2048 incorp. by ref. (limits liab.))
- Entrust.net Secure Server Certification Authority by Entrust.net ((c) 2000
- Entrust.net Limited, www.entrust.net/SSL_CPS incorp. by ref. (limits liab.))

In those cases, for example, if you attempt to use the Test section of the Beacon Performance page to test the HTTP Response of the secure URL, the following error appears in the **Status Description** column of the Response Metrics table on the URL Test Page:

```
javax.net.ssl.SSLException: SSL handshake failed:
X509CertChainIncompleteErr--https://mgmtsys.acme.com/OracleMyPage.Home
```

See Also: "Using Beacons to Monitor Remote URL Availability" in the Enterprise Manager online help

To correct this problem, you must allow the Beacon to recognize the Certificate Authority that was used by the Web Site to support HTTPS. You must add the Certificate of that Certificate Authority to the list of Certificate Authorities recognized by Beacon.

To configure the Beacon to recognize the Certificate Authority:

1. Obtain the Certificate of the Web Site's Certificate Authority, as follows:
 - a. In Microsoft Internet Explorer, connect to the HTTPS URL of the Web Site you are attempting to monitor.
 - b. Double-click the lock icon at the bottom of the browser screen, which indicates that you have connected to a secure Web site.

The browser displays the Certificate dialog box, which describes the Certificate used for this Web site. Other browsers offer a similar mechanism to view the Certificate detail of a Web Site.
 - c. Click the **Certificate Path** tab and select the first entry in the list of certificates as shown in [Figure 2-16](#).
 - d. Click **View Certificate** to display a second Certificate dialog box.
 - e. Click the **Details** tab on the Certificate window.
 - f. Click **Copy to File** to display the Certificate Manager Export wizard.
 - g. In the Certificate Manager Export wizard, select **Base64 encoded X.509 (.CER)** as the format you want to export and save the certificate to a text file with an easily-identifiable name, such as `beacon_certificate.cer`.
 - h. Open the certificate file using a text editor.

The content of the certificate file will look similar to the content shown in [Example 2-24](#).

2. Update the list of Beacon Certificate Authorities as follows:
 - a. Locate the `b64InternetCertificate.txt` file in the following directory of Agent Home of the Beacon host:

`agent_home/sysman/config/`

This file contains a list of Base64 Certificates.
 - b. Edit the `b64InternetCertificate.txt` file and add the contents of the Certificate file you just exported to the end of the file, taking care to include all the Base64 text of the Certificate including the BEGIN and END lines.
3. Restart the Management Agent.

After you restart the Management Agent, the Beacon detects your addition to the list of Certificate Authorities recognized by Beacon and you can successfully monitor the availability and performance of the secure Web site URL.

Example 2-24 Sample Content of an Exported Certificate

```
-----BEGIN CERTIFICATE-----  
MIIDBzCCAnCgAwIBAgIQTs4NcImNY3JAs5edi/5RkTANBgkqhkiG9w0BAQQFADCB  
... base64 certificate content...  
-----END CERTIFICATE-----
```

2.8.4 Using ORACLE_HOME Credentials

Oracle Enterprise Manager 11g Release 1 introduces the concept of `ORACLE_HOME` credentials to designate the owner of the `ORACLE_HOME` with special credentials for the `ORACLE_HOME`. The operating system user who installs the software will also need to perform the patching. In Oracle Enterprise Manager 11g Release 1, one can explicitly set the `ORACLE_HOME` credentials and store it in the Management

Repository. While patching, the user can use existing operating system credentials or override it under special circumstances. The user can specify `ORACLE_HOME` credentials and in the same interface choose to store it in the Management Repository for future use.

The Enterprise Manager Command line interface (EMCLI) also provides a facility to set `ORACLE_HOME` credentials. This is useful in cases where the Super Administrator sets the credentials and the user who initiates the patching job is unaware of the actual credentials. For auditing in security-hardened data centers, the owner of the software is usually different from the user who initiates the patching job. The patching application internally switches the user context to the owner of the software and patches the software. To emulate such a case, the patch administrator will set the `ORACLE_HOME` credentials to the owner of the `ORACLE_HOME`. The Grid Control user who executes the patching job will be unaware of the credentials. The patching job will internally execute as the owner of the `ORACLE_HOME`. Grid Control will audit the patching job and capture the name of the Grid Control user who initiated the job. For example, if the owner of the `ORACLE_HOME` is "X", the patch super administrator in Grid Control is "Y" and the target administrator in Grid Control is "Z". "Y" will set the `ORACLE_HOME` credential to "X" with the password, using EMCLI. "Z" will submit the patching job using the already stored preferred credentials. Grid Control will audit the job as submitted by "Z".

The following is an example for setting the Oracle Home credentials using command line:

```
emcli set_credential -target_type=host -target_name=val1 -credential_set=OHCreds
-column="OHUsername:val2;OHPassword:val3"
-oracle_homes="val4"
```

where:

val1 = Hostname

val2 = Oracle Home user name

val3 = Oracle Home password

val4 = Oracle Home location

You can also set credentials for multiple Oracle Homes on the same host using the following command:

```
emcli set_credential -target_type=host -target_name=val1 -credential_set=OHCreds
-column="OHUsername:val2;OHPassword:val3"
-oracle_homes="val4;val5"
```

where

val1 = Hostname

val2 = Oracle Home user name

val3 = Oracle Home password

val4 = Oracle Home location 1

val5 = Oracle Home location 2

Note: Only one host can be passed to the verb.* If one wants multiple Oracle Home credentials on multiple hosts, then you will need Shell or Perl script to read lines, one at a time, from a file containing the host, credential values, and home location, and call the `emcli set_credential verb` for each row in the file.

The `emcli set_credential` command sets preferred credentials for given users. The following table describes the input values to the `emcli set_credential` command.

Table 2-6 *emcli set_credential Parameters*

Parameter	Input Value	Description
-target_type	-target_type="ttype"	Type of target. Must be "host" in case the "-oracle_homes" parameter is specified.
-target_name	[-target_name="tname"]	Name of target. Omit this argument to set enterprise preferred credentials. Must be hostname in case "-oracle_homes" parameter is specified
-credential_set	-credential_set="cred_set"	Credential set affected.
-user	[-user="user"]	Enterprise Manager user whose credentials are affected. If omitted, the current user's credentials are affected.
-columns	-columns="col1:newval1;col2:newval2;..."	The name and new value of the column(s) to set. Every column of the credential set must be specified. Alternatively, a tag from the -input_file argument may be used so that the credential values are not seen on the command line. This argument may be specified more than once.
-input_file	[-input_file="tag1:file_path1;tag2:file_path2;..."]	Path of file that has -columns argument(s). This option is used to hide passwords. Each path must be accompanied by a tag which is referenced in the -columns argument. This argument may be specified more than once.
-oracle_homes	[-oracle_homes="home1;home2"]	Name of Oracle Homes on the target host. Credentials will be added/updated for all specified home

2.8.5 Patching Oracle Homes When the User is Locked

To patch an Oracle Home used by a user "Oracle" and the user is locked:

1. Edit the default patching script and prepend `sudo` or `sudo -u` or `pbrun -u` to the default patching step. You need to set a policy (by editing the `sudoers` file) to allow the user submitting the job (who must be a valid operating system user) to be able to run `sudo` or `pbrun` without being prompted for password.

Note: You cannot patch Oracle Homes without targets. This must be done by using the Patching wizard.

2.8.6 Cloning Oracle Homes

The cloning application is wizard-driven. The source of the Oracle Home being cloned may be either an installed Oracle Home or a Software Library. Following are the steps in the cloning process:

1. If the source is an installed Oracle Home, then, after selecting the Oracle Home, a user will need to specify the Oracle Home credentials. These credentials once specified for an Oracle Home are stored in the repository. The next time a user clones the same Oracle Home, these credentials are automatically populated. Other parameters queried from the user at this point is a temporary location (on the source computer) and the list of files to be excluded from the Oracle Home. If the cloning source is a Software Library, the source Oracle Home credentials will not be queried for.
2. The user needs to specify the target location and provide the required credentials for each target location. These credentials will be the Oracle Home credentials for each of these target locations. Subsequently, if a user selects any of these cloned Oracle Homes as a source, the Oracle Home credentials are automatically populated.
3. Depending on the product being cloned, the user can view the Enterprise Manager page where query parameters required for the particular product being cloned are displayed.
4. The user can, then, view the execution of user-supplied pre-cloning and post-cloning scripts and the root.sh script. The root.sh script will always be run with sudo privileges, but the user has the option to decide if the pre-cloning and post-cloning scripts run with sudo privileges.
5. Finally, the user can schedule the cloning job at a convenient time.

For more information about cloning, refer to the Enterprise Manager Online Help.

Notifications

The notification system allows you to notify Enterprise Manager administrators of alerts, policy violations, and the status changes of job executions. In addition to notifying administrators, the notification system can perform actions such as executing operating system commands (including scripts) and PL/SQL procedures when an alert is triggered. This capability allows you to implement automatically specific IT practices under particular alert conditions. For example, if an alert is generated when monitoring the operational (up/down) status of a database, 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. Some administrators may want to send third-party applications a notification when a certain metric has exceeded a threshold.

This chapter covers the following:

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

3.1 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.

3.1.1 Setting Up a Mail Server for Notifications

Before Enterprise Manager can send e-mail notifications, you must first specify the Outgoing Mail (SMTP) servers to be used by the notification system. Once set, you can

then define e-mail notifications for yourself or, if you have Super Administrator privileges, other Enterprise Manager administrators.

You specify the Outgoing Mail (SMTP) server on the Notification Methods page (Figure 3-1). Display the Notification Methods page by clicking **Setup** on any page in the Grid Control console and clicking **Notification Methods** in the vertical navigation bar.

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 3-1 shows sample notification method entries.

Example 3-1 Mail Server Settings

- **Outgoing Mail (SMTP) Server** - smtp01.mycorp.com:587, smtp02.mycorp.com
- **User Name** - myadmin
- **Password** - *****
- **Confirm Password** - *****
- **Identify Sender As** - Enterprise Manager
- **Sender's E-mail Address** - mgmt_rep@mycorp.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 3–1 Defining a Mail Server

Oracle Enterprise Manager (SYSMAN) - Notification Methods - Mozilla Firefox

Oracle Enterprise Manager (SYSMAN...)

ORACLE Enterprise Manager
Grid Control 11g

Home Targets Deployments Alerts Compliance Jobs Reports My Oracle Support

Enterprise Manager Configuration | Management Services and Repository | Agents

Overview of Setup
Roles
Administrators
Notification Methods
Email Customization
Patching Setup
Blackouts
Registration Passwords
Management Pack Access
Monitoring Templates
Corrective Action Library
Management Plug-ins
Management Connectors
Data Exchange
Client System Analyzer in Grid Control
Manage Privilege Delegation Settings

Notification Methods

Notification Methods allow you to globally define different mechanisms for sending notifications. These include e-mail, SNMP traps and running custom scripts. Once defined, Notification Methods are used by Notification Rules to send notifications to administrators for alerts, policy violations or job status changes. Each administrator has Notification Rules defined as a preference.

Mail Server

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

Revert Apply
Test Mail Servers

Outgoing Mail (SMTP) Server

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

Add OS Command Go

Name	Type	Support Repeat Notifications
No notification methods found.		

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

Repeat Notifications

Repeat notifications allow you to be notified repeatedly about the same metric or availability alert. Once enabled, you will still need to choose the Repeat Notifications option in each Notification 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

Home | Targets | Deployments | Alerts | Compliance | Jobs | Reports | My Oracle Support | Setup | Preferences | Help | Logout

Copyright © 1996, 2010, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.
[About Oracle Enterprise Manager](#)

Note: The e-mail address you specify on this page is not the e-mail address to which the notification is sent. You will have to specify the e-mail address (where notifications will be sent) from the Preferences General page.

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 and spread the load across multiple systems. The Management Service makes use of each mail server to send e-mails and the behavior is controlled by the following parameters found in the `$ORACLE_HOME/sysman/config/emoms.properties` file.

Example 3–2 Management Service Parameters

```
# The maximum number of emails that can be sent in a single connection to an
# email server
# em.notification.emails_per_connection=20
#
# The maximum number of emails that can be sent in a minute
```

```
# em.notification.emails_per_minute=250
```

Based on the defaults in [Example 3-2](#), the first mail server is used to send 20 e-mails before the Management Service switches to the next mail server which is used to send another 20 e-mails before switching to the next mail server. This prevents one mail server from becoming overloaded and should improve overall reliability and throughput.

3.1.1.1 Setting Up Repeat Notifications

Repeat notifications allow administrators to be notified repeatedly until an alert is either acknowledged or the number of **Maximum Repeat Notifications** has been reached. Enterprise Manager supports repeat notification for all notification methods (e-mail, OS command, PL/SQL procedure, and SNMP trap). To enable this feature for a notification method, select the **Send Repeat Notifications** option. In addition to setting the maximum number of repeat notifications, you can also set the time interval at which the notifications are sent.

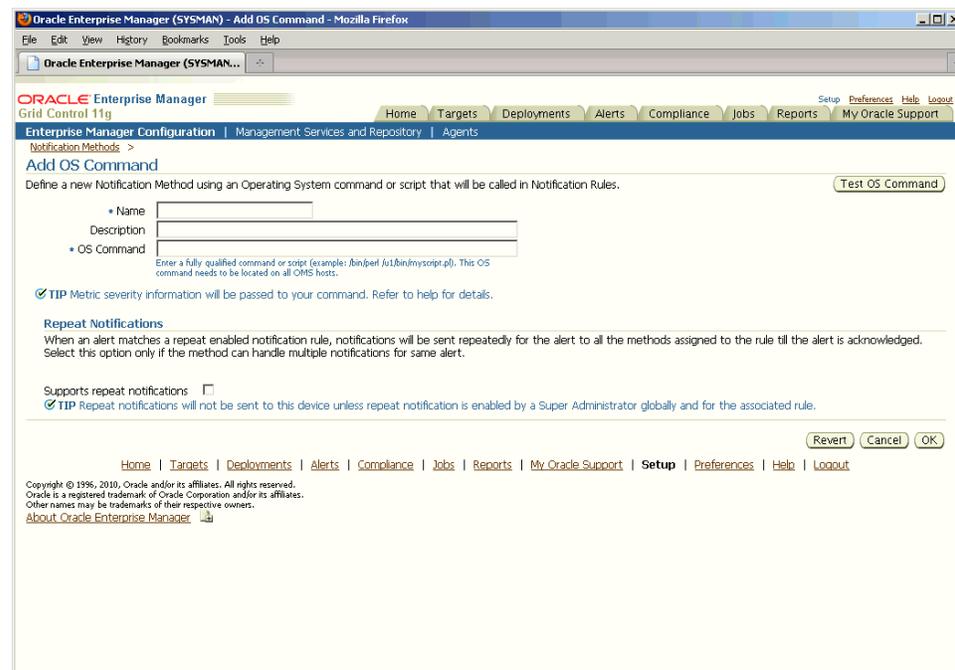
Important: If the Grid Control Repository database version is 9.2, the `aq_tm_processes` `init.ora` parameter must be set to at least 1 to enable repeat notification functionality.

Repeat Notifications for Rules

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

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

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

Figure 3–2 Enabling Repeat Notification for an OS Command Notification Method

3.1.2 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 General page under the Preferences link in the Grid Control console. In addition to defining notification e-mail addresses, you associate the notification message format (long or short) to be used for your e-mail address.

Setting up e-mail involves three steps:

Step 1: Define e-mail addresses.

Step 2: Set up a Notification Schedule.

Step 3: Subscribe to receive e-mail for notification rules.

3.1.2.1 Defining E-mail Addresses

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

To add an e-mail address:

1. From the Grid Control console, click **Preferences**. By default the General page is selected.
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 message format for your e-mail address. The Long Format sends a HTML formatted e-mail that contains detailed information. [Example 3–3](#) shows a typical notification that uses the long format.

The Short Format ([Example 3-4](#)) sends 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.

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

Example 3-3 Long E-mail Notification for Alerts

```
Name=myhost.com
Type=Host
Host=myhost.com
Metric=Filesystem Space Available (%)
Mount Point =/usr
Timestamp=06-OCT-2006 16:27:05 US/Pacific
Severity=Warning
Message=Filesystem / has only 76.07% available space
Rule Name=Host Availability and Critical States
Rule Owner=SYSMAN
```

Example 3-4 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 Short E-mail Format

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.

Entries in the `emoms.properties` file define the size and format of the short e-mail.

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

emoms.properties Entries for a Short E-mail Format

```
# The maximum size of a short format email
# em.notification.short_format_length=155
# The format of the short email. It can be set to subject, body or both.
#
# When set to subject the entire message is sent in the subject i.e.
# EM:<severity>:<target>:<message>:<timestamp>
# When set to body the entire message is sent in the body i.e.
# EM:<severity>:<target>:<message>:<timestamp>
# When set to both the message is split i.e. the subject contains
# EM:<severity>:<target>
# and the body contains
# <message>:<timestamp>
# In all cases the message is truncated to the length specified in the
# em.notification.short_format_length parameter
# em.notification.short_format=both
```

3.1.2.2 Setting Up a Notification Schedule

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

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

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

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

1. From the Enterprise Manager Grid Control, click **Preferences** at the top of the page. By default the General page is selected.
2. Click **Notification Schedule** in the vertical navigation bar. Your Notification Schedule page appears.
3. Follow the directions on the Notification Schedule page to specify when you want to receive e-mails.

3.1.2.3 Subscribe to Receive E-mail for Notification Rules

A notification rule is a user-defined rule that defines the criteria by which notifications should be sent for alerts, policy violations, corrective action execution status, and job execution status. Specifically, in each 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. During notification rule creation, you specify criteria such as the targets you are interested in, their monitored metrics, associated alert severity conditions (clear, warning, critical), policy violations, corrective action execution status, or job execution status, and the associated notification method.

To subscribe to a notification rule you create, while creating the rule, go to the Actions page and check the **Send Me E-mail** option.

Out-of-Box Notification Rules

Enterprise Manager Grid control comes with out-of-box notification rules that cover the most common alert conditions. When you install the Oracle Management Service, you are given the option to receive e-mail notifications for critical alerts. If you choose

this option, and if an e-mail address for the SYSMAN user was specified, then some default notification rules are created that cover the availability and critical states for common target types and would also be configured to send e-mail notifications to the SYSMAN e-mail address for the conditions defined in the notification rules.

You can access the out-of-box notification rules by clicking on Preferences on any page in the Enterprise Manager console and clicking Public Rules in the vertical navigation bar. If the conditions defined in the out-of-box notification rules meet your requirements, you can simply subscribe to receive e-mail notifications for the conditions defined in the rule by clicking on Subscribe column in the row of the Public Rules table that corresponds to the notification rule that you are interested in. Click **Apply** to save your changes.

[Table 3–1](#) lists all default notification rules. These are all owned by the SYSMAN user and are public rules.

Table 3–1 Default Notification Rules

Name	Description	Applies to Targets of the Type	Send Notification on the Following Availability States	Send Notification if Any of the Metrics is at CRITICAL Alert Severity
Agent Upload Problems	System-generated notification rule for monitoring Agents that may have problems uploading data to the Management Service.	Oracle Management Service and Repository	N/A	Count of targets not uploading data
Agents Unreachable	System-generated notification rule for monitoring Agents that lose contact with the Management Service due to network problems, host problems or Agents going down.	Agents	Agent Unreachable Agent Unreachable Resolved	N/A
Application Server Availability and Critical States	System-generated notification rule for monitoring Application Servers' availability, and critical metric statuses.	Application Servers	Down	CPU Usage (%)

Table 3–1 (Cont.) Default Notification Rules

Name	Description	Applies to Targets of the Type	Send Notification on the Following Availability States	Send Notification if Any of the Metrics is at CRITICAL Alert Severity
Database Availability and Critical States	System-generated notification rule for monitoring Databases' availability, and critical metric statuses.	Databases (single instance only)	Down	Process Limit Usage (%) Session Limit Usage (%) Blocking Session Count All Objects Archiver Hung Alert Log Error Status Data Block Corruption Alert Log Error Status Generic Alert Log Error Status Media Failure Alert Log Error Status Session Terminated Alert Log Error Status Archive Area Used (%) All Objects Segments Not Able to Extend Count All Objects Segments Approaching Maximum Extents Count All Objects Tablespace Space Used (%) All Objects Wait Time (%)
HTTP Server Availability and Critical States	System-generated notification rule for monitoring HTTP Server's availability, and critical metric statuses.	Oracle HTTP Server	Down	CPU Usage (%) Percentage of Busy Processes Active HTTP Connections Request Processing Time (seconds)
Host Availability and Critical States	System-generated notification rule for monitoring Hosts' availability, and critical metric statuses.	Hosts	Agent Unreachable Agent Unreachable Resolved	Average Disk I/O Service Time (ms) Disk Device Busy (%) Filesystem Space Available (%) CPU in I/O Wait (%) Run Queue Length (5 minute average) CPU Utilization (%) Memory Utilization (%) Memory Page Scan Rate (per second) Swap Utilization (%) Network Interface Combined Utilization (%)

Table 3–1 (Cont.) Default Notification Rules

Name	Description	Applies to Targets of the Type	Send Notification on the Following Availability States	Send Notification if Any of the Metrics is at CRITICAL Alert Severity
Listener Availability	System-generated notification rule for monitoring database Listeners' availability, and critical metric statuses.	Listeners	Down	N/A
Misconfigured Agents	System-generated notification rule for misconfigured agents.	Agent	Agent Unreachable Agent Unreachable Resolved	Consecutive severity upload failure count Consecutive heartbeat failure count MS Agent time skew (mins) Consecutive metadata upload failure count
OC4J Availability and Critical States	System-generated notification rule for monitoring OC4J instance's availability, and critical metric statuses.	OC4J	Down	CPU Usage (%) OC4J Instance - Request Processing Time (seconds) OC4J Instance - Active Sessions
OMS Service Initialization Errors	System-generated notification rule for monitoring OMS service initialization errors.	OMS and Repository	N/A	Service Status
PAF Status Notification	System-generated notification rule for Provisioning Advisor Framework: Notifies the instance creator of any status updates.	N/A	Up Down Corrective Actions on Target Down Agent Unreachable Agent Unreachable Resolved Metric Error Detected Metric Error Resolved Blackout Started Blackout Ended	N/A
Repository Operations Availability	System-generated notification rule for monitoring the availability of the DBMS jobs that are part of the Management Repository.	OMS and Repository	Critical	DBMS Job UpDown

Table 3–1 (Cont.) Default Notification Rules

Name	Description	Applies to Targets of the Type	Send Notification on the Following Availability States	Send Notification if Any of the Metrics is at CRITICAL Alert Severity
Violation Notification for Database Security Policies	System-generated notification rule for monitoring the secureness of the database configuration.	Databases	Critical	N/A
Web Cache Availability and Critical States	System-generated notification rule for monitoring Web Cache's instance's availability, and critical metric statuses.	Oracle Web Cache	Down	Hits (% of requests) Web Cache CPU Usage (%)

Creating Your Own Notification Rules

If you find that the default notification rules do not meet your needs, you can define your own custom rules. The following procedure documents the process of notification rule creation for non-Super Administrators.

To create your own notification rule:

1. From the Enterprise Manager Grid Control, click **Preferences**.
2. Click **My Rules** in the vertical navigation bar.

If you are not logged in as an administrator with Super Administrator privileges, you will see a link for **My Rules** instead of **Rules** as in the case of an administrator with Super Administrator privileges.

3. Click **Create**.

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

When you specify the notification rule properties, check **Make Public** in the General page if you want other non-privileged users to be able to view and share that rule. For example, it allows other administrators to later specify that they should receive e-mail for this rule.

When you specify the notification rule, you will only be able to choose from e-mail and SNMP traps. Specifying custom commands and PL/SQL procedures is an option that is only available to Super Administrators. To receive e-mail notifications for conditions defined in the rule, go to the Actions page and check the **Send Me E-Mail** option.

Specifying Additional Alert Duration Criteria

You can set additional alert duration criteria for a notification rule and have the rule apply only to alerts that have been open for at least a certain amount of time and have not been acknowledged. These criteria apply only to Target Down, Agent Unreachable, Metric alerts, Policy Violations, Blackout Started and Metric Error Start alerts.

Typical scenarios where you would use additional alert criteria are:

- For log alerts that have been open for at least 7 days, clear the alerts

- For alerts that have been open for at least 48 hours and have not been acknowledged, send e-mail to the DBA Manager

To specify additional alert duration criteria:

1. Create or edit a notification rule. Additional alert criteria can be added from the **Availability, Metrics, or Policy** tab.
2. From one of the aforementioned tabs, go to the **Additional Alert Criteria** section and click **Add**. The Additional Alert Criteria page appears.
3. Specify the alert duration criteria and click **Continue**.

3.1.3 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. Click **Setup**.
2. Click **Administrators** from the vertical navigation bar.
3. 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. Click **Setup**.
2. From the vertical navigation bar, click **Schedules** (under **Notification**). The **Notification Schedule** page appears.
3. Specify the administrator whose notification schedule you wish to edit and click **Change**.
4. Click **Edit Schedule Definition**. The **Edit Schedule Definition: Time Period** page appears. If necessary, modify the rotation schedule.
5. Click **Continue**. The **Edit Schedule Definition: E-mail Addresses** page appears.
6. Follow the directions on the **Edit Schedule Definition: E-mail Addresses** page to modify the notification schedule.
7. Click **Finish** when you are done.
8. Repeat steps three through seven for each administrator.

Step 3: Assign Notification Rules to Administrators

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

1. Click **Setup**.
2. From the vertical navigation bar, click **Administrators**.
3. Select the desired administrator.
4. Click **Subscribe to Rules**. The **Subscribe <administrator> to Public Notification Rules** page appears.
5. Select the desired notification rules and click **Subscribe**.
6. Click **OK** when you are finished.
7. Repeat steps three through six for each administrator.

3.1.4 E-mail Customization

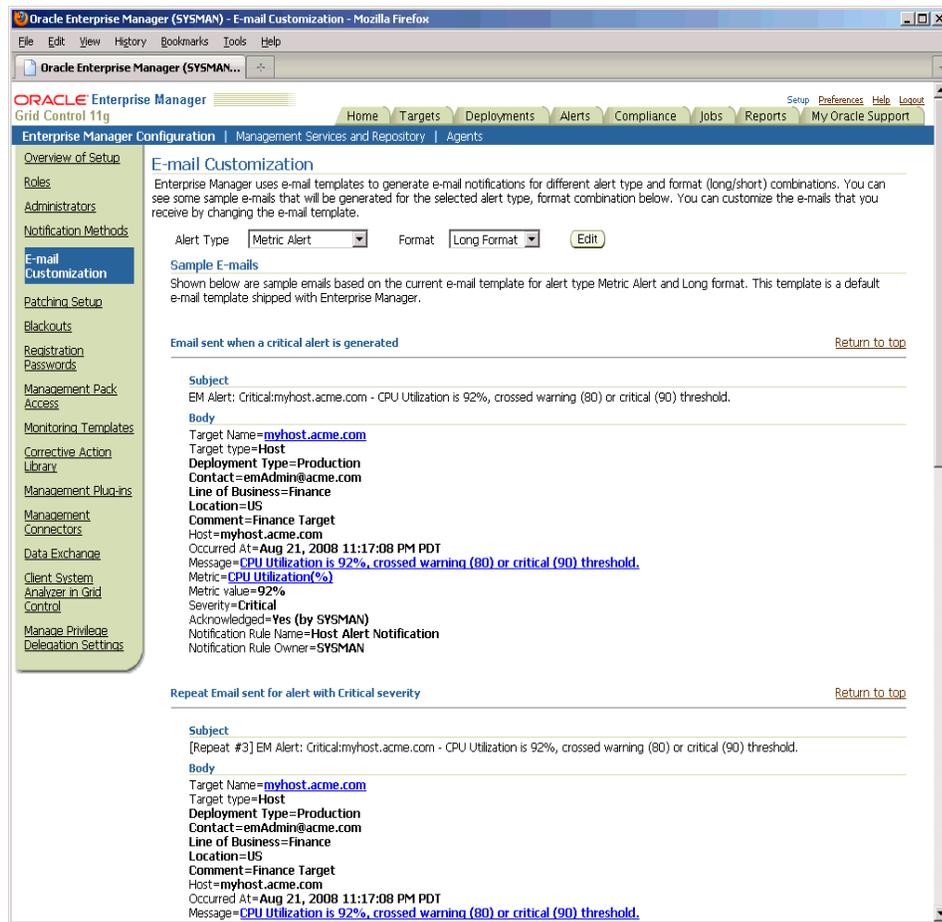
Enterprise Manager allows Super Administrators to customize global e-mail notifications for the four alert types (Metric Alert, Target Availability, Policy Violation, and Job Status Change). Using predefined building blocks (called attributes and labels) contained within a simple script, Super Administrators can customize alert e-mails by selecting from a wide variety of information content.

To customize an e-mail:

1. Access the E-mail Customization page. Setup-->E-mail Customization
2. Choose the **Alert Type** and **Format**.
3. Click **Edit**. The Edit E-mail Template page is displayed.

From the Edit 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 3–3 E-mail Customization



3.1.4.1 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 3–2 Reserved Words and Operators

Reserved Word/Operator	Description
IF, ELSIF, ENDIF, ELSE	Used in IF-ELSE constructs.
AND, OR	Boolean operators – used in IF-ELSE constructs only.
NULL	To check NULL value for attributes - used in IF-ELSE constructs only.
	Pipe operator – used to show the first non-NULL value in a list of attributes. For example: <code>METRIC_NAME POLICY_NAME</code>
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 [JOB_NAME]. For a metric e-mail 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 3–3 Control Structures

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

Table 3-3 (Cont.) Control Structures

Control Structure	Description
IF	<p data-bbox="764 260 1446 338">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 data-bbox="764 352 1446 430">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 data-bbox="764 445 1446 552">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 data-bbox="764 567 1446 619">[IF REPEAT_COUNT EQ "1" AND SEVERITY EQ "CRITICAL" THEN]</pre> <p data-bbox="764 634 1446 682">The statement above will be true when the attributes of the alert match the following condition:</p> <ul data-bbox="764 697 1446 842" style="list-style-type: none"> ■ Attribute Name: REPEAT_COUNT ■ Attribute Value: 1 ■ Attribute Name: SEVERITY ■ Attribute Value: CRITICAL <p data-bbox="764 856 954 882">Example IF Block:</p> <pre data-bbox="764 896 1446 1010">[IF JOB_NAME NEQ NULL] [JOB_NAME_LABEL] = [JOB_NAME] [JOB_OWNER_LABEL] = [JOB_OWNER] [ENDIF]</pre> <pre data-bbox="764 1045 1446 1211">[IF SEVERITY EQ CRITICAL] [MTRIC_NAME_LABEL] = [METRIC_NAME] [METRIC_VALUE_LABEL] = [METRIC_VALUE] [TARGET_NAME_LABEL] = [TARGET_NAME] [KEY_VALUES] [ENDIF]</pre> <p data-bbox="764 1264 1084 1289">Example IF and ELSEIF Block:</p> <pre data-bbox="764 1304 1446 1684">[IF SEVERITY EQ CRITICAL] statement1 [ELSIF SEVERITY EQ WARNING] statement2 [ELSIF SEVERITY 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: |

3.2 Extending Notification Beyond E-mail

Notification Methods are the mechanisms by which alerts 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 setup as part of the Oracle Management Service installation.

Enterprise Manager Super Administrators can also define other custom notification methods. For example, alerts 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 that is defined, any time an administrator needs to send alerts to the trouble-ticketing system, he merely 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) to their notification rules.

Through the Notification Methods page, you can:

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

3.2.1 Custom Notification Methods Using Scripts and SNMP Traps

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

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

Complete the following four steps to define a notification method based on an OS command or script.

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

Step 1: Define your OS command or script.

You can specify an OS command or script that will be called by the notification system. You can use target and alert or policy violation context information, corrective action execution status and job execution status within the body of the script. Passing metric severity attributes (severity level, type, notification rule, rule owner, or rule owner, and so on) or policy violation information to OS commands/scripts allows you to customize automated responses to alerts or policy violations. 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 Alert and Policy Violation Information to an OS Command or Script"](#) on page 3-20
- ["Passing Corrective Action Execution Status to an OS Command or Script"](#) on page 3-29
- ["Passing Job Execution Status to an OS Command or Script"](#) on page 3-30

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

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

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

[Example 3-5](#) shows the parameter in emoms.properties that controls how long the OS Command can execute without being killed by the Management Service. This is to prevent OS Commands from running for an inordinate length of time and blocking the delivery of other notifications. By default the command is allowed to run for 30 seconds before it is killed.

Example 3-5 Parameter in emoms.properties File

```
# The amount of time in seconds after which an OS Command started by the
# Notification System will be killed if it has not exited
# em.notification.os_cmd_timeout=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 Notification Rules. Log in as a Super Administrator, click Setup and then Notification Methods from the vertical navigation bar. From this page, you can define a new notification based on the 'OS Command' type. See ["Adding a Notification Method based on an OS Command or Script"](#) on page 3-19.

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 3-6](#) shows information required for the notification method.

Example 3-6 OS Command Notification Method

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

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

Step 4: Assign the notification method to a rule.

You can edit an existing rule (or create a new notification rule), then go to the Methods page. In the list of Advanced Notification Methods, select your notification method and click 'Assign Method to Rule'. To assign multiple rules to a method or methods to a single rule, see ["Assigning Rules to Methods"](#) on page 3-36 or ["Assigning Methods to Rules"](#) on page 3-35.

Passing Alert and Policy Violation Information to an OS Command or Script

The notification system passes severity 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.

Table 3–4 Environment Variables

Environment Variable	Description
TARGET_NAME	Name of the target on which the severity occurred.
TARGET_TYPE	Type of target on which the severity occurred. Targets are defined as any monitorable entity, such as Host, Database, Listener, or Oracle HTTP Server. You can view the type of a monitored target on the All Targets page.
HOST	Name of the machine on which the target resides.
METRIC	Metric generating the severity. This variable is not set for policy violations.
METRIC_VALUE	The value of the metric when the threshold was exceeded. Not set for policy violations
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.
POLICY_RULE	The name of the policy when the threshold was exceeded. Not set for metric severities
KEY_VALUE	For metrics that monitor a set of objects, the KEY_VALUE indicates the specific object that triggered the severity. For example for the Tablespace Space Used (%) metric that monitors tablespace objects, the KEY_VALUE is 'USERS' if the USERS tablespace triggered at warning or critical severity.
KEY_VALUE_NAME	For metrics that monitor a set of objects, the KEY_VALUE_NAME indicates the type of object monitored. For example for the Tablespace Space Used (%) metric that monitors tablespace objects, the KEY_VALUE_NAME is 'Tablespace Name'.
VIOLATION_CONTEXT	A comma-separated list of name-value pairs that shows the alert context for a policy violation.
TIMESTAMP	Time when the severity occurred.
SEVERITY	Type of severity. For example, severity for a target's (availability) status metric are: <ul style="list-style-type: none"> ■ UP ■ DOWN ■ UNREACHABLE CLEAR ■ UNREACHABLE START ■ BLACKOUT END ■ BLACKOUT START Other metrics can have any of the following severities: <ul style="list-style-type: none"> ■ WARNING ■ CRITICAL ■ CLEAR ■ METRIC ERROR CLEAR ■ METRIC ERROR START
MESSAGE	Message for the alert that provides details about what triggered the condition.
RULE_NAME	Name of the notification rule to which the OS Command notification method was assigned.

Table 3-4 (Cont.) Environment Variables

Environment Variable	Description
RULE_OWNER	Name of the Enterprise Manager administrator who owns the notification rule.

Your script may reference some or all of these variables.

The sample OS script shown in [Example 3-7](#) 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 3-7 Sample OS Command Script

```
#!/bin/ksh

LOG_FILE=/net/myhost/logs/severity.log
if test -f $LOG_FILE
then
echo $TARGET_NAME $MESSAGE $TIMESTAMP >> $LOG_FILE
else
    exit 100
fi
```

[Example 3-8](#) shows an OS script that logs alert information to the file 'alertmsg.txt'. The file is saved to the /u1/results directory.

Example 3-8 Alert Logging Script

```
#!/usr/bin/sh
echo "Alert logged:" > /u1/results/alertmsg.txt
echo "\n" >> /u1/results/alertmsg.txt
echo "target name is " $TARGET_NAME >> /u1/results/alertmsg.txt
echo "target type is " $TARGET_TYPE >> /u1/results/alertmsg.txt
echo "target is on host " $HOST >> /u1/results/alertmsg.txt
echo "metric in alert is " $METRIC >> /u1/results/alertmsg.txt
echo "metric index is " $KEY_VALUE >> /u1/results/alertmsg.txt
echo "timestamp is " $TIMESTAMP >> /u1/results/alertmsg.txt
echo "severity is " $SEVERITY >> /u1/results/alertmsg.txt
echo "message is " $MESSAGE >> /u1/results/alertmsg.txt
echo "notification rule is " $RULE_NAME >> /u1/results/alertmsg.txt
echo "rule owner is " $RULE_OWNER >> /u1/results/alertmsg.txt
exit 0
```

[Example 3-9](#) shows a script that sends an alert to an HP OpenView console from Enterprise Manager Grid Control. When a metric alert is triggered, the Enterprise Manager Grid Control displays the alert. The HP OpenView script is then called, invoking opcmmsg and forwarding the information to the HP OpenView management server.

Example 3-9 HP OpenView Script

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

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

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

Step 1: Define the PL/SQL procedure.

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

For alerts and policy violations:

```
PROCEDURE p(severity IN MGMT_NOTIFY_SEVERITY)
```

For job execution status changes:

```
PROCEDURE p(job_status_change IN MGMT_NOTIFY_JOB)
```

For corrective action status changes:

```
PROCEDURE p(ca_status_change IN MGMT_NOTIFY_CORRECTIVE_ACTION)
```

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 notification rule.

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

["Passing Alert and Policy Violation Information to a PL/SQL Procedure"](#) on page 3-24

["Passing Corrective Action Status Change Information"](#) on page 3-28

["Passing Job Execution Status Information"](#) on page 3-32

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 p(severity IN MGMT_NOTIFY_SEVERITY)
```

```
PROCEDURE p(job_status_change IN MGMT_NOTIFY_JOB)
```

```
PROCEDURE p(ca_status_change IN MGMT_NOTIFY_CORRECTIVE_ACTION)
```

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. See [Example 3-11, "PL/SQL Procedure Using a Severity Code"](#). The procedure initially retries 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, click Setup and then Notification Methods from the vertical navigation bar. From this page, you can define a new notification based on 'PL/SQL Procedure'. See ["Adding a Notification Method Based on a PL/SQL Procedure"](#) on page 3-23.

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 3–10](#).

Example 3–10 PL/SQL Procedure Required Information

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

Step 4: Assign the notification method to a rule.

You can edit an existing rule (or create a new notification rule), then go to the Methods page. In the list of Advanced Notification Methods, select your notification method and click 'Assign Method to Rule'. To assign multiple rules to a method or methods to a single rule, see "[Assigning Rules to Methods](#)" on page 3-36 or "[Assigning Methods to Rules](#)" on page 3-35.

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

Information about the severity types that relate to a target's availability, and how metric severity and policy violation information is passed to the PLSQL procedure is covered in the next section.

Passing Alert and Policy Violation Information to a PL/SQL Procedure

Passing metric severity attributes (severity level, type, notification rule, rule owner, or rule owner, and so on) or policy violation information to PL/SQL procedures allows you to customize automated responses to alerts or policy violations.

The notification system passes information about metric severities or policy violations to a PL/SQL procedure using the MGMT_NOTIFY_SEVERITY object. An instance of this object is created for every alert or policy violation. When an alert or policy violation occurs, the notification system calls the PL/SQL procedure associated with the notification rule and passes the populated object to the procedure. The procedure is then able to access the fields of the MGMT_NOTIFY_SEVERITY object that has been passed to it.

The following table lists all metric severity attributes that can be passed:

Table 3–5 Metric Severity Attributes

Attribute	Datatype	Additional Information
TARGET_NAME	VARCHAR2(256)	Name of the target on which the severity occurred.

Table 3-5 (Cont.) Metric Severity Attributes

Attribute	Datatype	Additional Information
TARGET_TYPE	VARCHAR2(64)	Type of target on which the severity occurred. Targets are defined as any monitorable service.
TIMEZONE	VARCHAR2(64)	The target's regional timezone
HOST_NAME	VARCHAR2(128)	Name of the machine on which the target resides.
METRIC_NAME	VARCHAR2(64)	Metric or policy generating the severity.
METRIC_DESCRIPTION	VARCHAR2(128)	Meaningful description of the metric that can be understood by other administrators.
METRIC_COLUMN	VARCHAR2(64)	For table metrics, the metric column contains the name of the column in the table that is being defined. If the metric that is being defined is not a table metric, the value in this column is a single space. This attribute is not used for policy violations.
METRIC_VALUE	VARCHAR2(1024)	The value of the metric.
KEY_VALUE	VARCHAR2(1290)	For metrics that monitor a set of objects, the KEY_VALUE indicates the specific object that triggered the severity. For example for the Tablespace Space Used (%) metric that monitors tablespace objects, the KEY_VALUE is 'USERS' if the USERS tablespace triggered at warning or critical severity.
KEY_VALUE_NAME	VARCHAR2(512)	For metrics that monitor a set of objects, the KEY_VALUE_NAME indicates the type of object monitored. For example for the Tablespace Space Used (%) metric that monitors tablespace objects, the KEY_VALUE_NAME is 'Tablespace Name'.
KEY_VALUE_GUID	VARCHAR2(256)	GUID associated with a composite key value name.
CTXT_LIST	MGMT_NOTIFY_COLUMNS	Details of the alert context.
COLLECTION_TIMESTAMP	DATE	The time when the target status change was last detected and logged in the management repository.
SEVERITY_CODE	NUMBER	Numeric code identifying the severity level. See Severity Code table below.
MESSAGE	VARCHAR2(4000)	An optional message that is generated when the alert is created that provides additional information about the alert condition.
SEVERITY_GUID	RAW(16)	Severity global unique identifier.
METRIC_GUID	RAW(16)	Metric global unique identifier.
TARGET_GUID	RAW(16)	Target global unique identifier.
RULE_OWNER	VARCHAR2(64)	Name of the Enterprise Manager administrator who owns the rule.
RULE_NAME	VARCHAR2(132)	Name of the notification rule that resulted in the severity.

When a severity occurs for the target, the notification system creates an instance of the MGMT_NOTIFY_SEVERITY object and populates it with values from the severity. The severity codes in Table 3-6 have been defined as constants in the MGMT_GLOBAL package and can be used to determine the type of severity in the severity_code field of the MGMT_NOTIFY_SEVERITY object.

Table 3-6 Severity Codes

Name	Datatype	Value
G_SEVERITY_COMMENT	NUMBER(2)	10
G_SEVERITY_CLEAR	NUMBER(2)	15
G_SEVERITY_WARNING	NUMBER(2)	20
G_SEVERITY_CRITICAL	NUMBER(2)	25
G_SEVERITY_UNREACHABLE_CLEAR	NUMBER(3)	115
G_SEVERITY_UNREACHABLE_START	NUMBER(3)	125
G_SEVERITY_BLACKout_END	NUMBER(3)	215
G_SEVERITY_BLACKout_START	NUMBER(3)	225
G_SEVERITY_ERROR_END	NUMBER(3)	315
G_SEVERITY_ERROR_START	NUMBER(3)	325
G_SEVERITY_NO_BEACONS	NUMBER(3)	425
G_SEVERITY_UNKNOWN	NUMBER(3)	515

Example 3-11 PL/SQL Procedure Using a Severity Code

```
CREATE TABLE alert_log (target_name VARCHAR2(64),
alert_msg VARCHAR2(4000),
occured DATE);

PROCEDURE LOG_CRITICAL_ALERTS(severity IN MGMT_NOTIFY_SEVERITY)
IS
BEGIN
-- Log all critical severities
IF severity.severity_code = MGMT_GLOBAL.G_SEVERITY_CRITICAL
THEN
BEGIN
INSERT INTO alert_log (target_name, alert_msg, occured)
VALUES (severity.target_name, severity.message,
severity.collection_timestamp);
EXCEPTION
WHEN OTHERS
THEN
-- If there are any problems then get the notification retried
RAISE_APPLICATION_ERROR(-20000, 'Please retry');
END;
COMMIT;
END IF;
END LOG_CRITICAL_ALERTS;
```

3.2.1.3 Adding a Notification Method Based on an SNMP Trap

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

The trap is an SNMP Version 1 trap and is described by the MIB definition shown at the end of this chapter. See "[Management Information Base \(MIB\)](#)" on page 3-37.

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

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

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

Log in to Enterprise Manager as a Super Administrator. Click Setup and then click Notification Method from the vertical navigation bar to access the Notification Methods page. From this page you can add a new method based on an SNMP trap.

You must provide the name of the host (machine) on which the SNMP master agent is running and other details as shown in the following example. In [Example 3-12](#), the SNMP host will receive your SNMP traps.

Example 3-12 *SNMP Trap Required Information*

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

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

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

An example SNMP Trap is shown in [Example 3-13](#). Each piece of information is sent as a variable embedded in the SNMP Trap.

Example 3-13 *SNMP Trap*

```
Tue Oct 28 05:00:02 2006

Command: 4
Enterprise: 1.3.6.1.4.1.111.15.2
Agent: 138.1.6.200
Generic Trap: 6
Specific Trap: 1
Time Stamp: 8464:39.99
Count: 11

Name: 1.3.6.1.4.1.111.15.1.1.1.2.1
Kind: OctetString
Value: "mydatabase"

Name: 1.3.6.1.4.1.111.15.1.1.1.3.1
Kind: OctetString
Value: "Database"
```

```
Name: 1.3.6.1.4.1.111.15.1.1.1.4.1
Kind: OctetString
Value: "myhost.com"

Name: 1.3.6.1.4.1.111.15.1.1.1.5.1
Kind: OctetString
Value: "Owner's Invalid Object Count"

Name: 1.3.6.1.4.1.111.15.1.1.1.6.1
Kind: OctetString
Value: "Invalid Object Owner"

Name: 1.3.6.1.4.1.111.15.1.1.1.7.1
Kind: OctetString
Value: "SYS"

Name: 1.3.6.1.4.1.111.15.1.1.1.8.1
Kind: OctetString
Value: "28-OCT-2006 04:59:10 (US/Eastern GMT)"

Name: 1.3.6.1.4.1.111.15.1.1.1.9.1
Kind: OctetString
Value: "Warning"

Name: 1.3.6.1.4.1.111.15.1.1.1.10.1
Kind: OctetString
Value: "12 object(s) are invalid in the SYS schema."

Name: 1.3.6.1.4.1.111.15.1.1.1.11.1
Kind: OctetString
Value: "Database Metrics"

Name: 1.3.6.1.4.1.111.15.1.1.1.12.1
Kind: OctetString
Value: "SYSMAN"
```

Step 2: Assign the notification method to a rule.

You can edit an existing rule (or create a new notification rule), then go to the Methods page. In the list of Advanced Notification Methods, select your notification method and click 'Assign Method to Rule'. To assign multiple rules to a method or methods to a rule, see ["Assigning Methods to Rules"](#) on page 3-35 or ["Assigning Rules to Methods"](#) on page 3-36.

3.3 Passing Corrective Action Status Change Information

Passing corrective action status change attributes (such as new status, job name, job type, notification rule, 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.

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

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

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

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

Table 3–7 Environment Variables

Environment Variable	Description
JOB_NAME	The name of the corrective action.
JOB_OWNER	The owner of the corrective action.
JOB_TYPE	The type of corrective action.
JOB_STATUS	The corrective action status.
TIMESTAMP	Time when the severity occurred.
NUM_TARGETS	The number of targets.
TARGET_NAME _n	The name of the <i>n</i> th target on which the corrective action ran. Example: TARGET_NAME1, TARGET_NAME2.
METRIC	The name of the metric in the alert that caused the corrective action to run. Not set for policy violations.
POLICY_RULE	The name of the policy rule in the alert that caused the corrective action to run. Not set for metric severities.
METRIC_VALUE	The value of the metric column in the alert that caused the corrective action to run.
VIOLATION_CONTEXT	A comma-separated list of name-value pairs that show the policy violation context.
KEY_VALUE_NAME	For metrics that monitor a set of objects, the KEY_VALUE_NAME indicates the type of object monitored. For example for the Tablespace Space Used (%) metric that monitors tablespace objects, the KEY_VALUE_NAME is 'Tablespace Name'.
KEY_VALUE	For metrics that monitor a set of objects, the KEY_VALUE indicates the specific object that triggered the severity. For example for the Tablespace Space Used (%) metric that monitors tablespace objects, the KEY_VALUE is 'USERS' if the USERS tablespace triggered at warning or critical severity.

Table 3–7 (Cont.) Environment Variables

Environment Variable	Description
SEVERITY	Type of alert severity. For example, severity types that relate to a target's availability are: <ul style="list-style-type: none"> ▪ UP ▪ DOWN ▪ UNREACHABLE CLEAR ▪ UNREACHABLE START ▪ BLACKOUT END ▪ BLACKOUT START Other metrics can have any of the following severities: <ul style="list-style-type: none"> ▪ WARNING ▪ CRITICAL ▪ CLEAR ▪ METRIC ERROR CLEAR ▪ METRIC ERROR START
RULE_NAME	Name of the notification rule that resulted in the execution of the corrective action.
RULE_OWNER	Name of the Enterprise Manager administrator who owns the notification rule.

3.3.2 Passing Corrective Action Execution Status to a PLSQL Procedure

The notification system passes corrective action status change information to a PL/SQL procedure via the MGMT_NOTIFY_CORRECTIVE_ACTION object. An instance of this object is created for every status change. When a corrective action executes, the notification system calls the PL/SQL procedure associated with the notification rule and passes the populated object to the procedure. The procedure is then able to access the fields of the MGMT_NOTIFY_CORRECTIVE_ACTION object that has been passed to it.

Table 3–8 lists all corrective action status change attributes that can be passed:

Table 3–8 Corrective Action Status Attributes

Attribute	Datatype	Additional Information
JOB_NAME	VARCHAR2(128)	The corrective action name.
JOB_OWNER	VARCHAR(256)	The owner of the corrective action.
JOB_TYPE	VARCHAR2(32)	The type of the corrective action.
JOB_STATUS	NUMBER	The new status of the corrective action. See Table 3–9, "Corrective Action Status Codes" for a list of possible status conditions.
STATE_CHANGE_GUID	RAW(16)	The GUID of the state change record.
JOB_GUID	RAW(16)	The unique id of the corrective action.
EXECUTION_ID	RAW(16)	The unique id of the corrective action execution.
TARGETS	SMP_EMD_NVPAIR_ARRAY	An array of the target name/target type pairs that the corrective action runs on.

Table 3–8 (Cont.) Corrective Action Status Attributes

Attribute	Datatype	Additional Information
METRIC_NAME	VARCHAR2(256)	The name of the metric/policy rule in the alert that caused the corrective action to run.
METRIC_COLUMN	VARCHAR2(64)	The name of the metric in the alert that caused the corrective action to run. This is not used for policy violations.
METRIC_VALUE	VARCHAR2(1024)	The value of the metric in the alert that caused the corrective action to run.
SEVERITY_CODE	NUMBER	The severity code of the alert that caused the corrective action to run. See Table 3–6, "Severity Codes" .
KEY_VALUE_NAME	VARCHAR2(512)	For metrics that monitor a set of objects, the KEY_VALUE_NAME indicates the type of object monitored. For example for the Tablespace Space Used (%) metric that monitors tablespace objects, the KEY_VALUE_NAME is 'Tablespace Name'.
KEY_VALUE	VARCHAR2(1290)	For table metrics, this column contains the value of the key column for the row in the table whose thresholds are being defined. If the thresholds are not for a table metric, or the thresholds apply for all rows in the metric column, then the value in this column will contain a single space.
KEY_VALUE_GUID	RAW(16)	GUID associated with a composite key value name.
CTXT_LIST	MGMT_NOTIFY_COLUMNS	Details of the corrective action status change context.
RULE_OWNER	VARCHAR2(64)	The owner of the notification rule that caused the PL/SQL notification to be sent.
RULE_NAME	VARCHAR2(132)	The name of the notification rule that caused the PL/SQL notification method to be invoked.
OCCURRED_DATE	DATE	The time and date when the status change happened.

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

Table 3–9 Corrective Action Status Codes

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

Table 3–9 (Cont.) Corrective Action Status Codes

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

Example 3–14 PL/SQL Procedure Using a Status Code

```

CREATE TABLE ca_log (jobid RAW(16),
    occured DATE);

CREATE OR REPLACE PROCEDURE LOG_PROBLEM_CAS(status_change IN MGMT_NOTIFY_
CORRECTIVE_ACTION)
IS
BEGIN
-- Log all failed corrective actions
    IF status_change.job_status = MGMT_JOBS.FAILED_STATUS
    THEN
        BEGIN
            INSERT INTO ca_log (jobid, occured)
            VALUES (status_change.job_guid, SYSDATE);
        EXCEPTION
        WHEN OTHERS
        THEN
            -- If there are any problems then get the notification retried
            RAISE_APPLICATION_ERROR(-20000, 'Please retry');
        END;
    COMMIT;
    END IF;
END LOG_PROBLEM_CAS;
    
```

3.4 Passing Job Execution Status Information

Passing job status change attributes (such as new status, job name, job type, notification rule, 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.

3.4.1 Passing Job Execution Status to a PLSQL Procedure

The notification system passes job status change information to a PL/SQL procedure via the MGMT_NOTIFY_JOB object. An instance of this object is created for every status change. When a job changes status, the notification system calls the PL/SQL procedure associated with the notification rule and passes the populated object to the procedure. The procedure is then able to access the fields of the MGMT_NOTIFY_JOB object that has been passed to it.

Table 3–10 lists all corrective action status change attributes that can be passed:

Table 3–10 Job Status Attributes

Attribute	Datatype	Additional Information
job_name	VARCHAR2(128)	The job name.
job_owner	VARCHAR2(256)	The owner of the job.
job_type	VARCHAR2(32)	The type of the job.
job_status	NUMBER	The new status of the job.
state_change_guid	RAW(16)	The GUID of the state change record.
job_guid	RAW(16)	The unique id of the job.
execution_id	RAW(16)	The unique id of the execution.
targets	SMP_EMD_ NVPAIR_ARRAY	An array of the target name/target type pairs that the job runs on.
rule_owner	VARCHAR2(64)	The name of the notification rule that cause the notification to be sent.
rule_name	VARCHAR2(132)	The owner of the notification rule that cause the notification to be sent.
occurred_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 MGMT_NOTIFY_JOB 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 MGMT_NOTIFY_JOB object.

Table 3–11 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

Table 3–11 (Cont.) Job Status Codes

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

Example 3–15 PL/SQL Procedure Using a Status Code (Job)

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

CREATE OR REPLACE PROCEDURE LOG_PROBLEM_JOBS(status_change IN MGMT_NOTIFY_JOB)
IS
BEGIN
-- Log all failed jobs
    IF status_change.job_status = MGMT_JOBS.FAILED_STATUS
    THEN
        BEGIN
            INSERT INTO job_log (jobid, occured)
            VALUES (status_change.job_guid, SYSDATE);
        EXCEPTION
        WHEN OTHERS
        THEN
            -- If there are any problems then get the notification retried
            RAISE_APPLICATION_ERROR(-20000, 'Please retry');
        END;
        COMMIT;
    END IF;
END LOG_PROBLEM_JOBS;
```

3.4.2 Passing Job Execution Status to an OS Command or Script

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

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

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

Table 3–12 Environment Variables

Environment Variable	Description
JOB_NAME	The name of the job.
JOB_OWNER	The owner of the job.

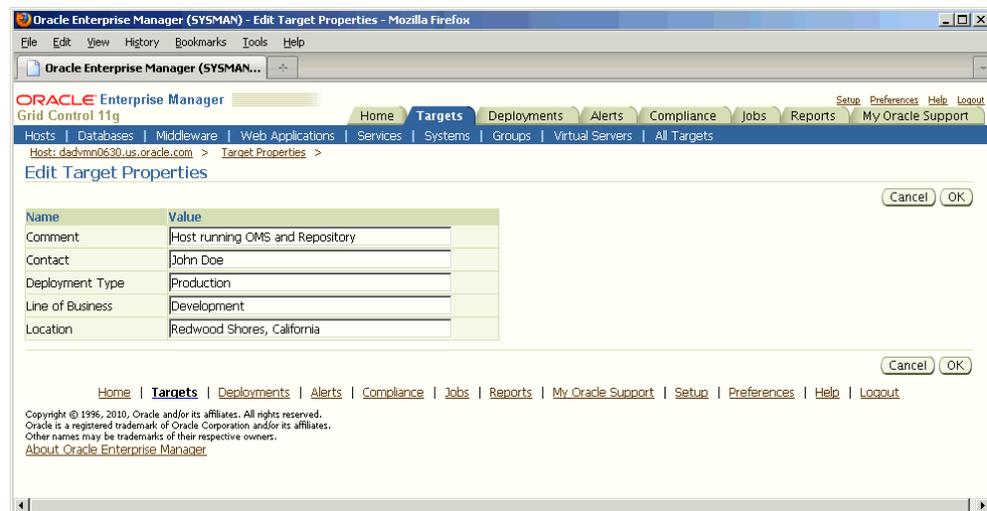
Table 3–12 (Cont.) Environment Variables

Environment Variable	Description
JOB_TYPE	The type of job.
JOB_STATUS	The job status.
TIMESTAMP	Time when the severity occurred.
NUM_TARGETS	The number of targets.
TARGET_NAME _n	The name of the <i>n</i> th target. For example, TARGET_NAME1, TARGET_NAME2.
TARGET_TYPE _n	The type of the <i>n</i> th target. For example TARGET_TYPE1, TARGET_TYPE2.
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.

3.5 Passing User-Defined Target Properties to Notification Methods

Enterprise Manager allows you to define target properties (accessed via Related Links on 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 3–4 Host Target Properties

3.6 Assigning Methods to Rules

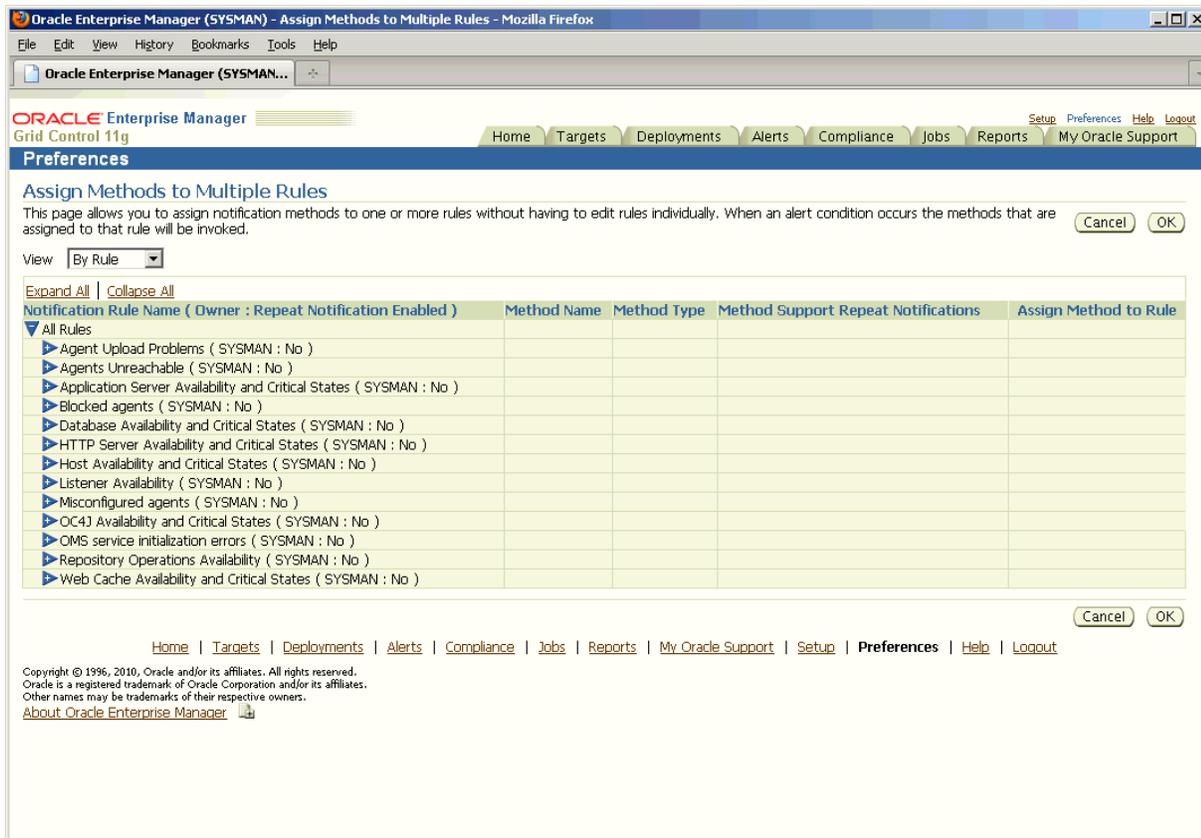
For each notification rule, you can assign one or more notification methods to be called when any of the criteria in the notification rule is met.

1. From the Enterprise Manager Grid Control, click **Preferences** at the top of the page.
2. Click **Notification Rules** in the vertical navigation bar.

The Enterprise Manager Grid Control displays the Notification Rules page. Any notification rules already created are listed in the **Notification Rules** table.

3. Click **Assign Methods to Multiple Rules**.
4. Perform your assignments.

Figure 3–5 Assigning Methods to Rules



3.7 Assigning Rules to Methods

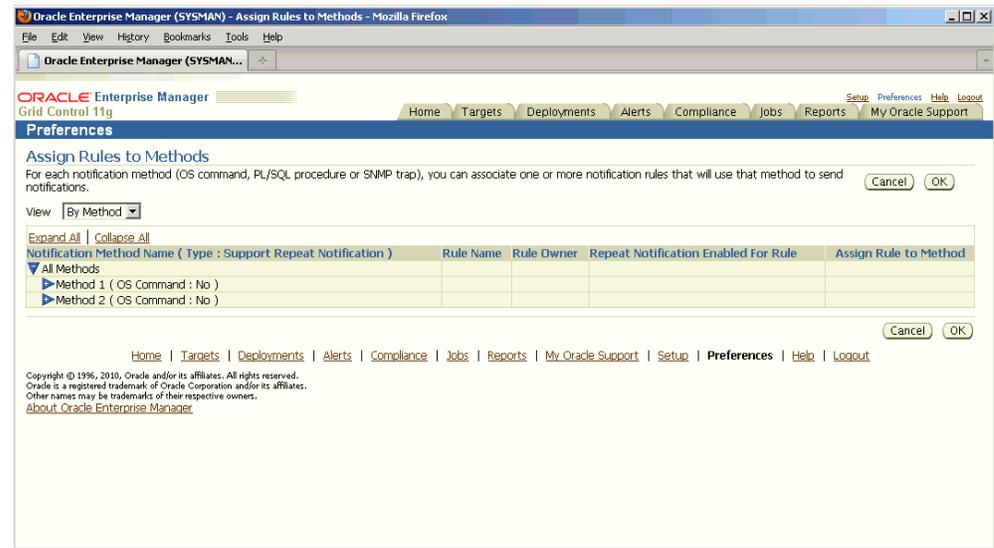
For each notification method, you can associate one or more notification rules that will use that method to send notifications.

1. From the Enterprise Manager Grid Control, click **Preferences** at the top of the page.
2. Click **Notification Rules** in the vertical navigation bar.

The Enterprise Manager Grid Control displays the Notification Rules page. Any notification rules already created are listed in the **Notification Rules** table.

3. Click **Assign Methods to Multiple Rules**.
4. From the **View** menu, select **By Method**.
5. Perform your assignments.

Figure 3–6 Assign Rules to Methods



3.8 Notification Coverage

To reduce the likelihood of an alert triggering and no administrator being notified because there was no notification rule covering that condition, you can use the Information Publisher (Enterprise Manager Report system) to view, for each target, the metrics monitored for that target and associated notification rules. Information Publisher provides an out-of-box report specifically designed for this purpose. You can run this report from the Report Definitions page (Reports tab) under *Monitoring-->Alerts and Policy Violations--> Notification Rule Coverage for Metric Alerts and Availabilities (Target)*

3.9 Management Information Base (MIB)

Enterprise Manager Grid 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.

3.9.1 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.

3.9.2 Reading the MIB Variable Descriptions

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

3.9.2.1 Variable Name

Syntax

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

Max-Access

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

Status

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

Explanation

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

Typical Range

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

Significance

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

Related Variables

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

Suggested Presentation

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

3.9.2.2 MIB Definition

[Example 3-16](#) shows a typical MIB definition used by Enterprise Manager.

Example 3-16 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
    }
oraEM4AlertIndex OBJECT-TYPE
    SYNTAX  INTEGER
    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 }
oraEM4Traps OBJECT IDENTIFIER ::= { oraEM4 2 }
oraEM4Alert TRAP-TYPE
ENTERPRISE oraEM4Traps
VARIABLES { oraEM4AlertTargetName, oraEM4AlertTargetType,
             oraEM4AlertHostName, oraEM4AlertMetricName,
             oraEM4AlertKeyName, oraEM4AlertKeyValue, oraEM4AlertTimeStamp,
             oraEM4AlertSeverity, oraEM4AlertMessage,
             oraEM4AlertRuleName, oraEM4AlertRuleOwner,
             oraEM4AlertMetricValue, oraEM4AlertContext }
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 OraEM4AlertEntry
    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
    }
oraEM4JobAlertIndex OBJECT-TYPE
    SYNTAX INTEGER
    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 }
oraEM4JobAlert TRAP-TYPE
    ENTERPRISE  oraEM4Traps
    VARIABLES   { oraEM4JobAlertJobName, oraEM4JobAlertJobOwner,
                  oraEM4JobAlertJobType, oraEM4JobAlertJobStatus,
                  oraEM4JobAlertTargets, oraEM4JobAlertTimeStamp,
                  oraEM4JobAlertRuleName, oraEM4JobAlertRuleOwner,
                  oraEM4JobAlertMetricName, oraEM4JobAlertMetricValue,
                  oraEM4JobAlertContext, oraEM4JobAlertKeyName,
                  oraEM4JobAlertKeyValue, oraEM4JobAlertSeverity }
DESCRIPTION
    "The variables included in the oraEM4JobAlert trap."
    ::= 2
END

```

3.10 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.

3.10.1 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 setup under General page of Preferences. 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 under the Schedule page of Preferences. No e-mails will be sent unless a Notification Schedule has been defined.
- Make sure a Notification Rule is defined to match the target, metric, severity and availability states you are interested and make sure e-mail and notification methods are assigned to the rule. A summary of the notification rule can be checked by going to the Rules page under Setup and clicking the rule name.

3.10.2 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. These errors can be seen in the Errors page under Management Services and Repository under Setup.
- Check for any delivery errors. From the Alerts section of a target home page, click on the alert message to access the metric details page. In the Alert History section, click on the Details icon for more information about the alert. The details will give the reason why the notification was not delivered. Delivery errors are stored in MGMT_NOTIFICATION_LOG with the DELIVERED column set to 'N'.
- Severities will not be displayed in the Grid Control console if no metric values have been loaded for the metric associated with the severity.

3.10.3 Notification System Trace Messages

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

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

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

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

```
grep em.notification emoms.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.

```
2006-11-08 03:18:45,385 [Orion Launcher] INFO em.notification init.1279 - Short
format maximum length is 155
```

```
2006-11-08 03:18:45,386 [Orion Launcher] INFO em.notification init.1297 - Short
format is set to both subject and body
```

```
2006-11-08 03:18:45,387 [NotificationMgrThread] INFO em.notification run.1010 -
Waiting for connection to EM Repository...
```

```
2006-11-08 03:18:46,006 [NotificationMgrThread] INFO em.notification run.1041 -
Registering for Administrative Queue Name...
```

```
2006-11-08 03:18:46,250 [NotificationMgrThread] INFO em.notification run.1078 -
Administrative Queue is ADM21
```

```
2006-11-08 03:18:46,250 [NotificationMgrThread] INFO em.notification run.1089 -
Creating thread pool: min = 6 max = 24
```

```
2006-11-08 03:18:48,206 [NotificationMgrThread] INFO em.notification
handleAdminNotification.655 - Handling notifications for EMAIL1
```

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

```
2006-11-08 07:26:30,242 [NotificationMgrThread] ERROR em.notification
getConnection.237 - Failed to get a connection Io exception: The Network Adapter
could not establish the connection
```

3.10.4 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.us.oracle.com: Sending failed; nested exception is:
javax.mail.MessagingException: Unknown SMTP host: badhost.us.oracle.com;

Invalid e-mail address:

Failed to connect to gateway: rgmemeasntp.oraclecorp.com: Sending failed; nested exception is:
javax.mail.MessagingException: 550 5.7.1 <smpemailtest_ie@oracle.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

3.10.5 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 (stacb10.us.oracle.com_Management_Service) (SYSMAN, myrule)

No execute permission on executable:

Error calling /bin/myscript: java.io.IOException: /bin/myscript: cannot execute (stacb10.us.oracle.com_Management_Service) (SYSMAN, myrule)

Timeout because OS Command ran too long:

Timeout occurred running /bin/myscript (stacb10.us.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.

3.10.6 SNMP Trap Errors

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

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

3.10.7 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 p(s IN MGMT_NOTIFY_SEVERITY)

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 Grid Control console.

Always use the Test button to make sure the PL/SQL configuration is valid.

User-Defined Metrics

User-defined metrics allow you to extend the reach of Enterprise Manager's monitoring to conditions specific to particular environments via custom scripts or SQL queries and function calls. Once defined, user-defined metrics will be monitored, aggregated in the repository and trigger alerts like regular metrics.

This chapter covers the following topics:

- [Extending Monitoring Capability](#)
- [Creating OS-Based User-Defined Metrics](#)
- [Creating a SQL-Based User-Defined Metric](#)
- [Notifications, Corrective Actions, and Monitoring Templates](#)
- [Changing User-Defined Metric Credentials](#)

4.1 Extending Monitoring Capability

There are two types of user-defined metrics:

- **OS-Based User-Defined Metrics:** Accessed from Host target home pages, these user-defined metrics allow you to define new metrics using custom Operating System (OS) scripts.

To monitor for a particular condition (for example, check successful completion of monthly system maintenance routines), you can write a custom script that will monitor that condition, then create an OS-based user-defined metric that will use your custom script. Each time the metric is evaluated by Enterprise Manager, it will use the specified script, relying on that script to return the value of the condition.

- **SQL-Based User-Defined Metrics:** Accessed from Database target home pages, these user-defined metrics allow you to implement custom database monitoring using SQL queries or function calls.

SQL-based user-defined metrics do not use external scripts. You enter SQL directly into the Enterprise Manager user interface at the time of metric creation

Once a user-defined metric is created, all other monitoring features, such as alerts, notifications, historical collections, and corrective actions are automatically available to it.

Administrators who already have their own library of custom monitoring scripts can leverage these monitoring features by integrating their scripts with Enterprise Manager via user-defined metrics. Likewise, existing SQL queries or function calls

currently used to monitor database conditions can be easily integrated into Enterprise Manager's monitoring framework using the SQL-based user-defined metric.

4.2 Creating OS-Based User-Defined Metrics

Creating an OS-based user-defined metric involves two steps:

- Step 1: [Create Your OS Monitoring Script](#)
- Step 2: [Register the Script as a User-Defined Metric](#)

4.2.1 Create Your OS Monitoring Script

Using a scripting language of your choice, create a script that contains logic to check for the condition being monitored. For example, scripts that check for disk space or memory usage. All scripts to be run with user-defined metrics should be placed in a directory to which the Management Agent has full access privileges. Scripts themselves must have the requisite permissions set so that they can be executed by the Management Agent. The script runtime environment must also be configured: If your script requires an interpreter, such as a Perl interpreter, this must be installed on that host as well.

All monitoring scripts should contain code to perform the following basic functions:

- [Code to check the status of monitored objects](#)
- [Code to return script results to Enterprise Manager](#)

4.2.1.1 Code to check the status of monitored objects

Define logic in the code that checks the condition being monitored such as determining the amount of free space on a particular file system or level of memory usage.

After checking the monitored condition, the script should return the value associated with the monitored object.

When you choose to have the script return a specific value from the monitored object (for example, current disk space usage), you can also have Enterprise Manager evaluate the object's current value against specific warning and critical thresholds. You specify these warning and critical thresholds from the Grid Control console at the time you create the user-defined metric. Based on the evaluation of the metric's value against the thresholds, an alert may be triggered at one of the following severity levels:

Table 4–1 Metric Severity Levels

Severity Level	Status
Script Failure	The script failed to run properly.
Clear	No problems with the object monitored; status is clear. If thresholds were specified for the metric, then it means the thresholds were not reached.
Warning	The value of the monitored object reached the warning threshold.
Critical	The value of the monitored object reached the critical threshold.

4.2.1.2 Code to return script results to Enterprise Manager

After checking the monitored condition, the script should return the value associated with the monitored object. The script returns values back to Enterprise Manager by sending formatted information to standard output (stdout) using the syntax that is

consistent with the scripting language (the "print" statement in Perl, for example). Enterprise Manager then checks the standard output of a script for this formatted information; specifically it checks for the tags: `em_result` and `em_message` and the values assigned to these tags.

The script must assign the value of the monitored object to the tag `em_result`. The output must be written as a string delimited by new line characters. For example, if the value of the monitored object is 200, your script can return this to Enterprise Manager as shown in this Perl statement:

```
print "em_result=200\n"
```

You can also have Enterprise Manager evaluate the returned value against specified warning and critical thresholds. You specify these warning and critical thresholds when you register your script as a user-defined metric in the console.

If the comparison between the warning or critical threshold holds true, a warning or critical alert will be generated. The default message for this alert will be:

```
"The value is $em_result".
```

You can choose to override this default message with a custom message by assigning the string to be used to the tag `em_message`.

For example, if you want your alert message to say "Disk usage is high", your script can return this custom message as follows:

```
print "em_message=Disk usage is high\n"
```

Important: Script output tags **must be lower-case** in order for Enterprise Manager to recognize the script output as valid user-defined metric feedback. Messages or values associated with each tag can be mixed case.

- Valid tag output: `em_result=My Value\n`
 - Invalid tag output: `Em_Result=My Value\n`
-
-

For a successful script execution, the script output must start with the "em_result=" string in a new line. The message must start with the "em_message=" string in a new line.

The following table summarizes the script output tags.

Table 4–2 *Script Output Information Tags*

Tag	Definition
<code>em_result</code>	Use this tag to return script result values. Exactly one <code>em_result</code> tag must be found in STDOUT. If more than one <code>em_result</code> tag is found, the first tag encountered will be used; subsequent <code>em_result</code> tags will be ignored. Example: <pre>print "em_result=200\n"</pre> Returns 200 as the value of the monitored object.

Table 4–2 (Cont.) Script Output Information Tags

Tag	Definition
em_message	<p>Use this tag to specify a message with the script result value in STDOUT. For OS-based user-defined metrics, only one em_message tag is permitted. If you submit more than one em_message tag, only the first tag is used. Subsequent tags are ignored.</p> <p>Example:</p> <pre>print "em_result=200\nem_message=Disk usage is high\n"</pre> <p>Returns 200 as the value of the monitored object in addition to the message "Disk usage is high".</p> <p>If you want to include the value of em_result in the message, you can use the placeholder \$em_result.</p> <p>Example:</p> <pre>print "em_message=Disk usage is at \$em_result.\n"</pre> <p>If script execution is successful AND it does not contain a em_message string, a default em_message string is automatically generated. The following message format is used:</p> <pre>em_message=The value is \$em_result</pre> <p>Example:</p> <pre>print "em_result=200\n"</pre> <p>Returns 200 as the value of the monitored object and the generated message "The value is 200"</p>

The output of the user-defined monitoring script must be either em_result or em_message. In the event of system error, such as Perl aborting and writing information to STDERR pertaining to invalid commands, the script returns:

- Non-zero value
- STDOUT and STDERR messages are concatenated and sent to STDERR

This error situation results in a metric error for this user-defined metric. You can view metric errors in the Errors page of the Alerts tab in the Enterprise Manager console.

OS Script Location

Oracle recommends that user-defined metric OS scripts reside in a location outside the Agent Oracle Home. Doing so isolates scripts from any changes that may occur as a result of an Agent upgrade and ensures your scripts remain operational. When registering your script in the Grid Control console, you must specify the full path to the script. Do not use Available Properties (for example, %scriptsDir% or %emdRoot%) as part of the path specification.

4.2.1.3 Script Runtime Environment

When the user-defined metric is evaluated, it executes the script using the credentials (user name and password) specified at the time the user-defined metric was registered in the Enterprise Manager console. See ["Register the Script as a User-Defined Metric"](#) on page 4-5. Ensure that the user name and password you specify for the user-defined metric is an active account (on that machine) possessing the requisite permissions to run the script.

4.2.2 Register the Script as a User-Defined Metric

Once you have created the monitoring script, you are ready to add this monitoring functionality to Enterprise Manager as a user-defined metric.

Important: For OS-based user-defined metrics, make sure the Management Agent is up and running on the machine where the monitoring script resides before creating the user-defined metric. Operator privilege or higher is required on the host target.

Creating an OS-Based User-Defined Metric

1. From the home page of the Host that has your OS monitoring script (Related Links), choose User-Defined Metrics. The User-Defined Metrics summary page appears containing a list of previously defined User-Defined Metrics. From this page, you perform edit, view, delete, or create like functions on existing User-Defined Metrics.
2. Click Create. The Create User-Defined Metric page appears.
3. Enter the requisite metric definition, threshold, and scheduling information. For the Command Line field, enter the full path to your script, including any requisite shell or interpreters. For example, `/bin/sh myscript`. See the following section for more details.
4. Click OK. The User-Defined Metric summary page appears with the new User-Defined Metric appended to the list.

If the user-defined metric has been created and the severity has not been updated recently, it is possible that there are metric errors associated with the user-defined metric execution. In this situation, access the Errors subtab under Alerts tab to check.

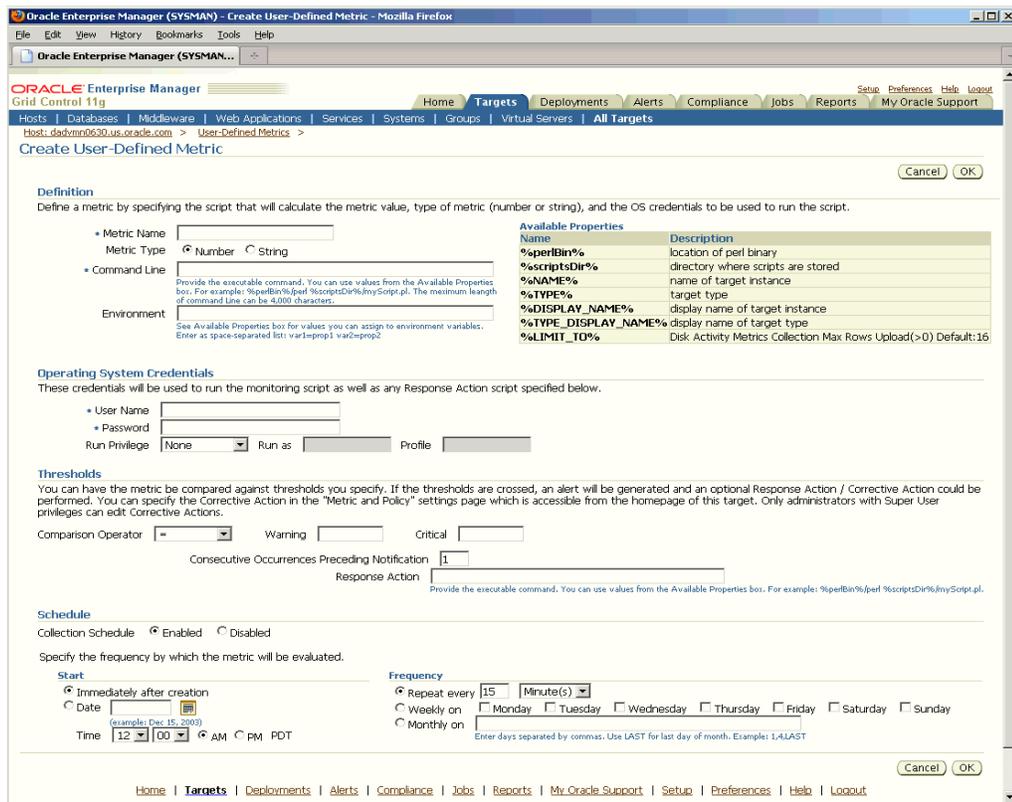
Create User-Defined Metric Page (OS-based User-Defined Metric)

The Create User-Defined Metric page allows you to specify the metric information required to successfully register the metric with Enterprise Manager. The page is divided into functional categories:

- **Definition:** You define the operational and environmental parameters required to run your script, in addition to the name of the script itself.
- **Operating System Credentials:** You enter the credentials used to run the monitoring script. See Enterprise Manager online help for more details on Response Actions. This functional area appears when creating OS-based user-defined metrics.
- **Thresholds:** To have the value returned by your script compared with set threshold values, enter the requisite threshold information. The value of the monitored metric returned by your script (as specified by `em_result`) will be compared against the thresholds you specify. If the comparison holds true, a warning or critical alert may be generated.
- **Schedule:** Specify the start time and frequency at which the user-defined script should be run. The time zone used is that of the Agent running on the monitored host.

The following figures show the Create User-Defined Metric pages for an OS-based user-defined metric. When accessing this page from any Host home page, the Create User-Defined Metric page appears as shown in [Figure 4-1](#).

Figure 4–1 Create User-Defined Metric Page (OS-Based)



Key elements of this page are described in the following tables.

Table 4–3 Create User-Defined Metric Page: Definition

User-Interface Element	Description
Metric Name	Metric name identifying the user-defined metric in the Enterprise Manager user interface. This name must be unique for all User-Defined Metrics created on that host.
Metric Type	Type of the value returned by the user-defined script. Choose "NUMBER" if your script returns a number. Choose "STRING" if your script returns an alphanumeric text string.
Command Line	Enter the complete command line entry required to execute the user-defined script. You must enter the full command path as well as full path to the script location. For example, to run a Perl script, you might enter something like the following in the Command Line entry field: <pre>/u1/bin/perl /u1/scripts/myScript.pl</pre> The content of the Command Line is passed as a literal string, so you may use any syntax, special characters, or parameters allowed by your operating system.

Table 4–3 (Cont.) Create User-Defined Metric Page: Definition

User-Interface Element	Description
Environment	<p>Optional. Enter any environmental variable(s) required to run the user-defined script. A list of predefined properties that can be passed to your script as variables is listed in the Available Properties box. You may also specify your own environment variables. Multiple variables can be defined as a space-separated list.</p> <p>Example: If your script uses three variables (var1, var2, var3) where var1 is the location of the Perl directory (predefined), var2 is the directory where your Perl scripts are stored (predefined), and var3 is an Oracle home, your entry in the Environment text entry field would appear as follows:</p> <pre>var1=%perlBin% var2=%scriptsDir% var3=/u1/orahome10</pre>

Table 4–4 Create User-Defined Metric Page: Operating System

User-Interface Element	Description
User Name	Enter the user name for a valid operating system account on the machine where the script is to be run. Make sure the specified account has the requisite privileges to access the script directory and execute the script.
Password	Enter the password associated with the User Name.

Table 4–5 Create User-Defined Metric Page: Threshold

User-Interface Element	Description																																				
Comparison Operator	<p>Select the comparison method Enterprise Manager should use to compare the value returned by the user-defined script to the threshold values.</p> <p>Available Comparison Operators</p> <table border="1"> <thead> <tr> <th>Operator</th> <th>Value</th> <th>Metric Type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>=</td> <td></td> <td>Number</td> <td>equal to</td> </tr> <tr> <td>></td> <td></td> <td>Number</td> <td>greater than</td> </tr> <tr> <td><</td> <td></td> <td>Number</td> <td>less than</td> </tr> <tr> <td>>=</td> <td></td> <td>Number</td> <td>greater than or equal to</td> </tr> <tr> <td><=</td> <td></td> <td>Number</td> <td>less than or equal to</td> </tr> <tr> <td>!=</td> <td></td> <td>Number</td> <td>not equal to</td> </tr> <tr> <td>CONTAINS</td> <td></td> <td>String</td> <td>contains at least</td> </tr> <tr> <td>MATCH</td> <td></td> <td>String</td> <td>exact match</td> </tr> </tbody> </table>	Operator	Value	Metric Type	Description	=		Number	equal to	>		Number	greater than	<		Number	less than	>=		Number	greater than or equal to	<=		Number	less than or equal to	!=		Number	not equal to	CONTAINS		String	contains at least	MATCH		String	exact match
Operator	Value	Metric Type	Description																																		
=		Number	equal to																																		
>		Number	greater than																																		
<		Number	less than																																		
>=		Number	greater than or equal to																																		
<=		Number	less than or equal to																																		
!=		Number	not equal to																																		
CONTAINS		String	contains at least																																		
MATCH		String	exact match																																		
Warning	<p>The value returned by the script is compared to the Warning threshold value using the specified comparison operator. If this comparison holds true, an alert triggers at the warning severity level.</p> <p>Specifically, an alert triggers at the warning severity level if the following comparison is true:</p> <pre><script_value> <comparison_operator> <warning_threshold></pre> <p>and if the consecutive occurrences preceding notification has been reached.</p>																																				

Table 4–5 (Cont.) Create User-Defined Metric Page: Threshold

User-Interface Element	Description
Critical	<p>The value returned by the script is compared to the Critical threshold value using the specified comparison operator. If this comparison holds true, an alert triggers at the critical severity level.</p> <p>Specifically, an alert triggers at the critical severity level if the following comparison is true:</p> <pre><script_value> <comparison_operator> <critical_threshold></pre> <p>and if the consecutive occurrences preceding notification has been reached.</p>
Consecutive Occurrences Preceding Notification	<p>Consecutive number of times a returned value reaches either the warning or critical thresholds before an alert triggers at a warning or critical severity. This feature is useful when monitoring for sustained conditions. For example, if your script monitors the CPU load for a particular machine, you do not want to be notified at every CPU load spike. Instead, you are only concerned if the CPU load remains at a particular threshold (warning or critical) level for a consecutive number of monitoring cycles.</p>
Response Action	<p>Optional. Specify a script or command that will be executed if the user-defined metric generates a warning or critical alert. The script or command will be executed using the credentials of the Agent owner. Important: Only an Enterprise Manager Super Administrator can create/edit response actions for metrics.</p> <p>For example, the Management Agent executes the response action if:</p> <pre>The Alert severity is Warning or Critical AND There is a change in severity (for example, warning -> critical, critical --> warning, clear --> warning or critical)</pre> <p>For more information, see Enterprise Manager online help.</p>

The User-Defined Metric Schedule interface lets you specify when the Management Agent should start monitoring and the frequency at which it should monitor the condition using your OS script.

4.2.3 OS-Based User-Defined Metric Example

The sample Perl script used in this example monitors the 5-minute load average on the system. The script performs this function by using the 'uptime' command to obtain the average number of jobs in the run queue over the last 5 minutes.

The script is written in Perl and assumes you have Perl interpreter located in /usr/local/bin on the monitored target.

This script, called `udmload.pl`, is installed in a common administrative script directory defined by the user. For example, `/u1/scripts`.

Important: Do not store user-defined metric monitoring scripts in the same location as Enterprise Manager system scripts.

Full text of the script:

```
#!/usr/local/bin/perl
```

```
# Description: 5-min load average.
# Sample User Defined Event monitoring script.

$ENV{PATH} = "/bin:/usr/bin:/usr/sbin";

$DATA = `uptime`;
$DATA =~ /average:\s+([\.\d]+),\s+([\.\d]+),\s+([\.\d]+)\s*$/;
```

```
if (defined $2) {
    print "em_result=$2\n";
} else {
    die "Error collecting data\n";
}
```

1. Copy the script (udmload.pl) to the monitored target. For example: /u1/scripts. Make sure you have an Enterprise Manager 10g Management Agent running on this machine.
2. Edit the script, if necessary, to point to the location of the Perl interpreter on the monitored target. By default, the script assumes the Perl interpreter is in /usr/local/bin.
3. As a test, run the script: udmload.pl You may need to set its file permissions so that it runs successfully. You should see output of this form:

```
em_result=2.1
```

4. In Create User-Defined Metric page, create a new user-defined metric as follows:

a. Definition Settings

- * **Metric Name:** Test User-Defined Metric
- * **Metric Type:** Number
- * **Command Line:** %perlBin%/perl /u1/scripts/udmload.pl
- * **Environment:** leave blank
- * **Operating System User Name:** <OS user able to execute the script>
- * **Password:** *****

b. Threshold Settings

- * **Comparison Operator:** >=
- * **Critical Threshold:** 0.005
- * **Warning Threshold:** 0.001
- * **Consecutive Occurrences Preceding Notification:** 1

In this example, we want the metric to trigger an alert at a Warning level if the 5-minute load average on the machine reaches 0.001, and trigger an alert at a Critical level if the 5-minute load average reaches 0.005. Feel free to change these thresholds depending on your system.

c. Schedule Settings:

- * **Start:** Immediately after creation
- * **Frequency:** Repeat every 5 minutes. You must specify at least a 5 minute interval.

Setting Up the Sample Script as a User-Defined Metric

When the 5-minute load reaches at least 0.001, you should see the metric trigger an alert.

4.3 Creating a SQL-Based User-Defined Metric

You can also define new metrics using custom SQL queries or function calls supported against single instance databases and instances on Real Application Clusters (RAC). To create this type of user-defined metric, you must have Enterprise Manager Operator privileges on the database:

1. From the Related Links area of any Database home page, choose User-Defined Metrics. The User-Defined Metrics summary page appears containing a list of previously defined user-defined metrics. From this page, you perform edit, view, delete, or create like functions on existing user-defined metrics.
2. Click Create. The Create User-Defined Metric page appears.
3. Enter the requisite metric definition, threshold, and scheduling information. For the SQL Query field, enter the query or function call. See the following section for more information.

Click Test to verify that the SQL query or function call can be executed successfully using the credentials you have specified

4. Click OK. The User-Defined Metric summary page appears with the new user-defined metric appended to the list.

If the user-defined metric has been created and the severity has not been updated recently, it is possible that there are metric errors associated with the user-defined metric execution. In this situation, access the Errors subtab under Alerts tab to check.

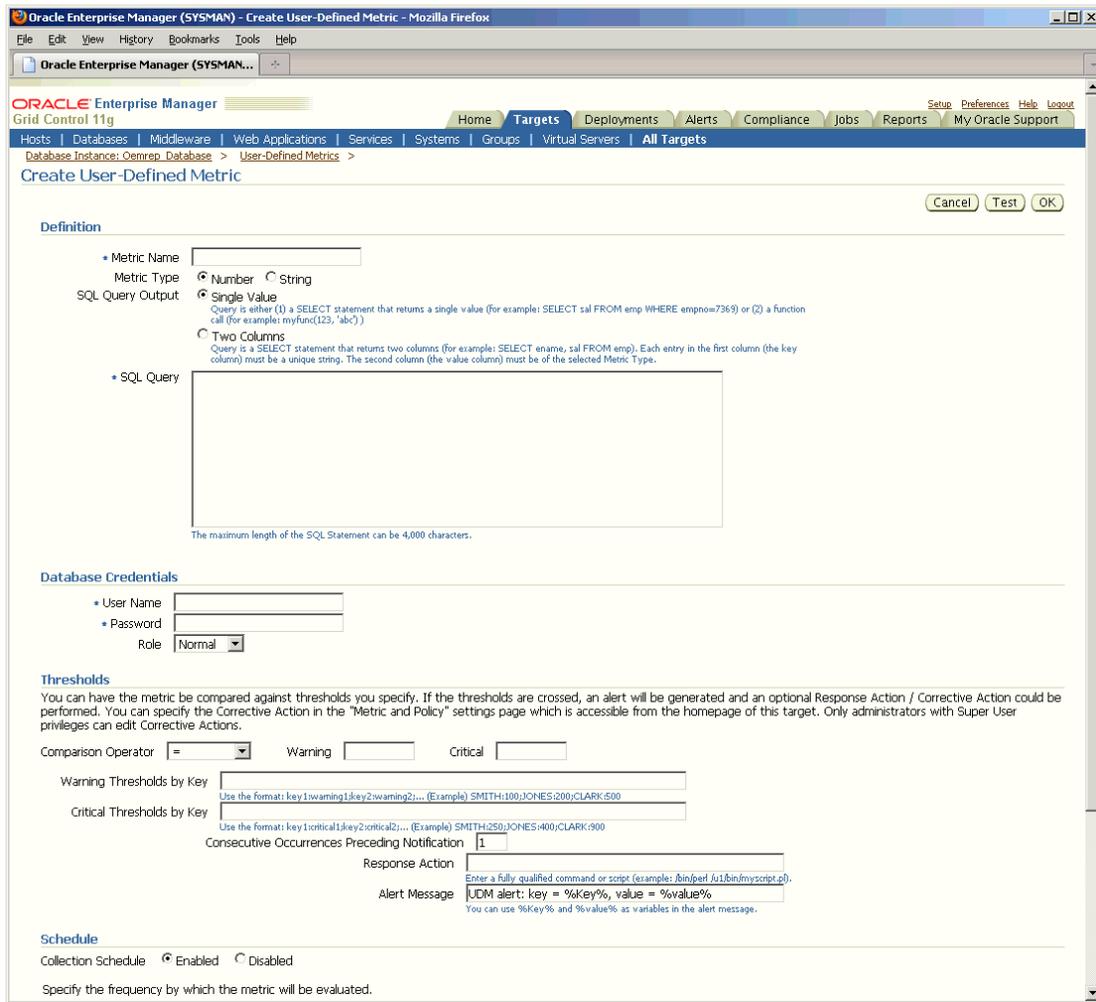
Create User-Defined Metric Page (SQL-Based User-Defined Metric)

The Create User-Defined Metric page allows you to specify the metric information required to successfully register the metric with Enterprise Manager. The page is divided into functional categories:

- **Definition:** You define the operational and environmental parameters required to run your script, in addition to the name of the script itself.
- **Database Credentials:** You enter the user name and password for a valid user account on the database where the SQL is to be run. Make sure the specified user account has the requisite administrative and access privileges to execute the SQL query or function call.
- **Thresholds:** To have the value returned by your SQL query or function call compared with set threshold values, enter the requisite threshold information. The value of the monitored metric returned by your query or function call will be compared against the thresholds you specify. If the comparison holds true, a warning or critical alert may be generated.
- **Schedule:** Specify the start time and frequency at which the user-defined SQL query or function call should be executed. The time zone used is that of the Agent running on the monitored machine.

The following figures show the Create User-Defined Metric pages for a SQL-based user-defined metric. When accessing this page from any Database home page, the Create User-Defined Metric page appears as shown in [Figure 4–2](#).

Figure 4–2 Create User-Defined Metric Page (SQL-Based)



Key elements of this page are described in the following tables.

Table 4–6 Create User-Defined Metric Page: Definition

User-Interface Element	Description
Metric Name	Metric name identifying the user-defined metric in the Enterprise Manager user interface.
Metric Type	Type of the value returned by the user-defined script. Choose "NUMBER" if your script returns a number. Choose "STRING" if your script returns an alphanumeric text string.

Table 4–6 (Cont.) Create User-Defined Metric Page: Definition

User-Interface Element	Description
SQL Query Output	<p>Specify whether the SQL script is to return a single value (one column) or a multiple rows (two columns).</p> <ul style="list-style-type: none"> Single Value: Query is one of the following types. <p>A <i>SELECT statement</i> returning a single value. Example: <code>SELECT sal FROM emp WHERE empno=7369</code></p> <p>A <i>function call</i> returning a single value. Example: <code>myfunc(123, 'abc')</code></p> Two Columns: Query is a <code>SELECT</code> statement that returns two columns and possibly multiple rows. Example: <code>SELECT ename, sal FROM emp</code>. Each entry in the first column (the key column) must be a unique string. The second column (the value column) must be of the selected Metric Type.
SQL Query	<p>Enter a SQL query or function call that returns values of the appropriate type (<code>STRING</code> or <code>NUMBER</code>). The SQL statement must return one or two column. If your SQL statement only returns one column, only one row can be returned. If you want multiple rows returned, your SQL statement must return two columns.</p>

Table 4–7 Create User-Defined Metric Page: Database Credentials

User-Interface Element	Description
User Name	Enter the user name for a valid database account on the database where the SQL query is to be run. Make sure that the specified account has the requisite privileges to run the SQL query.
Password	Enter the password associated with the User Name.

Table 4–8 Create User-Defined Metric Page: Threshold

User-Interface Element	Description																																				
Comparison Operator	<p>Select the comparison method Enterprise Manager should use to compare the value returned by the SQL query or function call to the threshold values. When the query returns two columns, the second column (value column) will be used for comparison against threshold values.</p> <p>Available Comparison Operators</p> <table border="1"> <thead> <tr> <th>Operator</th> <th>Value</th> <th>Metric Type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>=</td> <td></td> <td>Number</td> <td>equal to</td> </tr> <tr> <td>></td> <td></td> <td>Number</td> <td>greater than</td> </tr> <tr> <td><</td> <td></td> <td>Number</td> <td>less than</td> </tr> <tr> <td>>=</td> <td></td> <td>Number</td> <td>greater than or equal to</td> </tr> <tr> <td><=</td> <td></td> <td>Number</td> <td>less than or equal to</td> </tr> <tr> <td>!=</td> <td></td> <td>Number</td> <td>not equal to</td> </tr> <tr> <td>CONTAINS</td> <td></td> <td>String</td> <td>contains at least</td> </tr> <tr> <td>MATCH</td> <td></td> <td>String</td> <td>exact match</td> </tr> </tbody> </table>	Operator	Value	Metric Type	Description	=		Number	equal to	>		Number	greater than	<		Number	less than	>=		Number	greater than or equal to	<=		Number	less than or equal to	!=		Number	not equal to	CONTAINS		String	contains at least	MATCH		String	exact match
Operator	Value	Metric Type	Description																																		
=		Number	equal to																																		
>		Number	greater than																																		
<		Number	less than																																		
>=		Number	greater than or equal to																																		
<=		Number	less than or equal to																																		
!=		Number	not equal to																																		
CONTAINS		String	contains at least																																		
MATCH		String	exact match																																		

Table 4–8 (Cont.) Create User-Defined Metric Page: Threshold

User-Interface Element	Description
Warning	<p>The value returned by the SQL query or function call is compared to the Warning threshold value using the specified comparison operator. If this comparison holds true, an alert triggers at the warning severity level.</p> <p>Specifically, an alert triggers at the warning severity level if the following comparison is true:</p> <pre><query_value> <comparison_operator> <warning_threshold></pre> <p>and if the consecutive occurrences preceding notification has been reached.</p>
Critical	<p>The value returned by the SQL query or function call is compared to the Critical threshold value using the specified comparison operator. If this comparison holds true, an alert triggers at the critical severity level.</p> <p>Specifically, an alert triggers at the critical severity level if the following comparison is true:</p> <pre><query_value> <comparison_operator> <critical_threshold></pre> <p>and if the consecutive occurrences preceding notification has been reached.</p>
Warning Thresholds by Key and Critical Thresholds by Key	<p>For queries returning two columns (the first column is the key and the second column is the value), you can specify thresholds on a per key basis. The following example uses the following query:</p> <pre>SELECT ename FROM emp</pre> <p>Threshold settings for this example are shown.</p> <p>Use the format <i>key:value</i> . Keys are case-sensitive.</p> <ul style="list-style-type: none"> ■ Warning:500 ■ Critical:300 ■ Comparison Operator: < ■ Warning threshold by key: SMITH:250;JONES:400;CLARK:900 <p>The warning threshold is set to 250 for SMITH, 400 for JONES, and 900 for CLARK.</p> <ul style="list-style-type: none"> ■ Critical threshold by key: SMITH:100;JONES:200;CLARK:500 <p>The critical threshold is set to 100 for SMITH, 200 for JONES, and 500 for CLARK.</p> <p>All other keys will use the threshold values specified in the Warning and Critical fields.</p>
Consecutive Occurrences Preceding Notification	<p>Consecutive number of times a returned value reaches either the warning or critical thresholds before an alert triggers at a warning or critical severity. This feature is useful when monitoring for sustained conditions. For example, if your script monitors the CPU load for a particular machine, you do not want to be notified at every CPU load spike. Instead, you are only concerned if the CPU load remains at a particular threshold (warning or critical) level for a consecutive number of monitoring cycles.</p>

Table 4–8 (Cont.) Create User-Defined Metric Page: Threshold

User-Interface Element	Description
Response Action	<p>Optional. Specify a script or command that will be executed if the user-defined metric generates a warning or critical alert. The script or command will be executed using the credentials of the Agent owner. Important: Only an Enterprise Manager Super Administrator can create/edit response actions for metrics.</p> <p>For example, the Management Agent executes the response action if:</p> <p>The Alert severity is Warning or Critical</p> <p>AND</p> <p>There is a change in severity (for example, warning -> critical, critical --> warning, clear --> warning or critical)</p> <p>For more information, see Enterprise Manager online help.</p>
Alert Message	<p>Enter a custom message (up to 400 characters) to be used when an alert is sent. The default message uses %Key% and %value% variables to display the metric key and its returned value. The %Key% and %value% variables are case-sensitive.</p> <p>For example, a payroll system alert for underpayment of salary might be defined as:</p> <p>Underpaid Employee: %Key% has salary of %value%</p> <p>If the SQL query returns 2 columns, you can use the %Key% variable to represent the key value and the %value% variable to represent the return value.</p> <p>If the SQL query returns 1 column, only the %value% variable is applicable in the alert message.</p>

The User-Defined Metric Schedule interface lets you specify the frequency at which the SQL query or function should be run.

4.3.1 SQL-Based User-Defined Metric Examples

For a database version 9i and higher, you can run the example queries as db snmp, which is the default monitoring user account for the Management Agent. On a 8.1.7 database (which does not have SELECT ANY DICTIONARY system privilege), you must grant db snmp the following privileges in order for the queries to run successfully:

For example #1:

```
grant select on sys.dba_tablespaces to db snmp;
grant select on sys.dba_data_files to db snmp;
grant select on sys.dba_free_space to db snmp;
```

For example #2:

```
grant select on sys.dba_extents to db snmp;
```

The above grant statements can be run as SYSDBA after logging in via "connect internal". The queries can also be run by any user who has been granted the DBA role.

4.3.1.1 Example 1: Query Returning Tablespace Name and Percent Used

This sample user-defined metric monitors the percentage of space used for dictionary managed permanent tablespaces. A DBA can use this as a reference on when to add datafiles for the tablespace.

Oracle recommends setting a polling frequency of 30 minutes, warning threshold at 75, and critical threshold at 85.

Example 1 SQL

```
SELECT d.tablespace_name,
       round(((a.bytes - NVL(f.bytes,0))*100/a.maxbytes),2) used_pct
FROM   sys.dba_tablespaces d,
       (select tablespace_name, sum(bytes) bytes, sum(greatest(maxbytes,bytes))
        maxbytes
        from sys.dba_data_files group by tablespace_name) a,
       (select tablespace_name, sum(bytes) bytes
        from sys.dba_free_space group by tablespace_name) f
WHERE  d.tablespace_name = a.tablespace_name(+)
AND    d.tablespace_name = f.tablespace_name(+)
AND    NOT (d.extent_management = 'LOCAL' AND d.contents = 'TEMPORARY')
```

4.3.1.2 Example 2: Query Returning Segment Name/Type and Extent Count

This sample user-defined metric checks for non-system table and index segments that are reaching a high number of extents. A high number of extents could indicate a segment with fragmentation and/or performance problems. A DBA can use this as a reference on when to call Segment Shrink or the Reorganization Wizard in Enterprise Manager.

Oracle recommends setting a polling frequency of 24 hours, warning threshold at 1000, and critical threshold at 2000.

Example 2 SQL

```
SELECT decode(nvl(partition_name, ' '),
              ' ', owner || '.' || segment_name || ' ' || segment_type,
              owner || '.' || segment_name || '.' || partition_name || ' ' ||
segment_type) as segment,
       count(extent_id) as extent_count
FROM   dba_extents
WHERE  (segment_type like 'TABLE%' OR segment_type like 'INDEX%') AND
       (owner != 'SYSTEM' AND owner != 'SYS')
GROUP BY owner, segment_name, partition_name, segment_type
ORDER BY EXTENT_COUNT DESC
```

4.3.1.3 Example 3: Embed a Long SQL statement in a PL/SQL Routine

In situations where the SQL statement forming the SQL user-defined metric exceeds 1024 characters, you must embed the SQL statement in a PL/SQL routine. This must be carried out in three steps. In this example, a long SQL statement is used to track tablespaces & free space in them and raise alerts if the free space falls below a user specified threshold. A 2-column SQLUDM is created using this query.

[Example 4-1](#) of a long (more than 1024 characters) SQL statement that returns two values: `tablespace_name` (key) and `free_mb` (value)

Example 4-1 Long SQL Statement

```
select Tablespace, case when (MxAvail <= 15) and (MxFreeMB < 20000) then
'CRITICAL, '||MxUsed||'%' when (MxAvail <= 20) and (MxFreeMB < 20000) then
'WARNING, '||MxUsed||'%' else 'OK' end Error_Level,
MxAvail,MxUsed,MxFreeMB,MxExdMB from (select nvl(b.tablespace_name,
nvl(a.tablespace_name,'UNKNOWN')) as Tablespace, data_files as NumDBFs, mbytes_
alloc as AllocMB, Round( mbytes_alloc-nvl(mbytes_free,0),0) as UsedMB,
```

```

Round(nvl(mbytes_free,0),0) "AllocFreeMB", Round((( mbytes_alloc-nvl(mbytes_
free,0))/ mbytes_alloc)*100,0) as AllocUsed, MaxSize_Mbytes as MxExdMB,
Round(MaxSize_Mbytes - (mbytes_alloc-nvl(mbytes_free,0)),0) as MxFreeMB,
Round((mbytes_alloc/MaxSize_Mbytes*100),0) as MxUsed, Round((MaxSize_Mbytes -
(mbytes_alloc-nvl(mbytes_free,0)))/MaxSize_Mbytes *100,0) as MxAvail from ( select
sum(bytes)/1024/1024 mbytes_free, max(bytes)/1024/1024 largest,tablespace_name
from sys.dba_free_space group by tablespace_name ) a, ( select
sum(bytes)/1024/1024 mbytes_alloc,
sum(decode(maxbytes,0,bytes,maxbytes))/1024/1024 MaxSize_Mbytes, tablespace_
name,count(file_id) data_files from sys.dba_data_files group by tablespace_name )b
where a.tablespace_name (+) = b.tablespace_name order by 1)

```

Because a 2-column SQL UDM is being created (tablespace, free_space_in_MB), an array must be created for the data being returned from this query, as shown in

Example 4-2 Creating an Array of Returned Values

```

CREATE OR REPLACE TYPE tablespace_obj AS OBJECT
(
    tablespace_name VARCHAR2(256),
    free_mb NUMBER
);
/

CREATE OR REPLACE TYPE tablespace_array AS TABLE OF tablespace_obj;
/

```

The next step is to embed the long SQL statement shown in [Example 4-1](#) in a PL/SQL routine as shown in [Example 4-3](#)

Example 4-3 Embedded SQL in a PL/SQL Routine

```

CREATE OR REPLACE FUNCTION calc_tablespace_free_mb
RETURN tablespace_array
IS
    tablespace_data TABLESPACE_ARRAY := TABLESPACE_ARRAY();
BEGIN
    SELECT tablespace_obj(tablespace, mxfreemb)
        BULK COLLECT INTO tablespace_data
    FROM
    (
        select Tablespace, case when (MxAvail <= 15) and (MxFreeMB < 20000) then
        'CRITICAL, '|MxUsed||'%' when (MxAvail <= 20) and (MxFreeMB < 20000) then
        'WARNING, '|MxUsed||'%' else 'OK' end Error_Level,
        MxAvail,MxUsed,MxFreeMB,MxExdMB from (select nvl(b.tablespace_name,
        nvl(a.tablespace_name,'UNKNOWN')) as Tablespace, data_files as NumDBFs, mbytes_
        _alloc as AllocMB, Round( mbytes_alloc-nvl(mbytes_free,0),0) as UsedMB,
        Round(nvl(mbytes_free,0),0) "AllocFreeMB", Round((( mbytes_alloc-nvl(mbytes_
        _free,0))/ mbytes_alloc)*100,0) as AllocUsed, MaxSize_Mbytes as MxExdMB,
        Round(MaxSize_Mbytes - (mbytes_alloc-nvl(mbytes_free,0)),0) as MxFreeMB,
        Round((mbytes_alloc/MaxSize_Mbytes*100),0) as MxUsed, Round((MaxSize_Mbytes -
        (mbytes_alloc-nvl(mbytes_free,0)))/MaxSize_Mbytes *100,0) as MxAvail from (
        select sum(bytes)/1024/1024 mbytes_free, max(bytes)/1024/1024 largest,tablespace_
        _name from sys.dba_free_space group by tablespace_name ) a, ( select
        sum(bytes)/1024/1024 mbytes_alloc,
        sum(decode(maxbytes,0,bytes,maxbytes))/1024/1024 MaxSize_Mbytes, tablespace
        _name,count(file_id) data_files from sys.dba_data_files group by tablespace_name
        )b where a.tablespace_name (+) = b.tablespace_name order by 1)
    ) CUSTOMER_QUERY;

```

```

RETURN tablespace_data;
END calc_tablespace_free_mb;
/

```

The final step in the process is to create the 2-column UDM with the following:

- **Metric Type:** NUMBER
- **SQL Query Output:** Two Columns
- **SQL Query:**

```

SELECT tablespace_name, free_mb
FROM TABLE(CAST(calc_tablespace_free_mb as TABLESPACE_ARRAY))

```

4.4 Notifications, Corrective Actions, and Monitoring Templates

User-Defined Metrics, because they are treated like other metrics, can take advantage of Enterprise Manager's notification system, corrective actions and monitoring templates.

Note: Corrective actions and monitoring templates support both OS user-defined metrics and SQL-based user-defined metrics that return single scalar values.

4.4.1 Getting Notifications for User-Defined Metrics

As with regular metrics, you can receive e-mail notifications when user-defined metric critical or warning alert severities are reached. Assuming you have already defined your e-mail addresses and notification schedule, the remaining task is to set up a notification rule for the user-defined metric.

To set up notification rules:

1. Click Preferences.
2. From the vertical navigation bar, click Rules if you are a Super Administrator or My Rules if you are a regular Enterprise Manager administrator.
3. Click Create to define a new notification rule. The Create Notification Rule pages appear.
4. From the General page, enter the required rule definition information and choose Target Type Host for OS-based user-defined metrics or choose Database Instance or Cluster Database for SQL-based user-defined metrics.
5. On the Metrics page, click Add. A list of available metrics appears. To view all metrics simultaneously, choose Show All from the drop-down menu.
6. Select User-Defined Numeric Metric or User-Defined String Metric based on the type of value returned by your user-defined metric.
7. In the Objects column, choose whether you want to receive notification for all user-defined metrics (All Objects) or specific user-defined metrics (Select).

Add Metrics

Select the metrics and their severities for which you would like to receive notifications. Cancel Continue

Metrics

Search Go Clear

[Previous 10](#) 61-70 of 316 [Next 10](#)

Select All | Select None

Select Metric	Objects
<input type="checkbox"/> Data Not Received (logs)	<input checked="" type="radio"/> All Objects (Name) <input type="radio"/> Select <input type="text"/>
<input type="checkbox"/> Data Not Received (MB)	<input checked="" type="radio"/> All Objects (Name) <input type="radio"/> Select <input type="text"/>
<input type="checkbox"/> Database Block Changes (per second)	n/a
<input type="checkbox"/> Database Block Changes (per transaction)	n/a
<input type="checkbox"/> Database Block Gets (per second)	n/a
<input type="checkbox"/> Database Block Gets (per transaction)	n/a
<input type="checkbox"/> Database CPU Time (%)	n/a
<input type="checkbox"/> Database Time (centiseconds per second)	n/a
<input type="checkbox"/> Database Time Spent Waiting (%)	<input checked="" type="radio"/> All Objects (Wait Class) <input type="radio"/> Select <input type="text"/>
<input type="checkbox"/> Database Vault Attempted Violations - Command Rules	<input checked="" type="radio"/> All Objects (Database Vault Command Rule;Attempted Violation Time) <input type="radio"/> Select <input type="text"/>

[Previous 10](#) 61-70 of 316 [Next 10](#)

When choosing the Select option, enter the name of the user-defined metric, or specify multiple user-defined metrics separated by commas. You can use the wildcard character (%) to match patterns for specific user-defined metrics.

You can search for available user-defined metrics using the search function (flashlight icon). However, search results will only show user-defined metrics that have at least one collected data point. For metrics that have not yet collected at least one data point, as may be the case for a newly created user-defined metric, you must specify them in the Select text entry field.

The format used to specify individual user-defined metrics in the Select text entry field depends on the type of value(s) returned. The following three examples illustrate how to specify user-defined metrics for three possible return value situations.

Formats shown in the following table summarize how to specify multiple user-defined metrics returning single and double values. Specific examples follow.

Return Values	Select Text Entry Field Format
Single	MY_UDM1,MY_UDM3,MY_UDM5
Double (2-column), All Key Values	MY_UDM1;% ,MY_UDM2;% ,MY_UDM4;% User-defined metrics returning two columns are specified as <UDM Name>;% where "%" is a wildcard signifying ALL key values associated with that user-defined metric. <i>Note: 2-column values apply to SQL user-defined metrics only.</i>
Double (2-column), Specific Key Values	MY_UDM1;K1,MY_UDM1_K2 <i>Note: 2-column values apply to SQL user-defined metrics only.</i>

Example 1: Host User-Defined Metrics Returning a Single Value

In this example, you will receive notifications for two host user-defined metrics: MY_UDM9 and MY_UDM11. Both metrics return numeric values. To receive notifications for all host user-defined metrics, select the **All Objects (Script)** option where **Script** refers to the user-defined metric name.

Add Metrics
 Select the metrics and their severities for which you would like to receive notifications Cancel Continue

Metrics
 Search Go Clear

Select All | Select None Previous 10 71-76 of 76 Next

Select Metric	Objects
<input type="checkbox"/> Total Processes	n/a
<input type="checkbox"/> Total Users	n/a
<input type="checkbox"/> User Defined Numeric Metric	<input type="radio"/> All Objects (Script) <input checked="" type="radio"/> Select MY_UDM9, MY_UDM11
<input type="checkbox"/> Virtual IP Relocated	<input type="radio"/> All Objects (Virtual IP Name) <input type="radio"/> Select
<input type="checkbox"/> Windows event log message	<input type="radio"/> All Objects (Log Name;Source;Event ID) <input type="radio"/> Select
<input type="checkbox"/> Windows Event Severity	<input type="radio"/> All Objects (Log Name;Source;Event ID) <input type="radio"/> Select

Example 2: SQL User-Defined Metric Returning a Single Value

In this example, you will receive notifications for four SQL user-defined metrics: MY_UDM1 and MY_UDM3 (numeric value returned) and MY_UDM5 and MY_UDM7 (string value returned). To receive notifications for all SQL user-defined metrics, select the **All Objects (Script)** option where **Script** refers to the user-defined metric name.

Add Metrics
 Select the metrics and their severities for which you would like to receive notifications. Cancel Continue

Metrics
 Search Go Clear

Select All | Select None

Select Metric	Objects
<input type="checkbox"/> Enqueue: UL: User-defined - contention (%)	n/a
<input type="checkbox"/> User-Defined Numeric Metric	<input type="radio"/> All Objects (Metric ID;Key) <input checked="" type="radio"/> Select
<input type="checkbox"/> User-Defined Numeric Metric	<input type="radio"/> All Objects (Script) <input checked="" type="radio"/> Select MY_UDM1, MY_UDM3
<input type="checkbox"/> User-Defined String Metric	<input type="radio"/> All Objects (Metric ID;Key) <input checked="" type="radio"/> Select
<input type="checkbox"/> User-Defined String Metric	<input type="radio"/> All Objects (Script) <input checked="" type="radio"/> Select MY_UDM5, MY_UDM7

Example 3: SQL User-Defined Metrics Returning Two Values

To receive notifications for SQL user-defined metrics returning two values, the syntax is slightly different. In this situation, you specify the user-defined metric using the following format: *<Metric ID;Key>* where **Metric ID** is the name of the user-defined metric and **Key** is the specific key value you want returned. To receive notifications for all SQL user-defined metrics, select the **All Objects (Metric ID;Key)** option.

Add Metrics
 Select the metrics and their severities for which you would like to receive notifications. Cancel Continue

Metrics

Search

Select All | Select None

Select Metric	Objects
<input type="checkbox"/> Enqueue: UL: User-defined - contention (%)	n/a
<input type="checkbox"/> User-Defined Numeric Metric	<input type="radio"/> All Objects (Metric ID;Key) <input checked="" type="radio"/> Select[MY_UDM1;% , MY_UDM2;%]
<input type="checkbox"/> User-Defined Numeric Metric	<input type="radio"/> All Objects (Script) <input type="radio"/> Select
<input type="checkbox"/> User-Defined String Metric	<input type="radio"/> All Objects (Metric ID;Key) <input checked="" type="radio"/> Select[MY_UDM4;% , MY_UDM5;K1 , MY_UDM5;K2]
<input type="checkbox"/> User-Defined String Metric	<input type="radio"/> All Objects (Script) <input type="radio"/> Select

For SQL user-defined metrics returning two values, you have the option of receiving notifications for all key values or just specific values for a given metric. Use a wildcard character (%) to specify all key values. Example: *MY_UDM1;% , MY_UDM2;%*.

To receive notifications for specific key values, specify the exact key value to be returned. Example: *MY_UDM5;K1,MY_UDM5;K2* (only key values K1 and K2 are returned from MY_UDM5).

8. Select the severity or corrective action state for which you would like to receive the notification and then click Continue.
9. If you want to receive e-mail for the specified user-defined metric, go to the Notification Rule and check the "Send me E-mail" option.
10. Click OK to create the new notification rule. If you made the notification rule public, other administrators can subscribe to the same rule.

4.4.2 Setting Corrective Actions for User-Defined Metrics

Corrective actions allow you to specify automated responses to alerts ensuring that routine responses to alerts are automatically executed. Corrective actions can be defined for both SQL and OS-based user-defined metrics.

To set up corrective actions:

1. From a target home page, click Metric and Policy Settings from Related Links.
2. Locate and edit the user-defined metric.
3. From the Edit Advanced Settings page, click Add under Corrective Actions for the Critical or Warning alert severity and define the corrective action. Corrective actions can be defined for one or both alert severities.

4.4.3 Deploying User-Defined Metrics Across Many Targets Using Monitoring Templates

Monitoring Templates simplify the task of standardizing monitoring settings across your enterprise by allowing you to specify the monitoring settings once and applying them to your monitored targets. You can thus use Monitoring Templates as a way to propagate user-defined metrics across a large number of targets.

Assuming you have created the user-defined metric on the host or database target, you can use Monitoring Templates to propagate the user-defined metric to other hosts or database targets.

To create a Monitoring Template for the user-defined metric:

1. Click Setup.
2. From the vertical navigation bar, click Monitoring Templates
3. Click Create. The Copy Target Settings page appears.
4. Specify the host or database on which you defined the user-defined metric and click Continue.
5. Fill in the requisite information on the General page.
6. On the Metric Thresholds page, you can choose to keep or remove the other metrics that have been copied over from the target.
7. You can also edit the user-defined metric's thresholds, collection schedule, and corrective actions.
8. On the Policies page, you can choose to keep or remove any policy rules that have been copied over from the target.
9. Click OK to save the template settings to the Management Repository.

Once the template containing the user-defined metric has been created, you can propagate the user-defined metric by applying the template to other hosts or databases.

Important: For OS-based user-defined metrics, you will first need to separately deploy the OS Script used by the user-defined metric to all destination hosts. The OS Script should reside in the same location across all host targets.

For SQL-based user-defined metrics, if the SQL query specified is a function call, then you need to create this function across all databases on which the SQL-based user-defined metric will be created.

To apply the monitoring template:

1. On the Monitoring Templates page, select the monitoring template and click Apply.
2. On the Apply Monitoring Template page, add the targets on which the user-defined metric should be created.
3. If a two-column SQL-based user-defined metric is part of the template, the Metric with Multiple Thresholds option is applied according one of the following situations:
 - **Situation One:** The target to which the template will be applied does not contain the two-column SQL user-defined metric defined in the template. In this situation, regardless of which Metric with Multiple Thresholds option is chosen, the user-defined metric is copied to the target when you apply the template.
 - **Situation Two:** The target to which the template will be applied does contain the two-column SQL user-defined metric. The name (case insensitive) and return value (numeric, scalar, or two column) of both the target and template user-defined metrics must match. In this situation, you must select one of the Metric with Multiple Thresholds options:

- *Apply threshold settings for monitored objects common to both template and target:* For only those keys which the target has in common with the template, the target threshold values will be set to the values defined in the template. This option is chosen by default and is recommended for most situations.
- *Duplicate threshold settings on target:* For keys which are common between target and template, the thresholds will be set to the values defined in the template. Any extra keys (and their thresholds) that exist in the template but not on the target will be copied to the target in anticipation that these keys will be created in the target at some point in the future. Any extra keys (and their thresholds) that exist on the target but not in the template will be deleted from the target.

Important: If there are other metrics in the monitoring template, refer first to the Enterprise Manager online help for implications of what this option means for these other metrics.

4. Click Continue.
5. On the subsequent page, specify the credentials that should be used when running the user-defined metric on the destination targets.
6. Click Finish.
7. When you return back to the Monitoring Templates page, check that "Pending Apply Operations" count for your template is zero. This indicates the number of template apply operations that could be pending. Once they are all complete, the count should be zero.

Deploying User-Defined Metrics Using Scripts

An alternate method of deploying user-defined metrics to large numbers of targets is to use the Enterprise Manager Command Line Interface (EMCLI). Using the EMCLI "apply_template" verb, you can deploy user-defined metrics via custom scripts. For more information about the "apply_template" verb, see the *Oracle Enterprise Manager Command Line Interface* manual.

4.4.4 Deleting User-Defined Metrics Across Many Targets Using Monitoring Templates

Just as templates can be used to deploy user-defined metrics across targets, templates can also be used to delete these metrics across targets should these metrics no longer be in use.

To create a Monitoring Template for the user-defined metric:

1. Click Setup.
2. From the vertical navigation bar, click Monitoring Templates
3. Click Create. The Copy Target Settings page appears.
4. Specify the host or database on which there is a user-defined metric that needs to be deleted and click Continue.
5. Fill in the requisite information on the General page.
6. On the Metric Thresholds page, remove all metrics from the template except the user-defined metric to be deleted.
7. Click the pencil icon to access the Edit Advanced Settings page.

8. Check the Mark for Delete option and click Continue. The Mark for Deletion icon now appears next to the user-defined metric on the Metric Thresholds page.
9. On the Policies page, remove all policies.
10. Click OK to save the template settings to the Management Repository.

To apply the monitoring template:

1. On the Monitoring Templates page, select the monitoring template and click Apply.
2. On the Apply Monitoring Template page, add the targets on which the user-defined metric should be deleted.
3. Click Continue.
4. On the subsequent page, specify the credentials that should be used when running the user-defined metric on the destination targets.
5. Click Finish.
6. On the Monitoring Templates page, check that "Pending Apply Operations" count for your template is zero. This indicates the number of template apply operations that could be pending. Once they are all complete, the count should be zero.

Enterprise manager will delete all user-defined metrics found on the selected target that match the following criteria:

- Name of the user-defined metrics (case insensitive)
- Return value of the user-defined metric (numeric or scalar)
- For SQL-based user-defined metrics, the output of the query (single value or two columns). The match does not take into consideration the actual script used by the user-defined metric. For this reason, even though the script on the target user-defined metric may be different from that of the template user-defined metric, the target user-defined metric will still be deleted.
- Host User-Define Metrics using scripts: You must delete the script from the host on which you want the UDM deleted.
- SQL-based user-defined metrics using function calls: You must delete the function from the database on which you want the user-defined metric deleted.

4.5 Changing User-Defined Metric Credentials

As discussed earlier, user-defined metrics require valid credentials (username and password) in order to execute monitoring scripts/SQL queries. For this reason, both the monitored target's password and the password defined in the user-defined metric must match. This can be problematic if target passwords are changed frequently. For environments with a large number of targets, you can use the Enterprise Manager Command Line Interface (EMCLI) to change the target password and user-defined metric password simultaneously using scripts. Use the "update_password" verb to change the target password. This password change is then propagated to all features of Enterprise Manager that use the specified username, which includes preferred credentials, corrective actions, jobs, and both host and SQL-based user-defined metrics.

The following example changes the password associated with the OS user *sysUser* from *sysUserOldPassword* to *sysUserNewPassword*.

Example 4-4 Host Password Change

```
update_password -target_type=host -target_name=MyHost -credential_  
type=HostCreds -key_column=HostUserName:sysUser  
-non_key_column=HostPassword:sysUserOldPassword:sysUserNewPassword
```

The next example changes the password associated with the database user *sys* from *sysPassword* to *sysNewPassword*.

Example 4-5 Database Password Change

```
update_password -target_type=oracle_database -target_name=ORCL -credential_  
type=DBCreds -key_column=DBUserName:sys  
-non_key_column=DBPassword:sysPassword:sysNewPassword:DBAROLE
```

For more information about EMCLI, see the *Oracle Enterprise Manager Command Line Interface* guide.

Group Management

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

- [Introduction to Groups](#)
- [Managing Groups](#)
- [Out-of-Box Reports](#)
- [Redundancy Groups](#)
- [Privilege Propagating Groups](#)

5.1 Introduction to Groups

Today's IT operations can be responsible for managing a great number of components, such as databases, application servers, hosts, or other components, which can be time consuming and impossible to manage individually. The Enterprise Manager Grid Control group management system lets you combine components (called targets in Enterprise Manager) into logical sets, called groups. This enables you to organize, manage, and effectively monitor the potentially large number of targets in your enterprise.

Enterprise Manager Groups can include:

- Targets of the same type, such as:
 - All hosts in your data center
 - All of your production databases
- Targets of different types, such as:
 - The database, application server, listener, and host that are used in your application environment
 - Targets operating within a particular data center region

Note: An Enterprise Manager "System," used specifically to group the components on which a service runs, is a special kind of Enterprise Manager group. Many of the functions and capabilities for groups and systems are similar.

Typically you can gather together targets that you want to manage as a group. If you use the target properties (for example, *Line of Business* or *Deployment Type*) to put operational information about your targets in Enterprise Manager, you can use these

properties when creating groups to locate targets. For example, you could search for all databases of *Deployment Type = Production* and belonging to *Line of Business 'HCM'*. You can also create a group hierarchy and use nested groups.

5.2 Managing Groups

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

- View a summary status of the targets within the group.
- Monitor outstanding alerts and policy violations for the group collectively, rather than individually.
- Monitor the overall performance of the group through performance charts.
- Perform administrative tasks, such as scheduling jobs for the entire group, or blacking out the group for maintenance periods.

You can also customize the console to provide direct access to group management pages.

5.2.1 Group Home Page

The Group Home page, shown in [Figure 5-1](#), is the central location for monitoring information. The Group Home page provides the following features:

- Availability pie chart that provides at-a-glance information on the current status across all members so you can easily assess the percentage of members that are up, and the percentage of members that are unavailable. You can quickly drill down for information if any member target is down.
- Roll-up of open alerts and policy violations, categorized by severity, so you can quickly focus on the most critical problems first. Alerts and violations that have occurred within the last 24 hours highlight problems that recently occurred.
- Access to the Policy Trend Overview page, which provides a comprehensive view of the group for compliance with policy rules over a period of time. Using policy charts, you can assess trends such as increased or decreased number of policy violations, changes in the overall group compliance score, and the percentage of members in compliance with your enterprise's policy rules.
- Access to the Security at a Glance page, which provides an overview of the security health of the group. This shows statistics about security policy violations and critical security patches that have not been applied.
- Summary of recent configuration changes across all members in the group, so you can easily determine if a new problem is related to any recent changes.
- Summary of Critical Patch Advisories for Oracle homes within the group.

Figure 5-1 Group Home Page

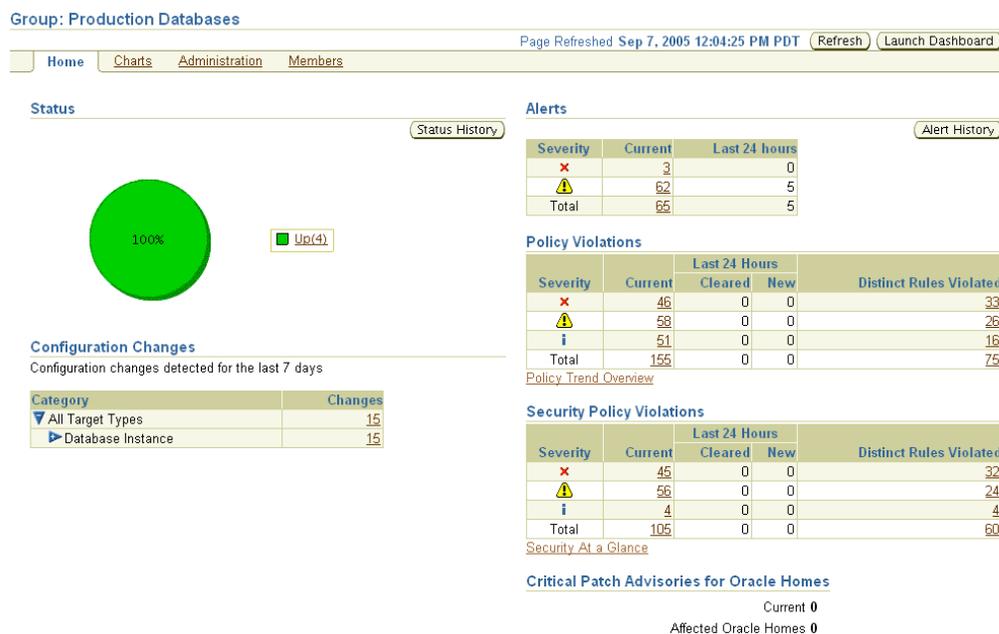
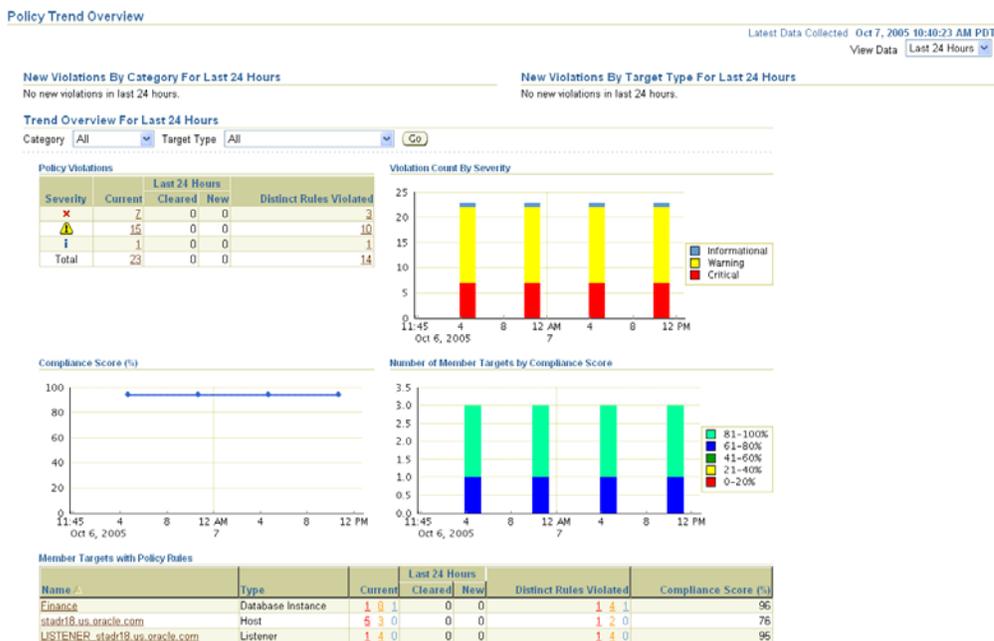


Figure 5-2 shows the Policy Trend Overview page that you can access from the Group Home page.

Figure 5-2 Policy Trend Overview Page

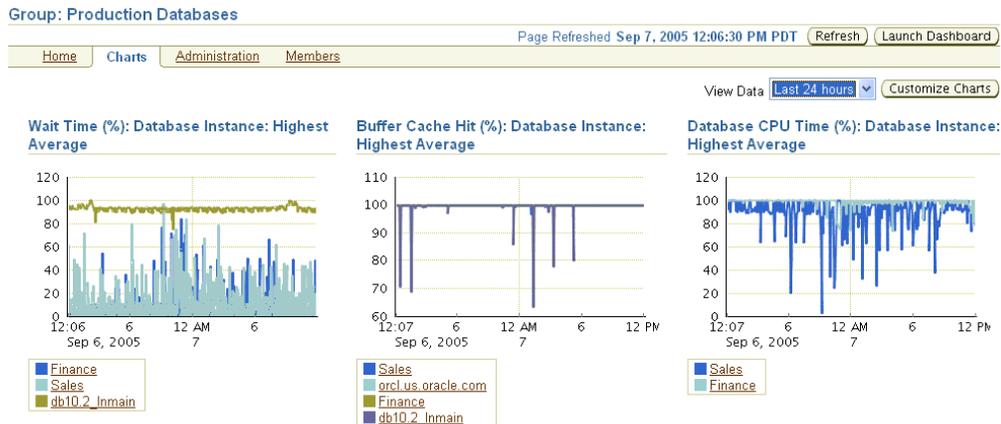


See Also: "Group Home Page" in the Enterprise Manager online help

5.2.2 Group Charts Page

The Group Charts page, shown in [Figure 5–3](#), enables you to monitor the collective performance of the group. Out-of-box performance charts are provided based on the type of members in the group. For example, when databases are part of the group, a Wait Time (%) chart is provided that shows the top databases with the highest wait time percentage values. You can view this performance information over the last 24 hours, last 7 days, or last 31 days. You can also add your own custom charts to the page.

Figure 5–3 Group Charts Page



See Also: "Group Charts Page" in the Enterprise Manager online help

5.2.3 Group Administration Page

The Group Administration page, shown in [Figure 5–4](#), provides a central point for performing common administrative tasks for the group. For example, you can:

- Run jobs or find out the status of currently running jobs against the group.
- Define planned outage windows, called blackouts, on the members of the group to perform maintenance tasks.
- Run SQL commands collectively against the database member targets of the group.
- View the most recent backup for each database in the group.
- View the last 100,000 bytes of the alert log for all databases in the group.
- Use a deployment summary to easily obtain hardware and software inventory information across all member targets.

Figure 5–4 Group Administration Page

Group: Production Databases Page Refreshed Sep 7, 2005 12:13:04 PM PDT [Refresh](#) [Launch Dashboard](#)

[Home](#) [Charts](#) [Administration](#) [Members](#)

Job Activity

Create Job: OS Command

Job executions scheduled to start no more than 7 days ago

Status	Submitted to the Group	Submitted to any member
Scheduled	4	4
Running	0	0
Suspended	0	0
Problem	0	0

Blackouts

Status	Submitted to the Group	Submitted to any Member
Scheduled	0	3
Active	0	0

Database Operations

[Execute SQL](#) [View Backup Report](#)
[Alert Log Content](#)

Configuration Searches

Database Feature Usage

Deployments Summary

View: Database Installations

Software Targets Without Inventory: 3 of 4

Database Installations	Targets	Installations	Interim Patches Applied
Oracle Database 10g 10.2.0.0.0	1	1	No

See Also: "Group Administration Page" in the Enterprise Manager online help

5.2.4 Group Members Page

The Group Members page, shown in Figure 5–5, summarizes information about the member targets in the group. It includes information on their current availability status, roll-up of open alerts and policy 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 sort on any of the columns to rank members by a certain criterion (for example, database targets in order of decreasing wait time percentage). Default key performance metrics are displayed based on the targets you select, but you can customize these to include additional metrics that are important for managing your group.

Figure 5–5 Group Members Page

Group: Production Databases Page Refreshed Sep 7, 2005 12:07:55 PM PDT [Refresh](#) [Launch Dashboard](#)

[Home](#) [Charts](#) [Administration](#) [Members](#)

Search: All [Customize Columns](#)

Name	Type	Status	Alerts	Policy Violations	Wait Time (%)	Compliance Score (%)	Deployment Type	Location
db10.2_inmain	Database Instance		0 4	21 26 41	92.97 ✓	84		
Finance	Database Instance		3 2	14 14 7	8.37 ✓	88	Production	Redwood Shores
orcl.us.oracle.com	Database Instance		0 56	1 8 1	✓	96		
Sales	Database Instance		0 1	10 10 2	0 ✓	93	Production	Redwood Shores

[Home](#) [Charts](#) [Administration](#) [Members](#)

See Also: "Group Members Page" in the Enterprise Manager online help

5.2.5 System Dashboard

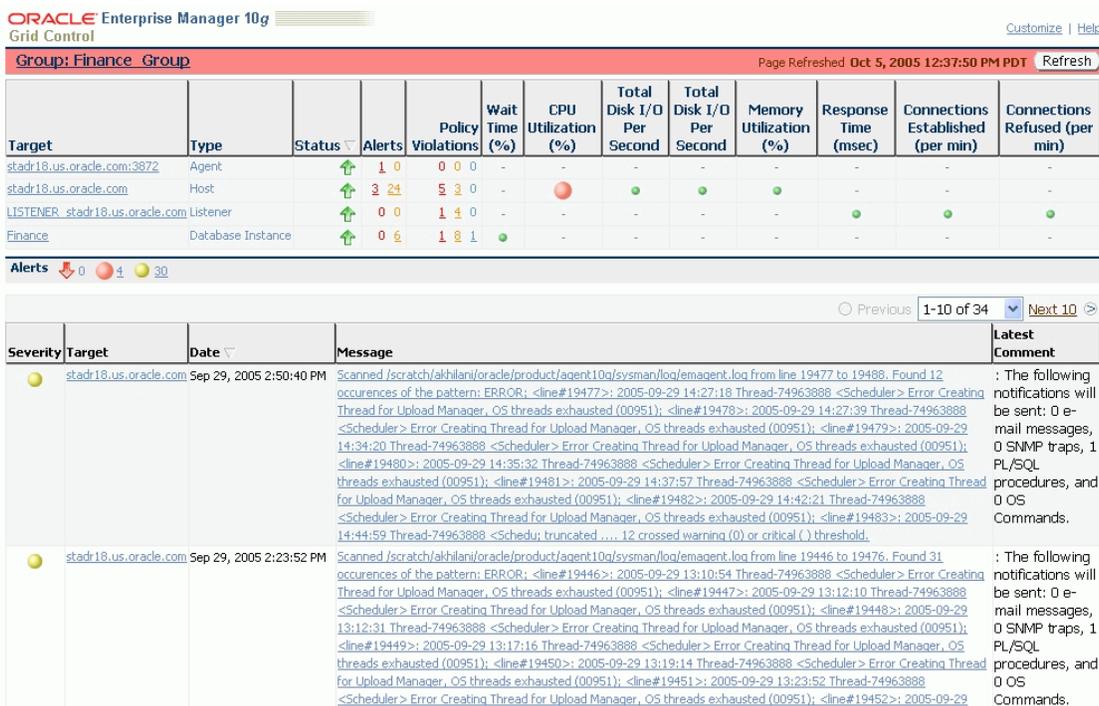
The System Dashboard, shown in Figure 5–6, enables you to proactively monitor the status and alerts in the group as they occur. The color-coded interface is designed to

highlight problem areas using the universal colors of alarm—targets that are down are highlighted in red, metrics in critical severity are shown as red dots, metrics in warning severity are shown as yellow dots, and metrics operating within normal boundary conditions are shown as green dots.

Using these colors, you can easily spot the problem areas for any target and drill down for details as needed. An alert table is also included to provide a summary for all open alerts in the group. The alerts in the table are presented in reverse chronological order to show the most recent alerts first, but you can also click on any column in the table to change the sort order.

You can customize the dashboard according to your needs. You can specify which key metrics should be included in the dashboard and the display names to be used for these metrics. You can also customize the refresh interval to ensure that you always receive timely information about alerts as they are detected. You can launch the System Dashboard in context from any Group Home page. However, using reporting framework features, you can also make the System Dashboard publicly available for any user that has access to a web browser and the Enterprise Manager Reports Web site.

Figure 5–6 System Dashboard



5.3 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 policy violations for a group.

You can access these reports from the **Reports** link in the Related Links section of all Group pages. Figure 5–7 shows the Availability History report for a specified group over the last 31 days.

Figure 5–7 Availability History Report

See Also: [Chapter 8, "Information Publisher"](#)

5.4 Redundancy Groups

A redundancy group is a group that contains members of the same type that function collectively as a unit. A type of redundancy group functions like a single logical target that supports a status (availability) metric. A redundancy group is considered up (available) if at least one of the member targets is up.

You can create and administer a redundancy group from the All Targets page. Redundancy groups support all group management features previously discussed.

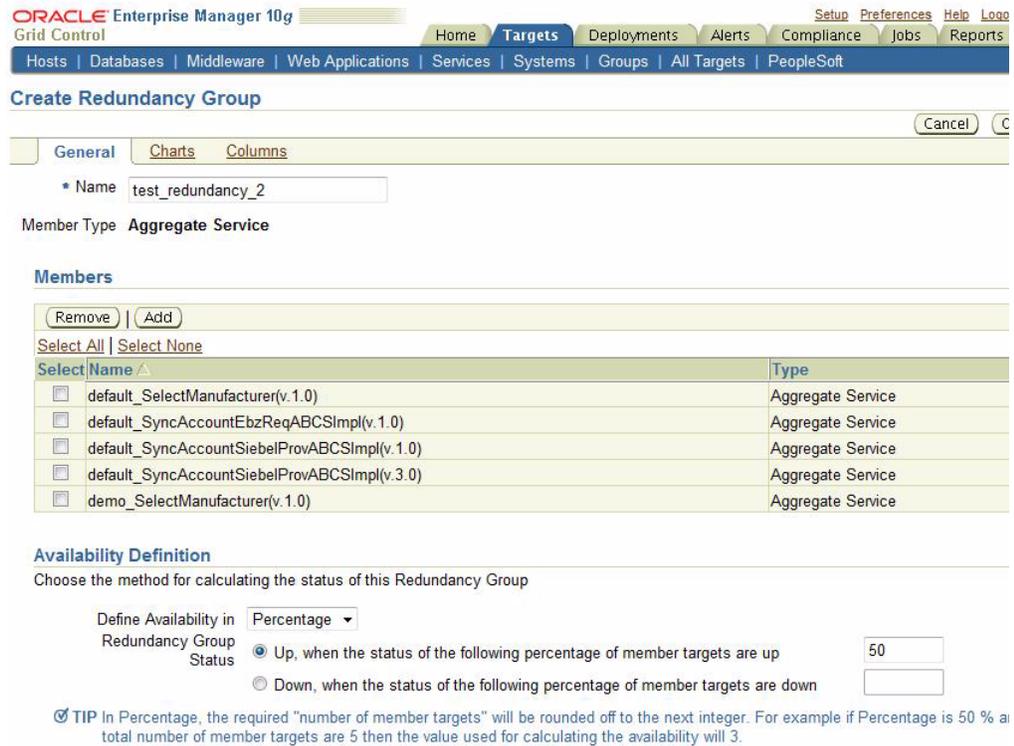
When you define the Redundancy Group, you must choose the member type for the members in the Redundancy Group.

You can define the options for how availability of the redundancy group is calculated by selecting either Number or Percentage:

- **Number** - When you choose Number, you can specify either the number of member targets that should be up in order for the group to be considered up, or the number of member targets that should be down in order for the group to be considered down.
- **Percentage** - When you choose Percentage, you can specify either the minimum percentage of member targets that must be up in order for the group to be considered up, or the minimum percentage of member targets that are down in order to consider the group to be down. If you choose Percentage, the required number of member targets will be rounded off to the next integer. For example if you define the Percentage as 50% and the total number of member targets is 5, then the value used for calculating the availability will be 3.

Figure 5–8 shows the Create Redundancy Group page while defining the group using Percentage availability.

Figure 5–8 Redundancy Group Example



Do not use redundancy groups if the group you want to model is an Oracle Real Application Clusters database, host cluster, HTTP server high availability group, or OC4J high availability group. Instead, you can use the following specialized target types for this purpose:

- Cluster
- Cluster Database
- HTTP HA Group
- OC4J HA Group

5.5 Privilege Propagating Groups

Privilege propagating groups enable administrators to grant privileges to other administrators in a manner in which new administrators get the same privileges as its member targets. For example, granting *operator* privilege on a group to an Administrator grants him the *operator* privilege on its member targets and also to any members that will be added in the future. Privilege propagating groups can contain individual targets or other privilege propagating groups.

Privileges on the group can be granted to an Enterprise Manager user or a role. Use a role if the privileges you want to grant are to be granted to a group of EM users. See [Figure 5–9, "Granting Privileges On a Group To a Role"](#).

Figure 5–9 Granting Privileges On a Group To a Role

Search and Select Administrator or Role Cancel Select

Search

Search

Type Role

Results

✔ **TIP** Owner has Full privilege on the target. Super Administrators have Full privilege on all targets.

Select All Select None			
Select	Name ▲	Type	Description
<input checked="" type="checkbox"/>	E2E_TEST	Role	
<input checked="" type="checkbox"/>	E2E_TEST2	Role	
<input checked="" type="checkbox"/>	PUBLIC	Role	
<input type="checkbox"/>	TEST_ROLE	Role	
<input type="checkbox"/>	TEST_ROLE2	Role	

Cancel Select

For example, suppose you create a large privilege propagating group and grant a privilege to a role which is then granted to administrators. If new targets are later added to the privilege propagating group, then the administrators receive the privileges on the target automatically. Additionally, when a new administrator is hired, you only need to grant the role to the administrator for the administrator to receive all the privileges on the targets automatically.

5.5.1 Creating Privilege Propagating Groups

The privilege propagating group creation function is a privileged activity. The privilege propagating group feature contains two new privileges:

- **Create Privilege Propagating Group**
This privileged activity allows the administrators to create the privilege propagating groups. Administrators with this privilege can create propagating groups and delegate the group administration activity to other users.
- **Group Administration**
This privilege can be granted to administrators on specific group targets and is used to delegate the group administration activities to other administrators. It is granted to both conventional and privilege propagating groups.

5.5.2 Using the Group Administration Privilege

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

5.5.3 Adding Members to Privilege Propagating Groups

The target privileges granted on a propagating group are propagated to member targets. The administrator grants target objects scoped to another administrator, the grantee maintains the same privileges on member targets. The propagating groups maintain the following features:

- The administrator with a Create Privilege Propagating Group privilege will be able to create a propagating group
- To add a target as a member of a propagating group, the administrator must have *Full* target privileges on the target

You can add any non-aggregated target as the member of a privilege propagating group. For aggregated targets in Grid Control version 10.2.0.5, cluster and RAC databases and other propagating groups can be added as members (cluster and RAC databases must be added via the *emcli* verb). There is no support for this through the Enterprise Manager interface in version 10.2.0.5. Grid Control version 11.1, however, supports more aggregated target types, such as redundancy groups, systems and services. These, along with cluster and RAC databases, can be added in version 11.1 via the Grid Control Console.

If you are not the group creator, you must have at least the *Full* target privilege on the group to add a target to the group.

5.5.4 Converting Conventional Groups to Privilege Propagating Groups

In Enterprise Manager version 11.1, you can convert conventional groups to privilege propagating groups (and vice-versa) through the use of the specified EMCLI verb. Two new parameters have been added in the *modify_group* EMCLI 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 EMCLI verbs: *modify_system*, *modify_redundancy_group*, and *modify_aggregate_service*.

The EMCLI 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 EMCLI verbs, see the *EMCLI Reference Manual*.

Job System

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 Grid 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](#)

6.1 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 6-1](#)) 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 6–1 Job Activity Page

Page Refreshed Jan 19, 2010 5:36:09 PM

Advanced Search

Name:

Owner: All

Status: Active

Scheduled Start: All

Job Type: All

Target Type: All Target Types against which jobs were executed

Target Name:

Go Simple Search

View: Runs

View Results Edit Create Like Copy To Library Suspend Resume Stop Delete Create Job OS Command Go

Select	Name	Status (Executions)	Scheduled	Targets	Target Type	Owner	Job Type
<input checked="" type="radio"/>	MY_JOB	1 Scheduled	Jan 19, 2010 6:20:00 PM PST	dadvmn0630.us.oracle.com	Host	SYSMAN	OS Command

Besides accessing the Job Activity page from the Jobs tab as shown in Figure 6–1, you can also access this page from any Database or Cluster Database property page (Home, Performance, Availability, and so forth) by clicking the Jobs link in the Related Links section. 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.

6.1.1 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 patch job against several databases. A typical question would be: Were all the patches successful, and if not, which patches failed? If this patch job runs every week, you would want to know which patches 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.

6.1.2 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 be able 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 Grid Control online help.

6.2 Preliminary Considerations

Before proceeding to the procedural information presented in [Section 6.3, "Creating Jobs"](#) on page 6-4, it would be beneficial to read these topics presented in the sections below:

- [Using Pre-defined Tasks](#)
- [Creating Scripts](#)
- [Sharing Job Responsibilities](#)
- [Jobs and Groups](#)

6.2.1 Using Pre-defined Tasks

Enterprise Manager provides predefined job tasks for database targets and deployments. A job task is used to contain predefined, unchangeable logic; for example: patch an application, back up a database, and so forth.

The predefined database jobs include backup, export, import, patch, and clone. These are available from the Availability, Data Movement, and Software and Support property pages you can access from the Database Instance Home page. The predefined jobs associated with deployments include patching, cloning Oracle homes, and cloning databases.

6.2.2 Creating Scripts

In addition to 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.

6.2.3 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.

6.2.4 Jobs and 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 its member targets 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.

For information on how to submit a job against a group of targets, see "[Specify General Job Information](#)" on page 6-5.

See Also: [Chapter 5, "Group Management"](#)

6.3 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](#)

6.3.1 Selecting a Job Type

Using the Job System, you can create a job by selecting one of the job types, including those below, from the Create Job drop-down in the Job Activity page.

- **Block Agent** — Blocks an Agent and submits it as a job function. See the discussion below for more information on blocked Agents.
- **Clone Home** — Copies the known state of an Oracle Home. For example, after you have an Oracle home in a known state (you have chosen particular install options for it, applied required patches to it, and tested it, you may want to clone this Oracle home to one or more hosts.
- **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 in much the same way as a regular job.
See [Section 6.3.4, "Creating a Multi-task Job"](#) for more information.
- **OS Command** — Runs an operating system command or script.
- **SQL Script** — Runs a user-defined SQL or PL/SQL script.

Blocked Agents

A blocked Agent is a condition where the 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, and no data will then penetrate the system.

6.3.2 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 Manager Grid Control Home page, click the **Jobs** tab. The Job Activity page appears.
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.
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. To submit a job against a group of targets, select Group as the Target Type.

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.

If you specify a composite (for example, a group or service) 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.

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`. The shell scripts size is limited to 2 GB.

To control the maximum step 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.

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 6.3.2.1, "Specifying a Single Operation"](#) and [Section 6.3.2.2, "Specifying a Script"](#).
3. Click the **Credentials** property page link.

Task 4 Specify Credentials - (optional)

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. You also do not need to provide input on this page if you are not sharing the job.

- **To use preferred credentials:**
 1. Select the **Use Preferred Credentials** radio button.

If the target for the OS Command job is a database or database group, the host credentials for the database target are used. These are specified on the Database Preferred Credentials page and are different from the host credentials for the host on which the database resides.

You can set preferred credentials by clicking the **Preferences** link at the top of the Enterprise Manager console, then clicking the **Preferred Credentials** link. The Preferred Credentials page appears, where you can click the Set Credentials icon of the target type for which you want to provide credentials.

2. Select either **Normal Host Credentials** or **Privileged Host Credentials** from the Host Credentials drop-down.

If you select host targets for your job, you can choose to use either Normal or Privileged credentials for the job. You specify these separately on the Preferred Credentials page.

- **To override preferred credentials:**

1. Select the **Override Preferred Credentials** radio button.

Note that override credentials apply to all targets.

2. Optionally provide Sudo or PowerBroker details, which 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.

Note: You cannot use Sudo or PowerBroker credentials with targets managed by pre-10.2.0.4 Agents or with Windows targets. See your super administrator about setting up these features if they are not currently enabled.

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)**

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.

- **Repeating**

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:

- **Indefinite** — Select to allow your job to continue running indefinitely. Otherwise, select Specified Date.
- **Specified Date** — Set the date for your job to stop running, and set the time for the stop date you specified.

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.

2. Specify 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 does not run if the job does not start before 2 p.m.

3. Click the **Access** property page link.

Task 6 Specify Who Can Access the Job - (optional)

You do not need to provide input for this page if you do not want to share 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. **Optional:** 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 right 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 right to a job. 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.

For more information on access levels, see [Section 6.3.2.3, "Access Level Rules"](#).

2. Optional: 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. Optional: Define a notification rule.

You can use the Notification system (rule creation) to easily associate specific jobs with a notification rule. The Grid 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 3, "Notifications"](#) for this information.

Task 7 Concluding 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 rights. 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

6.3.2.1 Specifying a Single Operation

The following information applies to step 2 in [Task 3, "Specify Parameters"](#) on page 6-6.

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
```

6.3.2.2 Specifying a Script

The following information applies to step 2 in [Task 3, "Specify Parameters"](#) on page 6-6.

The value you specify in the OS Script field is used as stdin for the command interpreter, which defaults to `/bin/sh` on Linux and `cmd /c` on Windows. You can override this with another interpreter; for example: `%perlbin%/perl`. The shell scripts size is limited to 2 GB.

To control the maximum output size, set the `mgmt_job_output_size_limit` parameter in `MGMT_PARAMETERS` to the required limit. Values less than 10 KB and greater than 2 GB are ignored. The default output size is 10 MB.

You can run a script in several ways:

- **OS Scripts** — Specify the path name to the script in the OS Script field. For example:

OS Script field: /path/to/mycommand

Interpreter field:

- **List of OS Commands** — You do not need to enter anything in the Interpreter field for the following example of standard shell commands for Linux or Unix systems. The OS's default shell of /bin/sh or cmd/c will be used.

```
/usr/local/bin/myProg arg1 arg2
mkdir /home/$USER/mydir
cp /dir/to/cp/from/file.txt /home/$USER/mydir
/usr/local/bin/myProg2 /home/$USER/mydir/file.txt
```

When submitting shell-based jobs, be aware of the syntax you use and the targets you choose. This script would 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
```

6.3.2.3 Access Level Rules

The following rules apply to the Create Job Access and Edit Job Access property page referenced in [Section 6.3.2, "Creating an OS Command Job"](#).

- 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 that job, ALL administrators will, by default, have VIEW access on the newly created job.

6.3.3 Creating a SQL Script Job

The basic process for creating a SQL script job is the same as described in [Section 6.3.2, "Creating an OS Command Job."](#) The following sections provide supplemental information specific to script jobs.

6.3.3.1 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.

6.3.3.2 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 plsqli_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.

6.3.3.3 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.

- **Override Preferred Credentials** —

Select this choice if you want to override the preferred credentials for all targets, then enter the preferred credentials you want the job to use on all targets. In addition to overriding credentials, you can also specify that Sudo or PowerBroker be used to run the job.

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). However, jobs and corrective actions for which you specified the Override Credentials option do not automatically pick up these changes. You need to edit such jobs and corrective actions, and change the password as needed.

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.

6.3.3.4 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;
```

Example 4

```
-- interactive mode only

col test new_value bye
SELECT 123 test
FROM dual;
exit &bye;
```

6.3.4 Creating a Multi-task Job

The basic process for creating a multi-task job is the same as described in [Section 6.3.2, "Creating an OS Command Job."](#) The following sections provide supplemental information specific to multi-task jobs.

6.3.4.1 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.

6.3.4.2 Specifying Targets for a Multi-task Job

You can run a multi-task job against any targets in which jobs are defined that can be used as tasks. Not all job types can be used as tasks.

The Targets 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 Targets drop-down, you then select the targets to run the job against by clicking Add in the Targets section.

6.3.4.3 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 the last task of the multi-task job fails. 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.

6.4 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 Grid Control Console Home page. [Figure 6-2](#) shows the All Targets Jobs information at the bottom of the Grid Control Console Home page.

Figure 6–2 Summary of Target Jobs on the Grid Control Console Home 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.

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, the Grid Control Console, the Fusion Middleware Control Console, and Database Control.

This chapter also explains the various emctl commands, exit codes, 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 Fusion Middleware Control](#)
- [Controlling the Database Control on UNIX](#)
- [Guidelines for Starting Multiple Enterprise Manager Components on a Single Host](#)
- [Starting and Stopping Oracle Enterprise Manager 11g Grid Control](#)
- [Additional Management Agent Commands](#)
- [emctl Commands](#)
- [Using emctl.log File](#)

7.1 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](#)

7.1.1 Starting, Stopping, and Checking the Status of the Management Agent on UNIX

To start, stop, or check the status of the Management Agent on UNIX systems:

1. Change directory to the `AGENT_HOME/bin` directory.
2. Use the appropriate command described in [Table 7-1](#).

For example, to stop the Management Agent, enter the following commands:

```
$PROMPT> cd AGENT_HOME/bin
```

```
$PROMPT> ./emctl stop agent
```

Table 7–1 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 7–1).

Example 7–1 Checking the Status of the Management Agent

```
$PROMPT> ./emctl status agent
Oracle Enterprise Manager 11g Release 1 Grid Control 11.1.0.1.0
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
-----
Agent Version      : 11.1.0.1.0
OMS Version       : 11.1.0.1.0
Protocol Version  : 11.1.0.0.0
Agent Home        : /ade/example_username/oracle
Agent binaries    : /ade/example_username/oracle
Agent Process ID  : 10677
Parent Process ID : 10593
Agent URL         : https://example.us.oracle.com:11865/emd/main/
Repository URL   : https://example.us.oracle.com:13123/em/upload
Started at       : 2010-01-05 10:35:58
Started by user  : example_username
Last Reload      : 2010-01-06 04:42:33
Last successful upload      : 2010-01-06 04:42:35
Total Megabytes of XML files uploaded so far : 1041.04
Number of XML files pending upload           : 0
Size of XML files pending upload (MB)       : 0.00
Available disk space on upload filesystem   : 71.53%
Last successful heartbeat to OMS            : 2010-01-06 04:42:44
-----
Agent is Running and Ready
$PROMPT>
```

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, set the following variables in the shell:

```
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.

7.1.2 Starting and Stopping the Management Agent on Windows

When you install the Oracle Management Agent on a Windows system, the installation procedure creates three new services in the Services control panel.

The procedure for accessing the Services control panel varies, depending upon the version of Microsoft Windows you are using. For example, on Windows 2000, locate the Services Control panel by selecting **Settings** and then **Administrative Tools** from the **Start** menu.

Note: The `emctl` utility described in [Section 7.2.1](#) 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 7-2](#) describes the Windows services that you use to control the Management Agent.

Table 7-2 Summary of Services 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.
Oracle SNMP Peer Encapsulator	Oracle<oracle_home>SNMPPeerEncapsulator For example: OracleOraHome1PeerEncapsulator	Use this service only if you are using the advanced features of the Simple Network Management Protocol (SNMP). For more information, see the <i>Oracle SNMP Support Reference Guide</i> .
Oracle Peer SNMP Master Agent	Oracle<oracle_home>SNMPPeerMasterAgent For example: OracleOraHome1PeerMasterAgent	Use this service only if you are using the advanced features of the Simple Network Management Protocol (SNMP). For more information, see the <i>Oracle SNMP Support Reference Guide</i> .

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.

7.1.3 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_HOME directory:

```
AGENT_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 7-1).

7.2 Controlling the Oracle Management Service

The following sections describe how to control the Oracle Management Service:

- [Controlling the Management Service on UNIX](#)
- [Controlling the Management Service on Windows](#)

7.2.1 Controlling the Management Service on UNIX

To start and stop the Oracle Management Service on UNIX systems, use a set of emctl commands.

- [Using emctl to Start, Stop, and Check the Status of the Oracle Management Service](#)

7.2.1.1 Using emctl to Start, Stop, and Check the Status of the Oracle Management Service

To start, stop, or check the status of the Management Service with the Enterprise Manager command-line utility:

1. Change directory to the ORACLE_HOME/bin directory in the Management Service home.
2. Use the appropriate command described in [Table 7-3](#).

For example, to stop the Management Service, enter the following commands:

```
$PROMPT> cd bin
$PROMPT> ./emctl stop oms
```

Table 7-3 Starting, Stopping, and Checking the Status of the Management Service

Command	Purpose
emctl start oms	Starts the Fusion Middleware components required to run the Management Service J2EE application. Specifically, this command starts HTTP Server and the EMGC_OMS1 domain where the Management Service is deployed. Note: The emctl start oms command does not start Fusion Middleware. Run the startWebLogic.sh script to start WebLogic Server and its managed services.

Table 7–3 (Cont.) Starting, Stopping, and Checking the Status of the Management

Command	Purpose
emctl stop oms	Stops the Management Service. 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.

7.2.2 Controlling the Management Service on Windows

When you install the Oracle Management Service on a Windows system, the installation procedure creates three new services in the Services control panel.

The procedure for accessing the Services control panel varies, depending upon the version of Microsoft Windows you are using. For example, on Windows 2000, locate the Services control panel by selecting **Settings** and then **Administrative Tools** from the **Start** menu.

Note: The `emctl` utility described in [Section 7.2.1](#) 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 7–4](#) describes the Windows services that you use to control the Oracle Management Service.

Table 7–4 Summary of Services Installed and Configured When Installing the Oracle Management Service on Windows

Component	Service Name Format	Description
WebLogic Server	OracleWeblogicNodeManager_EMGC_OMS1_1	Use this service to start and stop the node manager of the WebLogic Server that was installed and configured to deploy the Management Service J2EE application.
Oracle Management Server	OracleManagementServer_EMGC_OMS1_1	Use this service to start and stop all components that were installed and configured as part of the Management Service J2EE application.

7.3 Controlling Fusion Middleware Control

Fusion Middleware Control is a component of Oracle Fusion Middleware 11g that is installed as part of any WebLogic Server installation. For information about starting and stopping Fusion Middleware Control, see the chapter on Starting and Stopping Oracle Fusion Middleware in the *Oracle® Fusion Middleware Administrator's Guide* available on OTN.

7.4 Controlling the Database Control on UNIX

To control the Database Control, you use the `emctl` command-line utility that is available in the `ORACLE_HOME/bin` directory after you install Oracle Database 11g.

7.4.1 Starting the Database Control on UNIX

To start the Database Control, as well the Management Agent and the Management Service associated with the Database Control:

1. Set the following environment variables to identify the Oracle home and the system identifier (SID) for the database instance you want to manage:

- `ORACLE_HOME`
- `ORACLE_SID`

2. Change directory to the `ORACLE_HOME/bin` directory.

3. Enter the following command:

```
$PROMPT> ./emctl start dbconsole
```

7.4.2 Stopping the Database Control on UNIX

To stop the Database Control, as well the Management Agent and the Management Service associated with the Database Control:

1. Set the following environment variables to identify the Oracle home and the system identifier (SID) for the database instance you want to manage:

- `ORACLE_HOME`
- `ORACLE_SID`

2. Change directory to the `ORACLE_HOME/bin` directory.

3. Enter the following command:

```
$PROMPT> ./emctl stop dbconsole
```

7.4.3 Starting and Stopping the Database Control on Windows

To start or stop the Database Control on Windows systems:

1. Open the Services control panel.

For example, on Windows NT, select **Start**, point to **Settings**, select **Control Panel**, and then double-click the Services icon.

On Windows 2000, select **Start**, point to **Administrative Tools**, and select **Services**.

2. Locate the Database Control in the list of services.

The name of the service is usually consists of "Oracle", followed by the name of the home directory you specified during the installation and the database system identifier (SID), followed by the word "DBControl." For example, if you specified `DBd11g` as the Oracle Home, the Service name would be:

```
OracleDB11gDBControl
```

3. After you locate the service, you can use the Services control panel to start or stop the Database Control service.

By default, the Database Control service is configured to start automatically when the system starts.

7.5 Guidelines for Starting Multiple Enterprise Manager Components on a Single Host

Oracle Enterprise Manager 11g 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 the 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 all the OPMN-managed processes in the WebLogic Server home directory.
2. Wait 15 seconds.
3. Run the `StartWebLogic.sh` script to start HTTP services for WebLogic Server.
4. Wait 15 seconds.
5. 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.

7.6 Starting and Stopping Oracle Enterprise Manager 11g Grid Control

As described in the previous sections, you use separate commands to control the Oracle Management Service, Oracle Management Agent, and the Oracle Fusion Middleware components on which the Grid Control depends.

The following sections describe how to stop and start all the Grid Control components that are installed by the Oracle Enterprise Manager 11g Grid 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.

7.6.1 Starting Grid Control and All Its Components

The following procedure summarizes the steps required to start all the components of the Grid Control. For example, use this procedure if you have restarted the host computer and all the components of the Grid Control have been installed on that host.

To start all the Grid 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> ORACLE_HOME/bin/emctl start oms
```

See Also: ["Controlling the Oracle Management Service"](#) on page 7-4

3. Run the StartWebLogic.sh script to start HTTP services for WebLogic Server:

4. 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 7-1

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.

7.6.2 Stopping Grid Control and All Its Components

The following procedure summarizes the steps required to stop all the components of the Grid Control. For example, use this procedure if you have installed all the components of the Grid Control on the same host you want to shut down or restart the host computer.

To stop all the Grid Control components on a host, use the following procedure:

1. Stop the Oracle Management Service:

```
$PROMPT> $ORACLE_HOME/bin/emctl stop oms
```

See Also: ["Controlling the Oracle Management Service"](#) on page 7-4

2. Run the StartWebLogic.sh script to start HTTP services for WebLogic Server.
3. 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 7-1

Note: Be sure to run the `emctl stop agent` command in the Oracle Management Agent home directory and not in the Oracle WebLogic Server home directory.

4. If your Oracle Management Repository resides on the same host, change directory to the Oracle Home for the database where you installed the Management Repository and stop the database and the Net Listener for the database:
 - a. Set the ORACLE_HOME environment variable to the Management Repository database home directory.
 - b. Set the ORACLE_SID environment variable to the Management Repository database SID (default is asdb).
 - c. Stop the database instance:

```
$PROMPT> ORACLE_HOME/bin/sqlplus /nolog
SQL> connect SYS as SYSDBA
SQL> shutdown
SQL> quit
```

See Also: *Oracle Database Administrator's Guide* for information about starting and stopping an Oracle Database

- d. Stop the Net Listener:

```
$PROMPT> $ORACLE_HOME/bin/lsnrctl stop
```

7.7 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](#)

7.7.1 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.

However, there are two Enterprise Manager commands that can help you force an immediate upload of data to the Management Service or a reload of the target definitions and attributes stored in the Management Agent home directory.

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 7-5](#).

Table 7-5 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>	This command can be used to modify 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. This command can also be used when manual edits are made to the Management Agent configuration (.XML) files. For example, if changes are made to the <code>targets.xml</code> file, which defines the attributes of your managed targets, this command will upload the modified target information to the Management Service, which will then update the information in the Management Repository. 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.

7.7.2 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 Grid 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.

You can modify the Enterprise Manager monitoring credentials by using the Oracle Enterprise Manager 11g Grid Control Console or by using the Enterprise Manager command line utility (`emctl`).

7.7.2.1 Using the Grid Control Console to Modify the Monitoring Credentials

To modify the password for the DBSNMP account in the Oracle Enterprise Manager 11g Grid Control Console:

1. Click the **Targets** tab in the Grid Control Console.
2. Click the **Database** subtab to list the database targets you are monitoring.
3. Select the database and click **Configure**.
Enterprise Manager displays the Configure Database: Properties page.
4. Enter the new password for the DBSNMP account in the **Monitor Password** field.
5. Click **Test Connection** to confirm that the monitoring credentials are correct.
6. If the connection is successful, continue to the end of the Database Configuration wizard and click **Submit**.

7.7.2.2 Using the Enterprise Manager Command Line to Modify the Monitoring Credentials

To enter new monitoring credentials with the Enterprise Manager command-line utility:

1. Change directory to the `AGENT_HOME/bin` directory (UNIX) or the `AGENT_HOME\bin` directory (Windows).
2. Enter the following command to specify new monitoring credentials:
`$PROMPT>./emctl config agent credentials [Target_name[:Target_Type]]`

To determine the correct target name and target type, see ["Listing the Targets on a Managed Host"](#).

[Example 7-2](#) shows an example of the prompts and the output you receive from the command.

Example 7-2 Modifying the Database Monitoring Credentials

```
$PROMPT>./emctl config agent credentials example.com:oracle_database
Oracle Enterprise Manager 11g Release 11.1.0.0.0
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
Name = example.us.oracle.com, Type = oracle_database
Want to change for "UserName" (y/n):n
Want to change for "password" (y/n):y
Enter the value for "password" :*****
EMD reload completed successfully
```

7.7.3 Listing the Targets on a Managed Host

There are times when you need to provide the name and type of a particular target you are managing. For example, you must know the target name and type when you are setting the monitoring credentials for a target.

To list the name and type of each target currently being monitored by a particular Management Agent:

1. Change directory to the AGENT_HOME/bin directory (UNIX) or the AGENT_HOME\bin directory (Windows).
2. Enter the following command to specify new monitoring credentials:

```
$PROMPT>./emctl config agent listtargets
```

To list targets fully, use the following command:

```
$PROMPT>./emctl config agent listtargetsfully
```

[Example 7–3](#) shows the typical output of the commands.

Example 7–3 Listing the Targets on a Managed Host

```
ade:[ example_username_1208_qc_ag ] [example_username@example emagent]$ emctl
config agent listtargets
Oracle Enterprise Manager 11g Release 1 Grid Control 11.1.0.1.0
Copyright (c) 1996, 2010 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]$
```

Example 7–4 Listing the Targets Fully on a Managed Host

```
ade:[ example_username_1208_qc_ag ] [example_username@example emagent]$ emctl
config agent listtargetsfully
Oracle Enterprise Manager 11g Release 1 Grid Control 11.1.0.1.0
Copyright (c) 1996, 2010 Oracle Corporation. All rights reserved.
<Targets>
<Target TYPE="oracle_emd" NAME="example.us.oracle.com:11852">
</Target>
<Target TYPE="host" NAME="example.us.oracle.com" DISPLAY_
NAME="example.us.oracle.com">
</Target>
<Target TYPE="oracle_webcache" NAME="chronos_test">
  <Property NAME="HTTPProtocol" VALUE="http"/>
  <Property NAME="HTTPMachine" VALUE="localhost"/>
  <Property NAME="HTTPMachineForAdmin" VALUE="example.us.oracle.com"/>
  <Property NAME="OracleHome" VALUE="/ade/example_username_1208_
qc/oracle/work/middleware/oms"/>
  <Property NAME="HTTPPath" VALUE="/"/>
  <Property NAME="MonitorPort" VALUE="4002"/>
  <Property NAME="AdminPort" VALUE="4000"/>
  <Property NAME="authpwd" VALUE="6a6aeca8b6028643" ENCRYPTED="TRUE"/>
  <Property NAME="authrealm" VALUE="sysman"/>
  <Property NAME="authuser" VALUE="sysman"/>
  <Property NAME="logFileDir" VALUE="/ade/example_username_1208_
```

```

qc/oracle/work/middleware/oms/webcache/logs"/>
  <Property NAME="logFileName" VALUE="access_log"/>
  <Property NAME="version" VALUE="10.1.2.0.0"/>
  <AssocTargetInstance ASSOCIATION_NAME="ias" ASSOC_TARGET_
NAME="/ade/example_username_1208_
qc/oracle/work/middleware/oms.example.us.oracle.com" ASSOC_TARGET_TYPE="oracle_
ias"/>
</Target>
<Target TYPE="oracle_apache" NAME="chronos_apache_test" DISPLAY_NAME="chronos_
apache_test">
  <Property NAME="HTTPMachine" VALUE="example.us.oracle.com"/>
  <Property NAME="HTTPPort" VALUE="18828"/>
  <Property NAME="useDefaultProxy" VALUE="false"/>
  <Property NAME="version" VALUE="stdApache10.1.2"/>
  <Property NAME="OracleHome" VALUE="/ade/example_username_1208_
qc/oracle/work/middleware/oms"/>

  <Property NAME="logFileDir" VALUE="/ade/example_username_1208_qc/asg_
apache/apache2/apache2/logs"/>
  <Property NAME="logFileName" VALUE="eum_log"/>
</Target>
<Target TYPE="oracle_beacon" NAME="mytestBeacon">
  <Property NAME="proxyHost" VALUE="example.us.oracle.com"/>
  <Property NAME="proxyPort" VALUE="80"/>
  <Property NAME="dontProxyFor" VALUE="us.oracle.com"/>
</Target>
<Target TYPE="oracle_csa_collector" NAME="CSAcollector">
  <Property NAME="recvFileDir" VALUE="/ade/example_username_1208_
qc/oracle/work/user_projects/domains/EMGC_DOMAIN/em/EMGC_OMS1/sysman/oms_
csa/results"/>
</Target>
<Target TYPE="oracle_database" NAME="database" VERSION="1.0">
  <Property NAME="MachineName" VALUE="example.us.oracle.com"/>
  <Property NAME="OracleHome" VALUE="/ade/example_username_1208_
qct/oracle"/>
  <Property NAME="UserName" VALUE="94d6d81ed42c38ae" ENCRYPTED="TRUE"/>
  <Property NAME="Port" VALUE="25059"/>
  <Property NAME="SID" VALUE="t1208qc"/>
  <Property NAME="Role" VALUE="NORMAL"/>
  <Property NAME="password" VALUE="94d6d81ed42c38ae" ENCRYPTED="TRUE"/>
</Target>
<Target TYPE="oracle_database" NAME="database2" VERSION="1.0">
  <Property NAME="MachineName" VALUE="example.us.oracle.com"/>
  <Property NAME="OracleHome" VALUE="/ade/example_username_1208_
qct/oracle"/>
  <Property NAME="UserName" VALUE="94d6d81ed42c38ae" ENCRYPTED="TRUE"/>
  <Property NAME="Port" VALUE="25059"/>
  <Property NAME="SID" VALUE="t1208qc"/>
  <Property NAME="Role" VALUE="NORMAL"/>
  <Property NAME="password" VALUE="94d6d81ed42c38ae" ENCRYPTED="TRUE"/>
</Target>
<Target TYPE="oracle_database" NAME="database3" VERSION="1.0">
  <Property NAME="MachineName" VALUE="example.us.oracle.com"/>
  <Property NAME="OracleHome" VALUE="/ade/example_username_1208_
qct/oracle"/>
  <Property NAME="UserName" VALUE="94d6d81ed42c38ae" ENCRYPTED="TRUE"/>
  <Property NAME="Port" VALUE="25059"/>
  <Property NAME="SID" VALUE="t1208qc"/>
  <Property NAME="Role" VALUE="NORMAL"/>
  <Property NAME="password" VALUE="94d6d81ed42c38ae" ENCRYPTED="TRUE"/>

```

```

</Target>
<Target TYPE="oracle_listener" NAME="listener">
  <Property NAME="Machine" VALUE="example.us.oracle.com"/>
  <Property NAME="Port" VALUE="25059"/>
  <Property NAME="ListenerOraDir" VALUE="/ade/example_username_1208_
qct/oracle/work"/>
  <Property NAME="LsnrName" VALUE="LISTENER"/>
  <Property NAME="OracleHome" VALUE="/ade/example_username_1208_
qct/oracle"/>
</Target>
<Target TYPE="oracle_listener" NAME="listener2">
  <Property NAME="Machine" VALUE="example.us.oracle.com"/>
  <Property NAME="Port" VALUE="25059"/>
  <Property NAME="ListenerOraDir" VALUE="/ade/example_username_1208_
qct/oracle/work"/>
  <Property NAME="LsnrName" VALUE="LISTENER2"/>
  <Property NAME="OracleHome" VALUE="/ade/example_username_1208_
qct/oracle"/>
</Target>
<Target TYPE="oracle_listener" NAME="listener3">
  <Property NAME="Machine" VALUE="example.us.oracle.com"/>
  <Property NAME="Port" VALUE="25059"/>
  <Property NAME="ListenerOraDir" VALUE="/ade/example_username_1208_
qct/oracle/work"/>
  <Property NAME="LsnrName" VALUE="LISTENER3"/>
  <Property NAME="OracleHome" VALUE="/ade/example_username_1208_
qct/oracle"/>
</Target>
<Target TYPE="oracle_emrep" NAME="Management Services and Repository">
  <Property NAME="ConnectDescriptor" VALUE="(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=stbdg06) (PORT=1521))) (CONNECT_DATA=(SERVICE_
NAME=example.us.oracle.com)))"/>
  <Property NAME="UserName" VALUE="e34e191763acd13b" ENCRYPTED="TRUE"/>
  <Property NAME="password" VALUE="e34e191763acd13b" ENCRYPTED="TRUE"/>
</Target>
</Targets>

```

7.7.4 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.

See Also: The "Systems Monitoring" chapter in Oracle Enterprise Manager Concepts for more information about Enterprise Manager blackouts

You can control blackouts from the Oracle Enterprise Manager 11g Grid 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 Grid Control Console. Similarly, if you are controlling target blackouts from the Grid Control Console, do not attempt to control those blackouts from the command line.

See Also: "Creating, Editing, and Viewing Blackouts" in the Enterprise Manager online help for information about controlling blackouts from the Grid Control Console

From the command line, you can perform the following blackout functions:

- Starting Immediate Blackouts
- Stopping Immediate Blackouts
- Checking the Status of Immediate Blackouts

Note: When you start a blackout from the command line, any Enterprise Manager jobs scheduled to run against the blacked out targets will still run. If you use the Grid 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_HOME\bin directory (Windows).
2. Enter the appropriate command as described in [Table 7-6](#).

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 "[Listing the Targets on a Managed Host](#)".

Table 7-6 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 -d 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 -nodeLevel 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 -nodeLevel option is used to specify a blackout for all the targets on the host; in other words, all the targets that the Management Agent is monitoring, including the Management Agent host itself. The -nodeLevel option must follow the blackout name. If you specify any targets after the -nodeLevel option, the list is ignored.</p>

Table 7–6 (Cont.) Summary of Blackout Commands

Blackout Action	Command
Check the status of a blackout	<code>emctl status blackout [<Target_name>[:<Target_Type>]]</code>

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
```

7.7.5 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. In the Grid Control Console, go to the All Targets page under the Targets tab, and 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. In the Related Links section of the home page, click the **Blackouts** link. 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 Grid Control Console, stop all currently running blackouts on all targets monitored by that Management Agent.
2. From the Grid Control Console, 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.

7.7.6 Reevaluating Metric Collections

If you are running a Management Agent Release 10.2, then you can 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 `$ORACLE_HOME/sysman/admin/metadata` directory, where `$ORACLE_HOME` is the Oracle Home of the Management Agent.
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 \$ORACLE_HOME/sysman/admin/default_collection directory, where \$ORACLE_HOME 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 \$ORACLE_HOME/sysman/admin/metadata directory where \$ORACLE_HOME 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 \$ORACLE_HOME/sysman/admin/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.

7.8 emctl Commands

This section lists the `emctl` commands for the Enterprise Manager Agent and Management Service.

Table 7-7 *emctl Commands*

emctl Command	Description
<code>emctl start getversion oms</code>	Gets the version of the Management Service.
<code>emctl stop oms [-all]</code>	Stops the Management Service.
<code>emctl status oms</code>	Lists the status of the Management Service
<code>emctl status oms -details</code>	Lists the status of the Management Service in detail
<code>emctl config oms sso -host ssoHost -port ssoPort -sid ssoSid -pass ssoPassword -das dasURL -u user</code>	Configures the Management Service.
<code>emctl config oms loader -shared <yes no> -dir <loader dir></code>	Configures the Management Service.
<code>emctl config oms -list_repos_details</code>	Lists the Management Service repository details.
<code>emctl config oms -store_repos_details [-repos_host <host> -repos_port <port> -repos_sid <sid> -repos_connndesc <connect descriptor>] -repos_user <username> [-repos_pwd <pwd>] [-no_check_db]</code>	Configures the Management Service.
<code>emctl config oms -change_repos_pwd [-change_in_db] [-old_pwd <old_pwd>] [-new_pwd <new_pwd>] [-use_sys_pwd [-sys_pwd <sys_pwd>]]</code>	Configures the Management Service.
<code>emctl config oms -change_view_user_pwd [-sysman_pwd <sysman_pwd>] [-user_pwd <user_pwd>] [-auto_generate]</code>	Configures the Management Service.
<code>emctl start stop agent</code>	Starts or stops agent.
<code>emctl start stop status subagent</code>	Starts or stops subagent.
<code>emctl status agent</code>	Lists the status of agent.

Table 7-7 (Cont.) emctl Commands

emctl Command	Description
emctl status agent -secure [-omsurl <http://<oms-hostname>:< oms-unsecure-port>/em/*>]	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 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 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 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 Oracle Corporation. All rights reserved. ----- step id typ pid stat command line ----- --- --- --- ----- ----- Agent is Running and Ready</pre>

Table 7-7 (Cont.) emctl Commands

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 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 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 2009-07-09 02:38:07 DeferredTrans deftrans_count n/a CLEAR 2009-07-09 02:38:07 ha_recovery missing_media_files n/a CLEAR 2009-07-09 02:28:57 ha_recovery corrupt_data_blocks n/a CLEAR 2009-07-09 02:28:57 ha_recovery datafiles_need_recovery n/a CLEAR 2009-07-09 02:28:57 Response Status n/a CLEAR 2009-07-09 02:38:04 Response userLogon n/a CLEAR 2009-07-09 02:38:04 Response State n/a CLEAR 2009-07-09 02:38:04 OCMInstrumentation NeedToInstrument n/a CLEAR 2009-07-09 02:31:55 health_check Status n/a CLEAR 2009-07-09 02:40:00 health_check Unmounted n/a CLEAR 2009-07-09 02:40:00 health_check Mounted n/a CLEAR 2009-07-09 02:40:00 health_check Unavailable n/a CLEAR 2009-07-09 02:40:00 health_check Maintenance n/a CLEAR 2009-07-09 02:40:00 sql_response time n/a CLEAR 2009-07-09 02:38:50 sga_pool_wastage java_free_pct n/a CLEAR 2009-07-09 02:28: 58 UserAudit username DBSNMP_example CLEAR 2009-07-09 02:32:48 ----- Agent is Running and Ready </pre>

Table 7-7 (Cont.) emctl Commands

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 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 Oracle Corporation. All rights reserved. ----- Metric cache contains value for following metrics at 2009-07-09 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 status agent cpu	<p>Dumps the agent Thread CPU usage into a .trc file. This file contains the list of all the threads that are running at present and their CPU usage.</p> <p>Following is the sample output of emctl status agent cpu:</p> <pre>bash-3.00\$ emctl status agent cpu Oracle Enterprise Manager 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 Oracle Corporation. All rights reserved. ----- Agent Thread CPU Snapshot available in file: /ade/example_username_cpap4_ ag/oracle/sysman/emd/cputrack/emagent_1654792_ 2009-07-09_02-58-54_cpudiag.trc ----- Agent is Running and Ready</pre>

Table 7-7 (Cont.) emctl Commands

emctl Command	Description
emctl status agent mutex	<p>Gives the detailed status of the agent mutex contention for each thread. It gives the acquired, release, and wait time of mutexes for each thread.</p> <p>Following is a sample output:</p> <pre> bash-3.00\$ emctl status agent mutex Oracle Enterprise Manager 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 Oracle Corporation. All rights reserved. ----- Mutex status at 2009-07-09 03:03:46 (prev : 2009-07-09 03:03:13) Addr Name : TotAcq TotRel TotWT LSAcq LSRel LSWT MxWt OwnerTid 2346a578 CollState : 14 14 18 0 0 0 1 NULL-thread 232e7628 CollState : 0 0 0 0 0 0 0 NULL-thread 23386638 CollState : 0 0 0 0 0 0 0 NULL-thread 2329fd28 CollState : 0 0 0 0 0 0 0 NULL-thread 239521c8 CollItem : 0 0 0 0 0 0 0 NULL-thread 2328cd98 CollState : 0 0 0 0 0 0 0 NULL-thread 23a54948 CollItem : 0 0 0 0 0 0 0 NULL-thread 232d8b68 CollItem : 1 1 1 0 0 0 1 NULL-thread 2358ce38 MetricCacheItem : 2 2 2 0 0 0 1 NULL-thread </pre>
emctl status agent memory	Used for debugging agent memory. You will need to set "enableMemoryTracing=TRUE" in emd.properties for memory profiling agent.
emctl status agent memclean	<p>Clears the memory hash table. Note that by default, the memory tracing will not be enabled by the agent. So, if the memory tracing is not enabled, then emctl status agent memclean will not clear the hashtable. To enable memory tracing, you need to set enableMemoryTracing=true in emd.properties of the agent and then reload the agent.</p> <p>The following is a sample output:</p> <pre> bash-3.00\$ emctl status agent memclean Oracle Enterprise Manager 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 Oracle Corporation. All rights reserved. ----- Memory hashtable cleared. ----- Agent is Running and Ready </pre>

Table 7-7 (Cont.) emctl Commands

emctl Command	Description
emctl reload [agent]	Reloads the agent by reading the emd.properties and targets.xml files again. If you have changed any property in emd.properties file, for example, if you have changed the tracing level of collector in the emd.properties by changing tracelevel.collector=DEBUG (default will be WARN) then you need to reload the agent to make the agent takes this change into account. Note that the agent should be up and running for the reload to happen successfully.
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 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 Oracle Corporation. All rights reserved. ----- EMD recompute dynprops completed successfully</pre>
emctl upload	Uploads xml files that are pending to upload to the OMS under the upload directory.
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 <options>	Configures agent based on the options provided.
emctl config agent updateTZ	Updates the current timezone of the agent in emd.properties file.
emctl config agent getTZ	Prints the current timezone of the agent.
emctl config agent credentials [<Target_name>[:<Target_Type>]]	<p>Provides the option to change the credentials for a particular target. Through this, you can change the user name and password of the target. It will ask you when you run this command whether you want to change the user name or password. If you select yes, then you have to provide the new user name and password for the target which you want to configure. Then it will reload the agent.</p> <p>Sample output for oracle_database is as follows:</p> <pre>bash-3.00\$ emctl config agent credentials database Oracle Enterprise Manager 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 Oracle Corporation. All rights reserved. Name = database, Type = oracle_database Want to change for "UserName" (y/n): Want to change for "password" (y/n): EMD reload completed successfully</pre>
emctl config agent getSupportedTZ	Prints the supported timezones for the agent.
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>

Table 7-7 (Cont.) emctl Commands

emctl Command	Description
emctl config [agent] addtarget [-f -force] <fileloc> [<EM loc>]	Allows you to configure agent. <fileloc> contains a definition of target to add and where -f or -force allows to overwrite an existing target. If -f or -force is not specified, existing targets cannot be overwritten. <EM loc> is optional and can be used to operate on a different Oracle Home.
emctl config [agent] addtargets [-f -force] <fileloc> [<EM loc>]	Allows you to configure agent. <fileloc> contains a definition of targets to add and -f or -force allows to overwrite existing targets. If -f or -force is not specified, existing targets cannot be overwritten. <EM loc> is optional and can be used to operate on a different Oracle Home.
emctl config [agent] modifytarget <fileloc> [<EM loc>] [<EM State>] [-mergeProps]	Allows you to configure agent. <fileloc> contains a definition of target to modify and -mergeProps is used when only target properties are to be updated. <EM loc> is optional and can be used to operate on a different Oracle Home.
emctl config [agent] deletetarget <type> <name> [<EM loc>]	Allows you to delete target. <type>,<name> specify target type and name to delete. <EM loc> is optional and can be used to operate on a different Oracle Home.
emctl config [agent] listtargets [<EM loc>]	Lists all targets present in targets.xml. <EM loc> is optional and can be used to operate on a different Oracle Home.
emctl config agent listtargetsfully [<EM loc1>] [<EM loc2>] ...	Lists all targets present in targets.xml of the given Enterprise Manager location. <EM loc> is optional and can be used to operate on a different Oracle Home.
emctl config [agent] listcentralagents [<EM loc>]	Lists the central agents this home is associated with. The centralagent command does not apply in an agent-only home. <EM loc> is optional and can be used to operate on a different Oracle Home.
emctl config [agent] addcentralagent <centralAgentHomePath> [<EM loc>]	Associates this home with a new central agent. The centralagent command does not apply in an agent-only home. <EM loc> is optional and can be used to operate on a different Oracle Home.
emctl config [agent] removecentralagent <centralAgentHomePath> [<EM loc>]	Removes the association of this home with a central agent. The centralagent command does not apply in an agent-only home. <EM loc> is optional and can be used to operate on a different Oracle Home.

Table 7-7 (Cont.) emctl Commands

emctl Command	Description
emctl config [agent] upgradecentralagent <centralAgentHomePathOld> [<centralAgentHomePathNew>]	Upgrades all product homes being monitored by this central agent. The centralagent command does not apply in an agent-only home.
emctl config [agent] setcentralagents <centralAgent1> [<centralAgent2> ...]	Sets the list of central agents this home is associated with. The centralagent command does not apply in an agent-only home.
emctl config agent addTargetsToRepository <uploadFile> <update_on_dup (true false)>	Adds targets to repository. uploadFile contains definition for targets update_on_dup decides whether updating duplicate targets This function is for central agent.
emctl config agent addAssociationsToRepository <uploadFile>	Adds associations to repository. uploadFile contains definition(s) for association(s) This function is for central agent.
emctl config agent getLocalHost	Prints the local host where the agent is running.
emctl control agent runCollection <target_name>:<target_type> <metric_name>	Allows to manually run the collections for a particular metric of a target. Sample output is as follows: bash-3.00\$ emctl control agent runCollection example.us.oracle.com:host CPUUsage Oracle Enterprise Manager 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 Oracle Corporation. All rights reserved. ----- EMD runCollection completed successfully
emctl getcurdir agent	Prints the current working directory you are in (pwd).
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 resettzhost <hostname> <override_timezone>	Resets the timezone settings of the host where the agent is running.
emctl getversion	Prints the version of the agent. Sample output is as follows: bash-3.00\$ emctl getversion bash: emctl: command not found bash-3.00\$ emctl getversion Oracle Enterprise Manager 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 Oracle Corporation. All rights reserved. --- Standalone agent Enterprise Manager 10g Agent Version 10.2.0.5.0

Table 7-7 (Cont.) emctl Commands

emctl Command	Description
emctl dumpstate agent <component> . . .	Generates the dumps for the agent. This command allow you to analyze the memory/cpu issues of the agent. Sample output is as follows: <pre>bash-3.00\$ emctl dumpstate agent Oracle Enterprise Manager 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 Oracle Corporation. All rights reserved. dump file generated: /ade/example_username_cpap4_ ag/oracle/sysman/dump/emagent_1654792_ 20090709042448.diagtrc bash-3.00\$</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	Prints the agent home directory. The sample output is as follows: <pre>bash-3.00\$ emctl getemhome Oracle Enterprise Manager 10g Release 5 Grid Control 10.2.0.5.0. Copyright (c) 1996, 2009 Oracle Corporation. All rights reserved. EMHOME=/ade/example_username_cpap4_ag/oracle</pre>
emctl start blackout <Blackoutname> [-nodeLevel] [<Target_ name>[:<Target_Type>]]... [-d <Duration>]	Starts blackout on a target. <Target_name:Target_type> defaults to local node target if not specified. If -nodeLevel is specified after <Blackoutname>,the blackout will be applied to all targets and any target list that follows will be ignored. Duration is specified in [days] hh:mm
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> [-passwd_file <abs file loc>]	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.

Table 7-7 (Cont.) emctl Commands

emctl Command	Description
emctl deploy agent [-s <install-password>] [-o <omshostname:consoleSrvPort>] [-S] <deploy-dir> <deploy-hostname>:<port> <source-hostname>	'agent' creates and deploys only the agent. [-s <password>]: Install password for securing agent. [-S]: Password will be provided in STDIN. [-o <omshostname:consoleSrvPort>]: The OMS Hostname and console servlet port. Choose the unsecured port. <deploy-dir> : Directory to create the shared (state-only) installation port. <deploy-hostname:port> : Host name and port of the shared (state-only) installation. Choose unused port. <source-hostname>: The host name of the source install. Typically the machine where EM is installed. This is searched and replaced in targets.xml by the host name provided in argument <deploy-hostname:port>. <sid>: The instance of the remote database. Only specified when deploying "dbconsole".
emctl deploy dbconsole [-s <install-password>] <deploy-dir> <deploy-hostname>:<port> <source-hostname> <sid>	'dbconsole' creates and deploys both the agent and the dbconsole. [-s <password>]: Install password for securing agent. <deploy-dir> : Directory to create the shared (state-only) installation port. <deploy-hostname:port> : Host name and port of the shared(state-only) installation. Choose unused port. <source-hostname>: The host name of the source install. Typically the machine where EM is installed. This is searched and replaced in targets.xml by the host name provided in argument <deploy-hostname:port>. <sid>: The instance of the remote database. Only specified when deploying "dbconsole".
emctl ilint	Allows ilint support of agent.
emctl annotateconfigfiles agent [<template files dir> <config files dir>]	For annotating configuration files
emctl register oms targettype [-o <SQL Output filename>] <XML filename> [<rep user> <rep passwd> <rep host> <rep port> <rep sid>] OR	For registering target type
emctl register oms targettype [-o <SQL Output filename>] <XML filename> [<rep user> <rep passwd> <rep connect descriptor>]	
emctl switchOMS <reposUrl>	For switching OMS
emctl relocate_target agent <targetname> <targettype> [<name1>=<value1>]* [-force]	Used to relocate target.

This release introduces APIs that can be run instead of emctl commands for agent configuration. [Table 7–8](#) lists the APIs.

Table 7–8 APIs for Agent Configuration

API	Description
oracle.sysman.emd:addAssociationsToRepository	Adds definition(s) for association(s). This function is for central agent. Usage is emctl config agent oracle.sysman.emd:addAssociationsToRepository <uploadFile>
oracle.sysman.emd:upgradeCentralAgentHome	Modifies cagentPathOld (if it does not already exist) to curOraHome/sysman/emd/centralagents.lst if curOraHome is not null, otherwise to ORACLE_HOME/sysman/emd/centralagents.lst where ORACLE_HOME is the java property. Each member of centralAgentPaths has to be unique. Usage is emctl config agent oracle.sysman.emd:upgradeCentralAgentHome <cagentPathOld> [cagentPathNew]
oracle.sysman.emd:printLocalHost	Prints the hostname on standard output. Usage is emctl config agent oracle.sysman.emd:printLocalHost
oracle.sysman.emd:printTargets	Prints targets of type tType present in the targets.xml file. If tType is null, all targets are printed. Usage is emctl config agent oracle.sysman.emd:printTargets [tType]
oracle.sysman.emd:printRepositoryURL	Prints the REPOSITORY_URL value from emd.properties file. Usage is emctl config agent oracle.sysman.emd:printRepositoryURL
oracle.sysman.emd:reloadTargets	Causes the agent to reload the targets.xml file. If agent is not running then this call does nothing. Usage is emctl config agent oracle.sysman.emd:reloadTargets
oracle.sysman.emd:updateAgentTimeZone	Updates the agent time zone with the value of the property agentTZRegion from emd.properties file. Usage is emctl config agent oracle.sysman.emd:updateAgentTimeZone

Table 7–8 (Cont.) APIs for Agent Configuration

API	Description
oracle.sysman.emd:setCentralAgentHomes	<p>Adds all the centralAgentPaths to curOraHome/sysman/emd/centralagents.lst if curOraHome is not null, otherwise to ORACLE_HOME/sysman/emd/centralagents.lst where ORACLE_HOME is the java property. Adding is done as follows:</p> <ol style="list-style-type: none"> 1. If the entry is already in centralagents.lst, then the command does not do anything. 2. If the entry is not in the centralagents.lst, the command: <ol style="list-style-type: none"> a. adds it to the centralagents.lst b. reads the targets.xml file of the curOraHome (or ORACLE_HOME if curOraHome is null) and tries to add all these targets to the new central agent 3. If an entry was there but is not in centralAgentPaths, the command: <ol style="list-style-type: none"> a. deletes the entry from centralagents.lst b. reads the targets.xml file of the curOraHome (or ORACLE_HOME if curOraHome is null) and deletes these targets from the old central agent entry. <p>Note - Each member of centralAgentPaths must be unique.</p>
oracle.sysman.emd:deleteSingleTargetFromCentralHome	<p>Deletes the specific target instance from homeToRemove's targets.xml. The name and type are case-sensitive.</p> <p>Usage is emctl config agent oracle.sysman.emd:deleteSingleTargetFromCentralHome <centralHome> <parentHome> <targetType> <targetName></p>
oracle.sysman.emd:modifyTarget	<p>Modifies a specified target. The composite members of this target are not affected by this operation. Usage is emctl config agent modifyTarget <target></p>
oracle.sysman.emd:printProxyHostPort	<p>Prints the REPOSITORY_URL value from emd.properties file.</p> <p>Usage is emctl config agent oracle.sysman.emd:printRepositoryURL</p>
oracle.sysman.emd:modifyTargetFromFile	<p>Modifies a target as specified in a file. The file must contain an XML fragment describing the target to be modified. The composite members of this target are not affected by this operation.</p> <p>Usage is emctl config agent oracle.sysman.emd:modifyTargetFromFile <fileLoc></p>
oracle.sysman.emd:getProxyHostPort	<p>Gets the REPOSITORY_URL value from emd.properties file.</p> <p>Usage is emctl config agent oracle.sysman.emd:getRepositoryURL</p>

Table 7–8 (Cont.) APIs for Agent Configuration

API	Description
oracle.sysman.emd:registerWith904ForCentralMonitoring	Registers with the central agent for monitoring. Usage is emctl config agent oracle.sysman.emd:registerWith904ForCentralMonitoring <centralHome> <productOH>
oracle.sysman.emd:getRepositoryURL	Gets the REPOSITORY_URL value from emd.properties file. Usage is emctl config agent oracle.sysman.emd:getRepositoryURL
oracle.sysman.emd:removeCentralAgentHome	Removes the <cagentPath> (if it exists) from curOraHome/sysman/emd/centralagents.lst if curOraHome is not null, otherwise to ORACLE_HOME/sysman/emd/centralagents.lst where ORACLE_HOME is the java property. Usage is emctl config agent oracle.sysman.emd:removeCentralAgentHome <cagentPath>
oracle.sysman.emd:printTimeZone	Prints the time zone on standard output. Usage is emctl config agent oracle.sysman.emd:printTimeZone
oracle.sysman.emd:addCentralAgentHome	Adds the cagentPath (if it does not already exist) to curOraHome/sysman/emd/centralagents.lst if curOraHome is not null, otherwise to ORACLE_HOME/sysman/emd/centralagents.lst where ORACLE_HOME is the java property. NOTE - Requires each member of centralAgentPaths to be unique. Usage is emctl config agent oracle.sysman.emd:addCentralAgentHome <cagentPath>
oracle.sysman.emd:printSupportedTimeZone	Prints the time zone on standard output. Usage is emctl config agent oracle.sysman.emd:printTimeZone
oracle.sysman.emd:listTargetsFully	Reads all targets in the specified source ORACLE_HOME - source_oracle_home(if null - in the current, local to TargetInstaller, ORACLE_HOME) and writes them to stdout in the format suitable for the agent in the destination ORACLE_HOME - destination_oracle_home. It will also show the encrypted data. Usage is emctl config agent listTargetsFully [<destination_oracle_home> <source_oracle_home destination_oracle_home>]
oracle.sysman.emd:getCentralAgentHomes	Gets the list of central agents monitoring curOraHome. Usage is emctl config agent oracle.sysman.emd:getCentralAgentHomes
oracle.sysman.emd:getEmLoc	Gets the location of the active agent given, the Enterprise Manager ORACLE_HOME. Usage is emctl config agent oracle.sysman.emd:getEmLoc
oracle.sysman.emd:relocateTargetAgent	Relocates target to the caller agent
oracle.sysman.emd:addTargetFromFile	Adds a target to the list of targets monitored by the agent. Usage is emctl config agent oracle.sysman.emd:addTargetFromFile [forceOpt ((true force -f) false)] <fileLoc>

Table 7–8 (Cont.) APIs for Agent Configuration

API	Description
oracle.sysman.emd:addFileToTargetsXml	<p>Adds targets from a given file to the list of targets monitored by the agent.</p> <p>Usage is <code>emctl config agent oracle.sysman.emd:addFileToTargetsXml [forceOpt ((true force -f) false)] <xmlFileToAdd></code></p>
oracle.sysman.emd:addTargetsToRepository	<p>Adds targets to the list of targets in the repository monitored by the agent. This function is for central agent.</p> <p>Usage is <code>emctl config agent addTargetsToRepository <uploadFile> <update_on_dup (true false)></code></p>
oracle.sysman.emd:registerForCentralMonitoring	<p>Registers with the central agent for monitoring.</p> <p>Usage is <code>emctl config agent oracle.sysman.emd:registerForCentralMonitoring <centralHome> <productOH> <stateHome></code></p>
oracle.sysman.emd:listCentralAgentHomes	<p>Lists the centralAgentPaths (if it exists) from <code>curOraHome/sysman/emd/centralagents.lst</code> if <code>curOraHome</code> is not null, otherwise to <code>ORACLE_HOME/sysman/emd/centralagents.lst</code> where <code>ORACLE_HOME</code> is the java property.</p> <p>Usage is <code>emctl config agent oracle.sysman.emd:listCentralAgentHomes</code></p>
oracle.sysman.emd:addTarget	<p>Adds a target to the list of targets monitored by the agent.</p> <p>Usage is <code>emctl config agent addTarget [force (true false)] <target></code></p>
oracle.sysman.emd:getLocalHost	<p>Returns the hostname.</p> <p>Usage is <code>emctl config agent getLocalHost</code></p>
oracle.sysman.emd:deleteTarget	<p>Deletes the specific target instance. The name and type are case-sensitive.</p> <p>Usage is <code>emctl config agent oracle.sysman.emd:deleteTarget <targetType> <targetName></code></p>
oracle.sysman.emd:unregisterForCentralMonitoring	<p>Unregisters with the central agent.</p> <p>Usage is <code>emctl config agent oracle.sysman.emd:unregisterForCentralMonitoring <productOH> <stateHome></code></p>
oracle.sysman.emd:listTargets	<p>Enumerates targets of type <code>tType</code> present in the <code>targets.xml</code> file. If <code>tType</code> is null, all targets are returned.</p> <p>Usage is <code>emctl config agent listTargets [tType] [decryptLevel hideEncryptedData]</code></p>
oracle.sysman.emd:listTargets	<p>Adds new entry to the configuration file <code>sysman/config/classpath.lst</code> that lists the full path of jar files used by the <code>JavaWrapperFetchlet</code> (<code>classpath.lst</code>). Ignores entry that already exists in the file. Does not check validity of path name supplied.</p> <p>Usage is <code>emctl config agent oracle.sysman.emd:addToClasspathLst <jarName></code></p>

7.9 Using emctl.log File

The `emctl.log` file is a file that captures the results of all `emctl` commands you run. The log file resides in the `$ORACLE_HOME/sysman/log` directory of the Management

Agent, and 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:

```
<Oracle_Home>emctl status agent
```

After running the command, navigate to the log directory to view the following information in the emctl.log file:

```
1114306 :: Wed Jun 10 02:29:36 2009::AgentLifeCycle.pm: Processing status agent
1114306 :: Wed Jun 10 02:29:36 2009::AgentStatus.pm:Processing status agent
1114306 :: Wed Jun 10 02:29:37 2009::AgentStatus.pm:emdctl status returned 3
```

Here, the first column, that is, 1114306, is the thread 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.

Another example, run the following command from the Oracle home directory of the Management Agent to upload data:

```
<Oracle_Home>emctl upload agent
```

After running the command, navigate to the log directory to view the following information in the emctl.log file:

```
1286220 :: Tue Jun 9 07:13:09 2009::AgentStatus.pm:Processing upload
1286220 :: Tue Jun 9 07:13:10 2009::AgentStatus.pm:emdctl status agent returned 3
1286220 :: Tue Jun 9 07:13:41 2009::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.

[Table 7-9](#) describes the different exit codes of an emctl command:

Table 7-9 Exit Codes

Exit Code	Description
0	Operation was successful. No action required from user.
1	Failed to connect to the Management Agent. Maybe the Management Agent is not running. Restart the Management Agent, and then try the emctl command again.
2	Timed out while connecting to the Management Agent. Maybe the Management Agent is hanging. Restart the Management Agent, and then try the emctl command again.
3	Management Agent is up and running. No action required from user.
4	Management Agent is up but not ready. Wait for some more time, and then try the emctl command again.
5	Input/output error while sending Management Agent-related request or receiving Management Agent-related response.
6	Unable to upload. Check the status of the Management Agent, and then try the emctl command again.

Table 7–9 (Cont.) Exit Codes

Exit Code	Description
7	Management Agent is in an abnormal state. Check the status of the Management Agent, and then try the emctl command again.
8	Operation was incomplete. The command might have timed out. Try again.
9	Usage error. Check the command you are running and try again with the correct command.
10	SSL handshake error while communicating with the Management Agent. Secure the Management Agent, and then try the emctl command again.
11	Key mismatch while communicating with the Management Agent. Secure the Management Agent, and then try the emctl command again.
255	Unable to open the temporary file. Try the emctl command again.

Information Publisher

Information Publisher, Enterprise Manager's powerful 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.

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](#)

8.1 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.

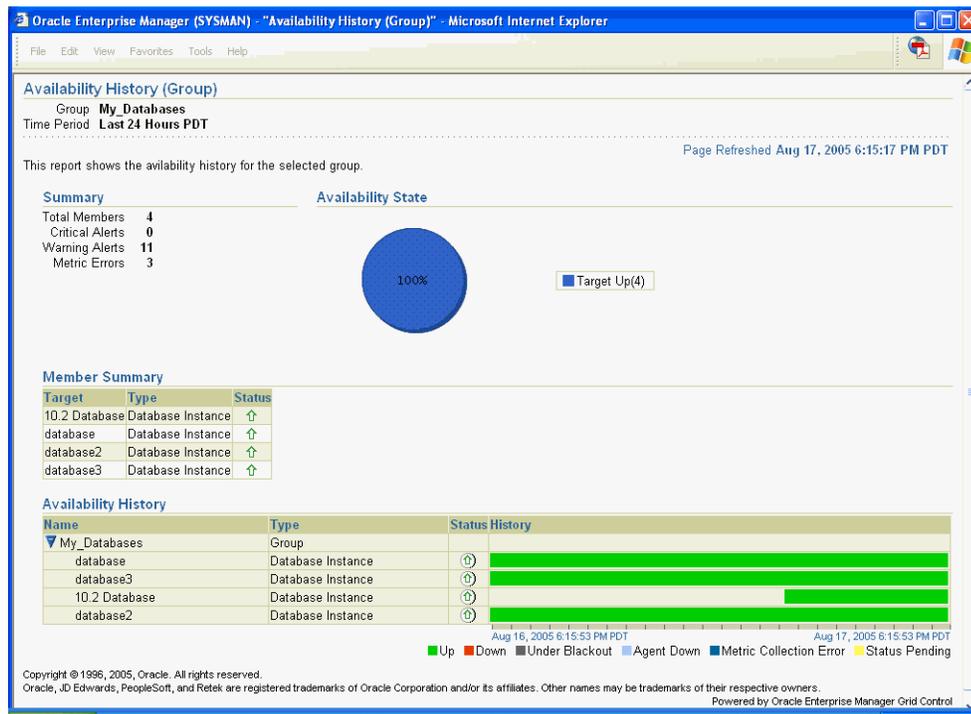
Information Publisher provides you with a feature-rich framework that is your central information source for your enterprise.

8.2 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. [Figure 8–1](#) shows an example of the Availability History (Group) report, displaying availability information for all members of a group.

Figure 8–1 Availability History (Group) Report



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. The following table lists the major functional categories and areas covered by out-of-box reports.

Table 8–1 Predefined Report Definitions

Functional Category	Areas Covered
Deployment and Configuration	Alerts and Policy Violations Application Server Configuration Client Configurations Hardware Linux Operating System Patching Operating System Oracle Database Configuration Oracle Database Software Oracle Fusion Middleware Software Oracle Home Patch Advisories Patching Automation Reports
Enterprise Manager Setup	Agent Management Pack Access
Monitoring	Aggregate Targets Alerts and Policy Violations Availability History BPEL Performance Reports Dashboards Disabled Policies Enterprise Manager Health JBoss Application Server Performance LDAP Directory Server OID Service Policy Groups Root Cause Analysis SOA Performance Reports Service Alerts Service Performance and Usage Service Tests Templates Web Application Page Performance Web Application Request Performance Web Application Transaction Performance
Security	Database Privileges <i>Note: For performance reasons, the number of records that are retrieved for any predefined report is restricted to 200. To retrieve more than 200, create a custom report. Example: Running the "Users and Roles with Access to Key Objects (Database)" report when more than 200 users and roles will be returned.</i> Database Targets Security Policy Overview VPD and OLS Policies

Table 8–1 (Cont.) Predefined Report Definitions

Functional Category	Areas Covered
Siebel Transaction	Transaction Reports
Storage	Oracle Database Space Issues Oracle Database Space Usage

8.3 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.

8.3.1 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 Create Like function.
2. Specify name, category, and sub-category. Grid 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.

8.3.2 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.

8.3.3 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 website. 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.

The following table lists the report elements that are supplied with Information Publisher.

Table 8–2 Report Elements

Report Element	Description
Generic Report Elements	Element Descriptions
Chart from SQL	Renders a line, pie or bar chart given a SQL or PL/SQL query.
Image Display	Displays a supplied image.
Separator	Displays a horizontal separator.
Styled Text	Displays text using a chosen style.
Table from SQL	Displays results of a SQL or PL/SQL query as a table.
Service Level Reporting Elements	Element Descriptions
Service Level Details	Displays Actual Service Level achieved over a time-period and violations that affected it.
Service Level Violation	Displays details on Service Level violations for a set of services over a given time-range.
Service Level Summary	Displays information on Service Levels over different time-ranges.
Services Monitoring Dashboard	Displays the Services Monitoring Dashboard showing status, performance, usage and Service-Level information for a set of Services.
Service Status Summary	Displays information on Services' Current status, Performance, Usage and Component Statuses.
Enterprise Manager Setup Element	Element Descriptions
Management Pack Access	Displays licensable targets with management pack access.
Monitoring Elements	Element Descriptions
Application Server Clusters	Displays monitoring and configuration information for Application Server Clusters.
Application Server Targets	Displays monitoring and configuration information for Application Server Targets.

Table 8–2 (Cont.) Report Elements

Report Element	Description
Availability Timeline (Group)	Displays availability of targets over a period of time. Groups, Systems, Redundancy Groups and Clusters are supported.
Metric Details	Displays a graph of a given metric for a set of targets of the same time, over a given time-period.
Open Alerts	Displays details for outstanding alerts for a user-customizable set of targets and severities.
Oracle HTTP Traffic	Displays Oracle HTTP/HTTPS Traffic information.
Service Metric Details	Displays graphs of Performance and Usage metrics for a given service.
System Monitoring Dashboard	Displays status and alert information for a Group or System.
Web Application Page Performance	Renders page performance information for a given Web Application.
Web Application Page Performance By Category	Web Application Page Performance By Category Renders page performance information by domain/region/visitor/web server.
Web Application Request Performance	Displays details about request performance of a web application.
Web Application Transaction Performance Details	Displays transaction performance details for a given transaction.
Web Application Transaction Performance Summary	Displays summary information about transaction performance.
Web Application URL Performance	Displays time series chart showing the performance for a given URL.

8.4 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.

8.4.1 Flexible Schedules

Grid 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.

8.4.2 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.

8.4.3 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.

8.5 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 website 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 website.

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.

Sizing Your Enterprise Manager Deployment

Oracle Enterprise Manager 10g Grid Control has the ability to scale for hundreds of users and thousands of systems and services on a single Enterprise Manager implementation.

This chapter describes techniques for achieving optimal performance using the Oracle Enterprise Manager application. It can also help you with capacity planning, sizing and maximizing Enterprise Manager performance in a large scale environment. By maintaining routine housekeeping and monitoring performance regularly, you insure that you will have the required data to make accurate forecasts of future sizing requirements. Receiving good baseline values for the Enterprise Manager Grid Control vital signs and setting reasonable warning and critical thresholds on baselines allows Enterprise Manager to monitor itself for you.

This chapter also provides practical approaches to backup, recovery, and disaster recovery topics while addressing different strategies when practical for each tier of Enterprise Manager.

This chapter contains the following sections:

- [Oracle Enterprise Manager Grid Control Architecture Overview](#)
- [Enterprise Manager Grid Control Sizing and Performance Methodology](#)
- [Oracle Enterprise Manager Backup, Recovery, and Disaster Recovery Considerations](#)

9.1 Oracle Enterprise Manager Grid Control Architecture Overview

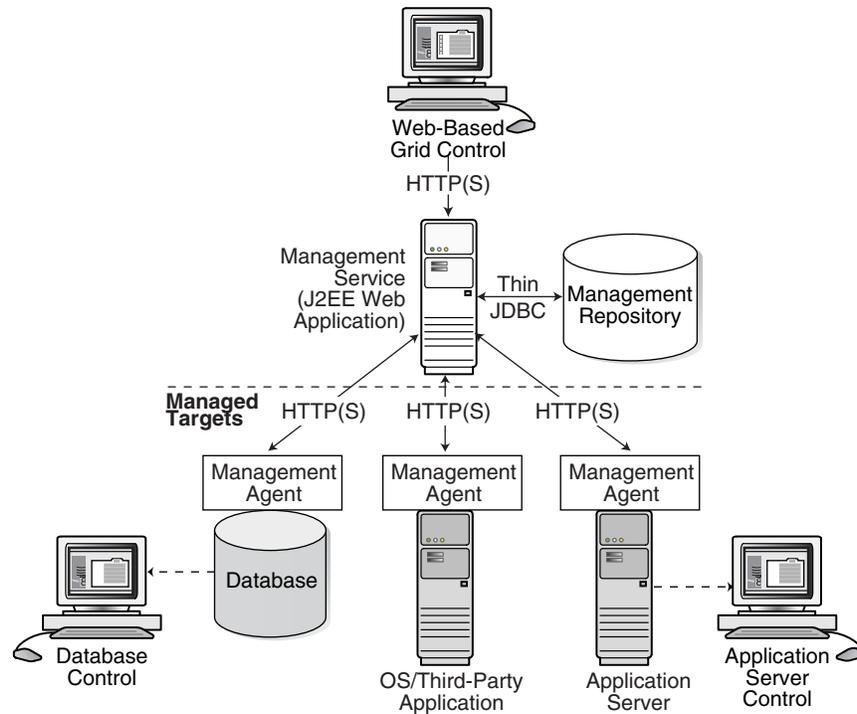
The architecture for Oracle Enterprise Manager 10g Grid Control exemplifies two key concepts in application performance tuning: distribution and parallelization of processing. Each component of Grid Control can be configured to apply both these concepts.

The components of Enterprise Manager Grid Control include:

- The Management Agent - A process that is deployed on each monitored host and that is responsible for monitoring all services and components on the host. The Management Agent is also responsible for communicating that information to the middle-tier Management Service and for managing and maintaining the system and its services.
- The Management Service - A J2EE Web application that renders the user interface for the Grid Control Console, works with all Management Agents to process monitoring and jobs information, and uses the Management Repository as its data store.

- The Management Repository - The schema is an Oracle Database that contains all available information about administrators, services, and applications managed within Enterprise Manager.

Figure 9-1 Overview of Enterprise Manager Architecture Components



For more information about the Grid Control architecture, see the Oracle Enterprise Manager 10g documentation:

- *Oracle Enterprise Manager Grid Control Installation and Basic Configuration*
- *Oracle Enterprise Manager Concepts*

The Oracle Enterprise Manager 10g documentation is available at the following location on the Oracle Technology Network (OTN):

<http://otn.oracle.com/documentation/oem.html>

9.2 Enterprise Manager Grid Control Sizing and Performance Methodology

An accurate predictor of capacity at scale is the actual metric trend information from each individual Enterprise Manager Grid Control deployment. This information, combined with an established, rough, starting host system size and iterative tuning and maintenance, produces the most effective means of predicting capacity for your Enterprise Manager Grid Control deployment. It also assists in keeping your deployment performing at an optimal level.

Here are the steps to follow to enact the Enterprise Manager Grid Control sizing methodology:

1. If you have not already installed Enterprise Manager Grid Control 10g, choose a rough starting host configuration as listed in [Table 9-1](#).

2. Periodically evaluate your site's vital signs (detailed later).
3. Eliminate bottlenecks using routine DBA/Enterprise Manager administration housekeeping.
4. Eliminate bottlenecks using tuning.
5. Extrapolate linearly into the future to plan for future sizing requirements.

Step one need only be done once for a given deployment. Steps two, three, and four must be done, regardless of whether you plan to grow your Enterprise Manager Grid Control site, for the life of the deployment on a regular basis. These steps are essential to an efficient Enterprise Manager Grid Control site regardless of its size or workload. You must complete steps two, three, and four before you continue on to step five. This is critical. Step five is only required if you intend to grow the deployment size in terms of monitored targets. However, evaluating these trends regularly can be helpful in evaluating any other changes to the deployment.

9.2.1 Step 1: Choosing a Starting Platform Grid Control Deployment

If you have not yet installed Enterprise Manager Grid Control on an initial platform, this step helps you choose a rough approximation based on experiences with real world Enterprise Manager Grid Control deployments. **If you have already installed Enterprise Manager Grid Control, proceed to Step 2.** Three typical deployment sizes are defined: small, medium, and large. The number and type of systems (or targets) it monitors largely defines the size of an Enterprise Manager Grid Control deployment.

Table 9–1 Management Server

Deployment Size	Hosts	CPUs/Hosts	Memory/Host (GB)
Small (100 monitored targets)	1	1 (3 GHz)	4
Medium (1,000 monitored targets)	1	2 (3 GHz)	Greater than or equal to 4
Large (10,000 monitored targets)	2	2 (3 GHz) 2	Greater than or equal to 6

In any OMS host box, OPMN processes, Admin Server Process, Node Manager processes, and/or DB processes will be running, so the minimum memory requirement is 4 GB per OMS host.

Table 9–2 Management Repository

Deployment Size	Hosts	CPUs/Host	Memory/Host (GB)
Small	Shares host with Management Server	Shares host with Management Server	Shares host with Management Server
Medium	1	2	4
Large	2	4	6

Table 9–3 Total Management Repository Storage

Deployment Size	Minimum Tablespace Sizes*				
	SYSTEM**	MGMT_TABLESPACE	MGMT_ECM_DEPOT_TS	MGMT_AD4J_TS	TEMP
Small	600 MB	50 GB	1 GB	100 MB	10 GB
Medium	600 MB	200 GB	4 GB	200 MB	20 GB
Large	600 MB	300 GB	Greater than 4 GB	400 MB	40 GB

*These are strictly minimum values and are intended as rough guidelines only. The actual size of the MGMT_TABLESPACE could vary widely from deployment to deployment due to variations in target type distribution, user customization, and several other factors. These tablespaces are defined with AUTOEXTEND set to ON by default to help mitigate space constraint issues. On raw file systems Oracle recommends using more than the minimum size to help prevent space constraint issues.

**The SYSTEM and TEMP tablespace sizes are minimums for Enterprise Manager only repositories. If Enterprise Manager is sharing the repository database with other application(s), these minimums may be too low.

Note: You cannot monitor tablespaces through the use of alerts with auto extended files in version 10g of Enterprise Manager. You can either set up TABLESPACE FULL alerts generate if you want to have greater control over the management of your tablespaces, or you can allow Oracle to grow your database and not alert you through the AUTOEXTEND feature. Therefore to exercise greater control of the TABLESPACE FULL alerts, you can turn off autoextend.

The previous tables show the estimated minimum hardware requirements for each deployment size. Management Servers running on more than one host, as portrayed in the large deployment above, will divide work amongst themselves.

Deploying multiple Management Servers also provides basic fail-over capabilities, with the remaining servers continuing to operate in the event of the failure of one. Use of a Server Load Balancer, or SLB, provides transparent failover for Enterprise Manager UI clients in the event of a Management Server host failure, and it also balances the request load between the available Management Servers. SLBs are host machines dedicated for load balancing purposes. SLBs can be clustered to provide fail-over capability.

Using multiple hosts for the Management Repository assumes the use of Oracle Real Application Clusters (RAC). Doing so allows the same Oracle database to be accessible on more than one host system. Beyond the storage required for the Management Server, Management Repository storage may also be required. Management Server storage is less impacted by the number of management targets. The numbers suggested in the Enterprise Manager Grid Control documentation should be sufficient in this regard.

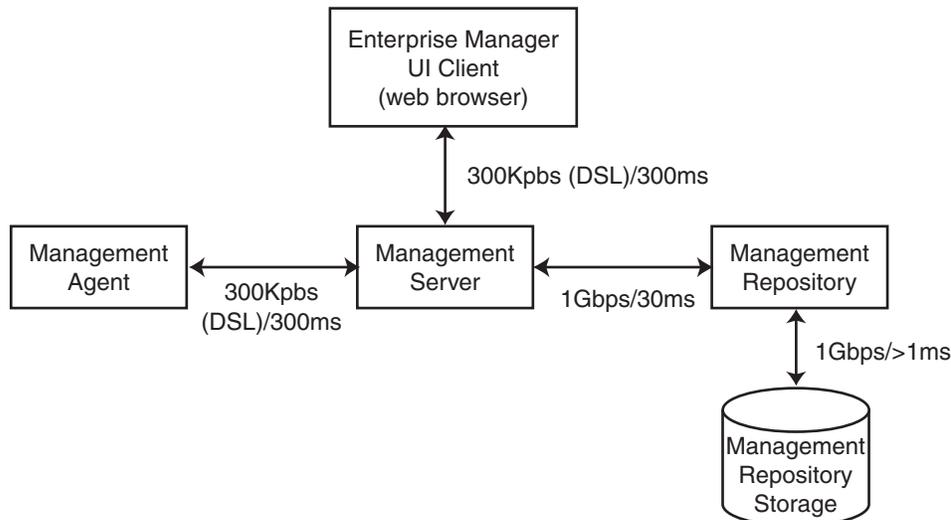
9.2.1.1 Network Topology Considerations

A critical consideration when deploying Enterprise Manager Grid Control is network performance between tiers. Enterprise Manager Grid Control ensures tolerance of network glitches, failures, and outages between application tiers through error tolerance and recovery. The Management Agent in particular is able to handle a less performant or reliable network link to the Management Service without severe impact to the performance of Enterprise Manager as a whole. The scope of the impact, as far as a single Management Agent's data being delayed due to network issues, is not likely to be noticed at the Enterprise Manager Grid Control system wide level.

The impact of slightly higher network latencies between the Management Service and Management Repository will be substantial, however. Implementations of Enterprise Manager Grid Control have experienced significant performance issues when the network link between the Management Service and Management Repository is not of

sufficient quality. The following diagram that displays the Enterprise Manager components and their connecting network link performance requirements. These are minimum requirements based on larger real world Enterprise Manager Grid Control deployments and testing.

Figure 9–2 Network Links Related to Enterprise Manager Components



You can see in [Figure 9–2](#) that the bandwidth and latency minimum requirements of network links between Enterprise Manager Grid Control components greatly impact the performance of the Enterprise Manager application.

9.2.2 Step 2: Periodically Evaluate the Vital Signs of Your Site

This is the most important step of the five. Without some degree of monitoring and understanding of trends or dramatic changes in the vital signs of your Enterprise Manager Grid Control site, you are placing site performance at serious risk. Every monitored target sends data to the Management Repository for loading and aggregation through its associated Management Agent. This adds up to a considerable volume of activity that requires the same level of management and maintenance as any other enterprise application.

Enterprise Manager has "vital signs" that reflect its health. These vital signs should be monitored for trends over time as well as against established baseline thresholds. You must establish realistic baselines for the vital signs when performance is acceptable. Once baselines are established, you can use built-in Oracle Enterprise Manager Grid Control functionality to set baseline warning and critical thresholds. This allows you to be notified automatically when something significant changes on your Enterprise Manager site. The following table is a point-in-time snapshot of the Enterprise Manager Grid Control vital signs for two sites:

Module	Metrics	EM Site 1	EM Site 2
Site URL		emsite1.acme.com	emsite2.acme.com
Target Counts	Database Targets	192 (45 not up)	1218 (634 not up)
	Host Targets	833 (12 not up)	1042 (236 not up)

Module	Metrics	EM Site 1	EM Site 2
	Total Targets	2580 (306 not up)	12293 (6668 not up)
Loader Statistics	Loader Threads	6	16
	Total Rows/Hour	1,692,000	2,736,000
	Rows/hour/load/thread	282,000	171,000
	Rows/second/load thread	475	187
	Percent of Hour Run	15	44
Rollup Statistics	Rows per Second	2,267	417
	Percent of Hour Run	5	19
Job Statistics	Job Dispatchers	2	4
	Job Steps/second/dispatcher	32	10
Notification Statistics	Notifications per Second	8	1
	Percent of Hour Run	1	13
Alert Statistics	Alerts per Hour	536	1,100
Management Service Host Statistics	Average % CPU (Host 1)	9 (emhost01)	13 (emhost01)
	Average % CPU (Host 2)	6 (emhost02)	17 (emhost02)
	Average % CPU (Host 3)	N/A	38 (em6003)
	Average % CPU (Host 4)	N/A	12 (em6004)
	Number of CPUs per host	2 X 2.8 (Xeon)	4 X 2.4 (Xeon)
	Memory per Host (GB)	6	6
Management Repository Host Statistics	Average % CPU (Host 1)	12 (db01rac)	32 (em6001rac)
	Average % CPU (Host 2)		
	Average % CPU (Host 3)		
	Average % CPU (Host 4)		
	Number of CPUs per host		
	Buffer Cache Size (MB)		
	Memory per Host (GB)	6	12
	Total Management Repository Size (GB)	56	98
	RAC Interconnect Traffic (MB/s)	1	4
	Management Server Traffic (MB/s)	4	4
	Total Management Repository I/O (MB/s)	6	27
Enterprise Manager UI Page Response/Sec	Home Page	3	6
	All Host Page	3	30+

Module	Metrics	EM Site 1	EM Site 2
	All Database Page	6	30+
	Database Home Page	2	2
	Host Home Page	2	2

The two Enterprise Manager sites are at the opposite ends of the scale for performance.

EM Site 1 is performing very well with high loader rows/sec/thread and high rollup rows/sec. It also has a very low percentage of hours run for the loader and the rollup. The CPU utilization on both the Management Server and Management Repository Server hosts are low. Most importantly, the UI Page Response times are excellent. To summarize, Site 1 is doing substantial work with minimal effort. This is how a well configured, tuned and maintained Oracle Enterprise Manager Grid Control site should look.

Conversely, EM Site 2 is having difficulty. The loader and rollup are working hard and not moving many rows. Worst of all are the UI page response times. There is clearly a bottleneck on Site 2, possibly more than one.

The following table outlines metric guidelines for the different modules based on tests that were run with the configurations outlined. It can serve as a reference point for you to extrapolate information and data based on the metrics and test environment used in the specified environment.

Table 9–4 Metric Guidelines for Modules Based On Test Environments

Module	Metrics	Value	Test Environment
Loader Statistics	Loader Threads	10	OMS Details
	Total Rows/Hour	4,270,652	# of OMS Hosts = 2
	Rows/Hour/loaderthread	427,065	# of CPU Per Host = 4 Intel Xeon Memory = 6 GB
	Rows/second/loaderthread	120	Repository Details # of Repository Nodes = 2 # of CPU per host = 4 Intel Xeon Memory = 6 GB
			EM Details Shared Recv Directory = Yes # of Agents = 867 # of Hosts = 867 Total Targets = 1803 The Metrics are collected for 5 hours after 2 OMSs were started and each agent had 50 MB of upload backlog files.
Rollup Statistics	Rows per second		
Job Statistics	Job Dispatchers	1 x Number of OMSs	
	Job Steps/second/dispatcher		

Table 9–4 (Cont.) Metric Guidelines for Modules Based On Test Environments

Module	Metrics	Value	Test Environment
Notification Statistics	Notifications per second	16	<p>OMS Details</p> <p># of OMS Hosts = 1 # of CPU Per Host = 4 Intel Xeon Memory = 6 GB</p> <p>Repository Details</p> <p># of Repository Nodes = 1 # of CPU per host = 4 Intel Xeon Memory = 6 GB</p> <p>EM Details</p> <p># of OMSs = 1 # of Repository Nodes = 1 # of Agents = 2474 # of Hosts = 2474 DB Total Targets = 8361</p>
Alert Statistics	Alerts per hour	7200	<p>OMS Details</p> <p># of OMS Hosts = 1 # of CPU Per Host = 4 Intel Xeon Memory = 6 GB</p> <p>Repository Details</p> <p># of Repository Nodes = 1 # of CPU per host = 4 Intel Xeon Memory = 6 GB</p> <p>EM Details</p> <p># of OMSs = 1 # of Repository Nodes = 1 # of Agents = 2474 # of Hosts = 2474 DB Total Targets = 8361</p>
Management Service Host Statistics	Average % CPU (Host 1)	31%	<p>OMS Details</p> <p># of OMS Hosts = 2 # of CPU Per Host = 4 Intel Xeon Memory = 6 GB</p> <p>Repository Details</p> <p># of Repository Nodes = 2 # of CPU per host = 4 Intel Xeon Memory = 6 GB</p> <p>EM Details</p> <p>Shared Recv Directory = Yes # of Agents = 867 # of Hosts = 867 Total Targets = 1803</p> <p>The Metrics are collected for 5 hours after 2 OMSs were started and each agent had 50 MB of upload backlog files.</p>
	Average % CPU (Host 2)	34%	
	Number of CPUs per host	4 (Xeon)	
	Memory per Host (GB)	6	

Table 9–4 (Cont.) Metric Guidelines for Modules Based On Test Environments

Module	Metrics	Value	Test Environment
Management Repository Host Statistics	Average % CPU (Host 1)	32%	OMS Details
	Average % CPU (Host 2)	26%	# of OMS Hosts = 2
	Number of CPUs per host	4	# of CPU Per Host = 4 Intel Xeon
	SGA Target	2 GB	Memory = 6 GB
	Memory per Host (GB)	6	Repository Details
	Total Management Repository Size (GB)	94	# of Repository Nodes = 2
	RAC Interconnect Traffic (MB/s)	1	# of CPU per host = 4 Intel Xeon
	Management Server Traffic (MB/s)		Memory = 6 GB
Enterprise Manager UI Page Response/Sec	Home Page	9.1 secs	EM Details
	All Host Page	9.8 secs	Shared Recv Directory = Yes
	All Database Page	5.7 secs	# of Agents = 867
	Database Home Page	1.7 secs	# of Hosts = 867
	Host Home Page	< 1 sec	Total Targets = 1803
			The Metrics are collected for 5 hours after 2 OMSs were started and each agent had 50 MB of upload backlog files.
			OMS Details
			# of OMS Hosts = 1
			# of CPU Per Host = 4 Intel Xeon
			Memory = 6 GB
			Repository Details
			# of Repository Nodes = 1
			# of CPU per host = 4 Intel Xeon
			Memory = 6 GB
			EM Details
			# of OMSs = 1
			# of Repository Nodes = 1
			# of Agents = 2474
			# of Hosts = 2474
			DB Total Targets = 8361

These vital signs are all available from within the Enterprise Manager interface. Most values can be found on the All Metrics page for each host, or the All Metrics page for Management Server. Keeping an eye on the trends over time for these vital signs, in addition to assigning thresholds for warning and critical alerts, allows you to maintain good performance and anticipate future resource needs. You should plan to monitor these vital signs as follows:

- Take a baseline measurement of the vital sign values seen in the previous table when the Enterprise Manager Grid Control site is running well.
- Set reasonable thresholds and notifications based on these baseline values so you can be notified automatically if they deviate substantially. This may require some iteration to fine-tune the thresholds for your site. Receiving too many notifications is not useful.
- On a daily (or weekly at a minimum) basis, watch for trends in the 7-day graphs for these values. This will not only help you spot impending trouble, but it will also allow you to plan for future resource needs.

The next step provides some guidance of what to do when the vital sign values are not within established thresholds. Also, it explains how to maintain your site's performance through routine housekeeping.

9.2.3 Step 3: Use DBA and Enterprise Manager Tasks To Eliminate Bottlenecks Through Housekeeping

It is critical to note that routine housekeeping helps keep your Enterprise Manager Grid Control site running well. The following are lists of housekeeping tasks and the interval on which they should be done.

9.2.3.1 Online Weekly Tasks

Analyze the three major tables in the Management Repository: *MGMT_METRICS_RAW*, *MGMT_METRICS_1HOUR*, and *MGMT_METRICS_1DAY*. If your Management Repository is in an Oracle 10g database, then these tables are automatically analyzed weekly and you can skip this task. If your Management Repository is in an Oracle version 9 database, then you will need to ensure that the following commands are run weekly:

- `exec dbms_stats.gather_table_stats('SYSMAN', 'MGMT_METRICS_RAW', null, .2, false, 'for all indexed columns', null, 'global', true, null, null, null);`
- `exec dbms_stats.gather_table_stats('SYSMAN', 'MGMT_METRICS_1HOUR', null, .2, false, 'for all indexed columns', null, 'global', true, null, null, null);`
- `exec dbms_stats.gather_table_stats('SYSMAN', 'MGMT_METRICS_1DAY', null, .2, false, 'for all indexed columns', null, 'global', true, null, null, null);`
- `exec dbms_stats.gather_table_stats('SYSMAN', 'MGMT_STRING_METRIC_HISTORY', null, .2, false, 'for all indexed columns', null, 'global', true, null, null, null);`

9.2.3.2 Offline Monthly Tasks

Rebuild and defragment indexes and reorganize tables as required. You may not actually need to rebuild any indexes or tables on a monthly basis. All you should do monthly is evaluate the Management Repository for tables and indexes that have grown significantly and been purged back down to a fraction of their allocated size. Unnecessarily building tables and indexes causes the Management Repository to work harder than necessary to reallocate needed space. The tables that require reorganization are easily identifiable. These tables will be large in allocated size with a relatively small number of rows, or actual size. In a real Management Repository, you may see one table that is approximately 800MB in size but only contains 6 rows. If the table is this badly oversized, it requires reorganization. Tables can be reorganized and rebuilt using commands similar to the following example:

- `exec dbms_redefinition.start_redef_table('SYSMAN','MGMT_VIOLATIONS','temp_viol');`
- `CREATE TABLE temp_viol AS
SELECT * FROM mgmt_violations WHERE 1=2;`
- `exec dbms_redefinition.start_redef_table
('SYSMAN','MGMT_VIOLATIONS','temp_viol');`

- DROP table temp_viol;

These commands rebuild the table and returns its physical structure to its clean initial state. The 800 MB table is an extreme case. Smaller disparities between actual size and row count may also indicate the need for reorganization. The Management Server(s) must be down when reorganizing a table. If you reorganize the table, the indexes must also be rebuilt. This helps make index range scans more efficient again. Indexes can be reorganized using a command similar to the following example:

- ALTER INDEX SEVERITY_PRIMARY_KEY REBUILD;

There are a few tables (along with their indexes) that may require rebuilding more frequently than others based on the higher volume of inserts and deletes they typically see. These tables are:

- MGMT_VIOLATIONS
- MGMT_CURRENT_VIOLATION
- MGMT_SYSTEM_ERROR_LOG
- MGMT_METRIC_ERRORS
- MGMT_CURRENT_METRIC_ERRORS
- MGMT_JOB_OUTPUT

These Management Repository tables can be prone to uneven growth patterns and can have large areas of allocated space that are unused due to infrequent, large spikes in data volume.

9.2.4 Step 4: Eliminate Bottlenecks Through Tuning

The most common causes of performance bottlenecks in the Enterprise Manager Grid Control application are listed below (in order of most to least common):

1. Housekeeping that is not being done (far and away the biggest source of performance problems)
2. Hardware or software that is incorrectly configured
3. Hardware resource exhaustion

When the vital signs are routinely outside of an established threshold, or are trending that way over time, you must address two areas. First, you must ensure that all previously listed housekeeping is up to date. Secondly, you must address resource utilization of the Enterprise Manager Grid Control application. The vital signs listed in the previous table reflect key points of resource utilization and throughput in Enterprise Manager Grid Control. The following sections cover some of the key vital signs along with possible options for dealing with vital signs that have crossed thresholds established from baseline values.

9.2.4.1 High CPU Utilization

When you are asked to evaluate a site for performance and notice high CPU utilization, there are a few common steps you should follow to determine what resources are being used and where.

1. The Management Server is typically a very minimal consumer of CPU. High CPU utilization in the Enterprise Manager Grid Control almost always manifests as a symptom at the Management Repository.

2. Use the Processes display on the Enterprise Manager Host home page to determine which processes are consuming the most CPU on any Management Service or Management Repository host that has crossed a CPU threshold.
3. Once you have established that Enterprise Manager is consuming the most CPU, use Enterprise Manager to identify what activity is the highest CPU consumer. Typically this manifests itself on a Management Repository host where most of the Management Service's work is performed. It is very rare that the Management Service itself is the source of the bottleneck. Here are a few typical spots to investigate when the Management Repository appears to be using too many resources.
 - a. Click on the CPU Used database resource listed on the Management Repository's Database Performance page to examine the SQL that is using the most CPU at the Management Repository.
 - b. Check the Database Locks on the Management Repository's Database Performance page looking for any contention issues.

High CPU utilization is probably the most common symptom of any performance bottleneck. Typically, the Management Repository is the biggest consumer of CPU, which is where you should focus. A properly configured and maintained Management Repository host system that is not otherwise hardware resource constrained should average roughly 40 percent or less total CPU utilization. A Management Server host system should average roughly 20 percent or less total CPU utilization. These relatively low average values should allow sufficient headroom for spikes in activity. Allowing for activity spikes helps keep your page performance more consistent over time. If your Enterprise Manager Grid Control site interface pages happen to be responding well (approximately 3 seconds) while there is no significant (constant) loader backlog, and it is using more CPU than recommended, you may not have to address it unless you are concerned it is part of a larger upward trend.

The recommended path for tracking down the root cause of high Management Repository CPU utilization is captured under step 3.b above. This allows you to start at the Management Repository Performance page and work your way down to the SQL that is consuming the most CPU in its processing. This approach has been used very successfully on several real world sites.

If you are running Enterprise Manager on Intel based hosts, the Enterprise Manager Grid Control Management Service and Management Repository will both benefit from Hyper-Threading (HT) being enabled on the host or hosts on which they are deployed. HT is a function of certain late models of Intel processors, which allows the execution of some amount of CPU instructions in parallel. This gives the appearance of double the number of CPUs physically available on the system. Testing has proven that HT provides approximately 1.5 times the CPU processing power as the same system without HT enabled. This can significantly improve system performance. The Management Service and Management Repository both frequently have more than one process executing simultaneously, so they can benefit greatly from HT.

9.2.4.2 Loader Vital Signs

The vital signs for the loader indicate exactly how much data is continuously coming into the system from all the Enterprise Manager Agents. The most important items here are the percent of hour runs and rows/second/thread. The (Loader) % of hour run indicates whether the loader threads configured are able to keep pace with the incoming data volume. As this value approaches 100%, it becomes apparent that the loading process is failing to keep pace with the incoming data volume. The lower this value, the more efficiently your loader is running and the less resources it requires

from the Management Service host. Adding more loader threads to your Management Server can help reduce the percent of hour run for the loader.

Rows/Second/Thread is a precise measure of each loader thread's throughput per second. The higher this number, the better. Rows/Second/Thread as high as 1200 have been observed on some smaller, well configured and maintained Enterprise Manager Grid Control sites. If you have not increased the number of loader threads and this number is trending down, it may indicate a problem later. One way to overcome a decreasing rows/second/thread is to add more loader threads.

The number of Loader Threads is always set to one by default in the Management Server configuration file. Each Management Server can have a maximum of 10 loader threads. Adding loader threads to a Management Server typically increases the overall host CPU utilization by 2% to 5% on a Enterprise Manager Grid Control site with many Management Agents configured. Customers can change this value as their site requires. Most medium size and smaller configurations will never need more than one loader thread. Here is a simple guideline for adding loader threads:

Max total (across all Management Servers) loader threads = 2 X number of Management Repository host CPUs

There is a diminishing return when adding loader threads. You will not yield 100% capacity from the second, or higher, thread. There should be a positive benefit, however. As you add loader threads, you should see rows/second/thread decrease, but total rows/hour throughput should increase. If you are not seeing significant improvement in total rows/hour, and there is a constantly growing loader file backlog, it may not be worth the cost of the increase in loader threads. You should explore other tuning or housekeeping opportunities in this case.

To add more loader threads, you can change the following configuration parameter where n is a positive integer [1-10]:

em.loader.threadPoolSize=n

The default is 1 and any value other than [1-10] will result in the thread pool size defaulting to 1. This property file is located in the `{ORACLE_HOME}/sysman/config` directory. Changing this parameter will require a restart of the Management Service to be reloaded with the new value.

The following two parameters are set for the Receiver module which receives files from agents.

1. *em.loader.maxDataRecvThreads=n (Default 75)*

Where n is a positive integer and default value is 75. This is used to limit the maximum number of concurrent data file receiver threads. So at the peak time only 75 receiver threads will be receiving files and an extra request will be rejected with a *Server Busy* error. These rejected requests will be resent by the agent after the default retry time.

Care should be taken while setting this value as too high a value will put an increased load on the OMS machine and shared receiver directory box. If too low a value is set then data file receive throughput will be low.

2. *oracle.sysman.emRep.dbConn.maxConnForReceiver=n (Default 25)*

Where n is a positive integer and default value is 25. This is used to set the maximum number of Repository Database connections for the receive threads. Oracle recommends you set this value equal to *em.loader.maxDataRecvThreads*, as each Receiver thread gets one DB session and there will be no wait for DB connections.

9.2.4.3 Rollup Vital Signs

The rollup process is the aggregation mechanism for Enterprise Manager Grid Control. Once an hour, it processes all the new raw data loaded into the Management Repository table `MGMT_METRICS_RAW`, calculates averages and stores them in the tables `MGMT_METRICS_1HOUR` and `MGMT_METRICS_1DAY`. The two vital signs for the rollup are the rows/second and % of hour run. Due to the large volume of data rows processed by the rollup, it tends to be the largest consumer of Management Repository buffer cache space. Because of this, the rollup vital signs can be great indicators of the benefit of increasing buffer cache size.

Rollup rows/second shows exactly how many rows are being processed, or aggregated and stored, every second. This value is usually around 2,000 (+/- 500) rows per second on a site with a decent size buffer cache and reasonable speedy I/O. A downward trend over time for this value may indicate a future problem, but as long as % of hour run is under 100 your site is probably fine.

If rollup % of hour run is trending up (or is higher than your baseline), and you have not yet set the Management Repository buffer cache to its maximum, it may be advantageous to increase the buffer cache setting. Usually, if there is going to be a benefit from increasing buffer cache, you will see an overall improvement in resource utilization and throughput on the Management Repository host. The loader statistics will appear a little better. CPU utilization on the host will be reduced and I/O will decrease. The most telling improvement will be in the rollup statistics. There should be a noticeable improvement in both rollup rows/second and % of hour run. If you do not see any improvement in any of these vital signs, you can revert the buffer cache to its previous size. The old Buffer Cache Hit Ratio metric can be misleading. It has been observed in testing that Buffer Cache Hit Ratio will appear high when the buffer cache is significantly undersized and Enterprise Manager Grid Control performance is struggling because of it. There will be times when increasing buffer cache will not help improve performance for Grid Control. This is typically due to resource constraints or contention elsewhere in the application. Consider using the steps listed in the High CPU Utilization section to identify the point of contention. Grid Control also provides advice on buffer cache sizing from the database itself. This is available on the database Memory Parameters page.

One important thing to note when considering increasing buffer cache is that there may be operating system mechanisms that can help improve Enterprise Manager Grid Control performance. One example of this is the "large memory" option available on Red Hat Linux. The Linux OS Red Hat Advanced Server™ 2.1 (RHAS) has a feature called big pages. In RHAS 2.1, `bigpages` is a boot up parameter that can be used to pre-allocate large shared memory segments. Use of this feature, in conjunction with a large Management Repository SGA, can significantly improve overall Grid Control application performance. Starting in Red Hat Enterprise Linux™ 3, big pages functionality is replaced with a new feature called huge pages, which no longer requires a boot-up parameter.

9.2.4.4 Rollup Process

The Rollup process introduces the concept of rollup participating instance; where rollup processing will be distributed among all participating instances. To add a candidate instance to the participating `EMROLLUP` group, the parameter `instance_groups` should be set on the instance level as follows:

- Add `EMROLLUP_1` to the `instance_group` parameter for node 1
Add `EMROLLUP_2` to the `instance_group` parameter for node 2
- Introduce the `PQ` and `PW` parallel processing modes where:

- PQ is the parallel query / parallel dml mode. In this mode, each participating instance will have one worker utilizing the parallel degree specified.
- PW is the parallel worker mode. In this mode, each participating instance will have a number of worker jobs equal to the parallel level specified
- Distribute the work load for all participating RAC instances as follows:
 - Each participating instance will be allocated equal number of targets. So for (n) number of participating instances with total workload (tl), each instance will be allocated (tl/n)
 - Each worker on any participating instance will be allocated equal number of targets of that instance workload. So for (il) number of targets per instance with (w) number of workers, each worker will be allocated (il/w)
 - For each worker, the load is further divided into batches to control the number of times the rollup SQL is executed. The number of rows per batch will be the total number of rows allocated for the worker divided by the number of batches.

Use the following recommendations as guidelines during the Rollup process:

- Use the parallel worker (PW) mode, and utilize the participating EMROLLUP_xx instance group.
- The recommendation is to use the parallel worker mode
- Splitting the work among more workers will improve the performance and scalability until a certain point where the diminishing returns rule will apply. This is dependent on the number of CPUs available on each RAC node. In this test case, running with 10 workers was the optimal configuration, balancing the response time, machine CPU and IO utilization
- It is important to set a proper batch size (10 recommended). The optimal run was the one with 10 batches, attributed to balancing the number of executions of the main SQL (calling EMD_1HOUR_ROLLUP) and the sort space needed for each individual execution
- Start by setting the number of batches to 10 bearing in mind the number of batches can be changed based on the data distribution

The recommendations above will yield the following results. Using the multi-instance parallel worker (8 PW) mode (with the redesigned code described earlier) improves the performance by a factor of 9-13 when utilizing two participating RAC instances.

Rollup row count (in millions) in MGMT_METRICS_1HOUR	Time (min)	Workers	Batch Size
29.5	30	8	1
9.4	5	8	10

** For the entire test there were 15779 distinct TARGET_GUID

** The test produced "29.5 Million" new rollup rows in MGMT_METRICS_1HOUR

Run **	Rows/Workers	Batches/Workers	Rows/Batch	Time (min)
8 PW /1 instance	3945	3945	1	40
8 PW /2 instances	1973	1973	1	30

9.2.4.5 Job, Notification, and Alert Vital Signs

Jobs, notifications, and alerts are indicators of the processing efficiency of the Management Service(s) on your Enterprise Manager Grid Control site. Any negative trends in these values are usually a symptom of contention elsewhere in the application. The best use of these values is to measure the benefit of running with more than one Management Server. There is one job dispatcher in each Management Server. Adding Management Servers will not always improve these values. In general, adding Management Servers will improve overall throughput for Grid Control when the application is not otherwise experiencing resource contention issues. Job, Notification, and Alert vital signs can help measure that improvement.

9.2.4.6 I/O Vital Signs

Monitoring the I/O throughput of the different channels in your Enterprise Manager Grid Control deployment is essential to ensuring good performance. At minimum, there are three different I/O channels on which you should have a baseline and alert thresholds defined:

- Disk I/O from the Management Repository instance to its data files
- Network I/O between the Management Server and Management Repository
- RAC interconnect (network) I/O (on RAC systems only)

You should understand the potential peak and sustained throughput I/O capabilities for each of these channels. Based on these and the baseline values you establish, you can derive reasonable thresholds for warning and critical alerts on them in Grid Control. You will then be notified automatically if you approach these thresholds on your site. Some Grid Control site administrators can be unaware or mistaken about what these I/O channels can handle on their sites. This can lead to Enterprise Manager Grid Control saturating these channels, which in turn cripples performance on the site. In such an unfortunate situation, you would see that many vital signs would be impacted negatively.

To discover whether the Management Repository is involved, you can use Grid Control to check the Database Performance page. On the Performance page for the Management Repository, click on the wait graph showing the largest amount of time spent. From this you can continue to drill down into the actual SQL code or sessions that are waiting. This should help you to understand where the bottleneck is originating.

Another area to check is unexpected I/O load from non-Enterprise Manager Grid Control sources like backups, another application, or a possible data-mining co-worker who engages in complex SQL queries, multiple Cartesian products, and so on.

Total Repository I/O trouble can be caused by two factors. The first is a lack of regular housekeeping. Some of the Grid Control segments can be very badly fragmented causing a severe I/O drain. Second, there can be some poorly tuned SQL statements consuming much of the site I/O bandwidth. These two main contributors can cause most of the Grid Control vital signs to plummet. In addition, the lax housekeeping can cause the Management Repository's allocated size to increase dramatically.

One important feature of which to take advantage is asynchronous I/O. Enabling asynchronous I/O can dramatically improve overall performance of the Grid Control application. The Sun Solaris™ and Linux operating systems have this capability, but may be disabled by default. The Microsoft Windows™ operating system uses asynchronous I/O by default. Oracle strongly recommends enabling of this operating system feature on the Management Repository hosts and on Management Service hosts as well.

Automatic Storage Management (ASM) is recommended for Enterprise Manager Grid Control repository database storage.

9.2.4.7 The Oracle Enterprise Manager Performance Page

There may be occasions when Enterprise Manager user interface pages are slow in the absence of any other performance degradation. The typical cause for these slow downs will be an area of Enterprise Manager housekeeping that has been overlooked. The first line of monitoring for Enterprise Manager page performance is the use of Enterprise Manager Beacons. These functionalities are also useful for web applications other than Enterprise Manager.

Beacons are designed to be lightweight page performance monitoring targets. After defining a Beacon target on an Management Agent, you can then define UI performance transactions using the Beacon. These transactions are a series of UI page hits that you will manually walk through once. Thereafter, the Beacon will automatically repeat your UI transaction on a specified interval. Each time the Beacon transaction is run, Enterprise Manager will calculate its performance and store it for historical purposes. In addition, alerts can be generated when page performance degrades below thresholds you specify.

When you configure the Enterprise Manager Beacon, you begin with a single predefined transaction that monitors the home page you specify during this process. You can then add as many transactions as are appropriate. You can also set up additional Beacons from different points on your network against the same web application to measure the impact of WAN latency on application performance. This same functionality is available for all Web applications monitored by Enterprise Manager Grid Control.

After you are alerted to a UI page that is performing poorly, you can then use the second line of page performance monitoring in Enterprise Manager Grid Control. This new end-to-end (or E2E) monitoring functionality in Grid Control is designed to allow you to break down processing time of a page into its basic parts. This will allow you to pinpoint when maintenance may be required to enhance page performance. E2E monitoring in Grid Control lets you break down both the client side processing and the server side processing of a single page hit.

The next page down in the Middle Tier Performance section will break out the processing time by tier for the page. By clicking on the largest slice of the Processing Time Breakdown pie chart, which is JDBC time above, you can get the SQL details. By clicking on the SQL statement, you break out the performance of its execution over time.

The JDBC page displays the SQL calls the system is spending most of its page time executing. This SQL call could be an individual DML statement or a PL/SQL procedure call. In the case of an individual SQL statement, you should examine the segments (tables and their indexes) accessed by the statement to determine their housekeeping (rebuild and reorg) needs. The PL/SQL procedure case is slightly more involved because you must look at the procedure's source code in the Management Repository to identify the tables and associated indexes accessed by the call.

Once you have identified the segments, you can then run the necessary rebuild and reorganization statements for them with the Management Server down. This should dramatically improve page performance. There are cases where page performance will not be helped by rebuild and reorganization alone, such as when excessive numbers of open alerts, system errors, and metric errors exist. The only way to improve these calls is to address (for example, clean up or remove) the numbers of these issues. After these numbers are reduced, then the segment rebuild and reorganization should be completed to optimize performance. These scenarios are covered in [Section 9.2.3](#). If

you stay current, you should not need to analyze UI page performance as often, if at all.

9.2.5 Step 5: Extrapolating Linearly Into the Future for Sizing Requirements

Determining future storage requirements is an excellent example of effectively using vital sign trends. You can use two built-in Grid Control charts to forecast this: the total number of targets over time and the Management Repository size over time.

Both of the graphs are available on the All Metrics page for the Management Service. It should be obvious that there is a correlation between the two graphs. A straight line applied to both curves would reveal a fairly similar growth rate. After a target is added to Enterprise Manager Grid Control for monitoring, there is a 31-day period where Management Repository growth will be seen because most of the data that will consume Management Repository space for a target requires approximately 31 days to be fully represented in the Management Repository. A small amount of growth will continue for that target for the next year because that is the longest default data retention time at the highest level of data aggregation. This should be negligible compared with the growth over the first 31 days.

When you stop adding targets, the graphs will level off in about 31 days. When the graphs level off, you should see a correlation between the number of targets added and the amount of additional space used in the Management Repository. Tracking these values from early on in your Enterprise Manager Grid Control deployment process helps you to manage your site's storage capacity proactively. This history is an invaluable tool.

The same type of correlation can be made between CPU utilization and total targets to determine those requirements. There is a more immediate leveling off of CPU utilization as targets are added. There should be no significant increase in CPU cost over time after adding the targets beyond the relatively immediate increase. Introducing new monitoring to existing targets, whether new metrics or increased collections, would most likely lead to increased CPU utilization.

9.3 Oracle Enterprise Manager Backup, Recovery, and Disaster Recovery Considerations

Enterprise Manager incorporates a portable browser-based interface to the Grid Control console, as well as the Oracle application server technology, to serve as the middle-tier Management Service tool. The foundation of the tool remains rooted in database server technology to manage both the Management Repository and historical data. This section provides practical approaches to these high availability topics and discusses different strategies when practical for each tier of Enterprise Manager.

9.3.1 Best Practices for Backup

For the Oracle database, the best backup practice is to use the standard database tools and do the following:

1. Have the database in archivelog mode
2. Perform regular online backups using the Oracle Suggested Backup strategy option available through Grid Control. This strategy uses Recovery Manager (RMAN).

This strategy creates a full backup and then creates incremental backups on each subsequent run. The incremental changes are then rolled up into the baseline, creating a new full backup baseline.

Using the Oracle Suggested Backup strategy also takes advantage of the capabilities of Grid Control to execute the backups. Backup jobs are automatically scheduled through the Grid Control Job subsystem. The history of the backups is available for review and the status of the backup displays in the Job Activity section of the database target's home page.

Use of this job along with archiving and flashback technologies provides a restore point in the event of the loss of any part of the Management Repository. This backup, along with archive and online logs, allows the Management Repository to be recovered to the last completed transaction.

To enable archiving and flashback technologies, use the Recovery Settings page and enable:

1. Archive Logging

Bounce the database and restart all Management Service processes

2. Flashback Database

Bounce the database and restart all Management Service processes

3. Block Change Tracking feature to speed up backup operations.

A summary of how to configure backups using Enterprise Manager is available in the *Oracle Database 2 Day DBA* manual.

For additional information on database high availability best practices, review the *Oracle Database High Availability Architecture and Best Practices* manual.

You can set the frequency of the backup job depending on how much data is generated in the Grid Control environment and how much outage time you can tolerate if a restore is required. If the outage window is small and the Service Level Agreement can not be satisfied by restoring the database, consider additional strategies for Management Repository availability such a Real Application Cluster (RAC) or Data Guard database. Additional High Availability options for the Management Repository are documented in the *Configuring Enterprise Manager for High Availability* paper available from the Maximum Availability Architecture (MAA) page on the Oracle Technology Network (OTN) at

<http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>

9.3.2 Best Practices for Recovery

For the Oracle Database, the best practice for recovery is to be prepared. Because in some situations the Management Repository, Management Service, and Management Agents will not have access to Grid Control, you will need to use the command-line interface to enter the RMAN commands.

9.3.2.1 Recovering the Management Repository

If something happens to affect the Management Repository, Grid Control will not be available to provide the management interface to RMAN.

A sample syntax for database recovery using RMAN follows. For detailed information, review the information on database recovery in the *Oracle Database Backup and Recovery User's Guide*.

```
RMAN> STARTUP MOUNT;
RMAN> RESTORE DATABASE;
RMAN> RECOVER DATABASE;
RMAN> ALTER DATABASE OPEN;
```

When considering recovery of the Management Repository, there are two cases to consider:

- Full recovery of the Management Repository is possible

There are no special considerations for Enterprise Manager. When the database is recovered, restart the database and Management Service processes. Management Agents will then upload pending files to the Management Repository.

- Only *point in time* and *incomplete recovery* is possible

Management Agents will be unable to communicate to the Management Repository correctly until they are reset. You must perform the following steps manually:

- a. Shut down the Management Agent
- b. Delete the `agntstmp.txt` and `lastupld.xml` files in the `$AGENT_HOME/sysman/emd` directories
- c. Go the `/state` and `/upload` subdirectories and clear their contents
- d. Restart the Management Agent.

You must repeat these steps for each Management Agent.

In the case of incomplete recovery, Management Agents may not be able to upload data until the previous steps are completed. Additionally, there is no immediate indication in the interface that the Management Agents cannot communicate with the Management Service after this type of recovery. This information would be available from the Management Agent logs or command line Management Agent status. If incomplete recovery is required, you must perform this procedure for each Management Agent.

9.3.2.2 Recovering the Oracle Management Service

Because the Management Service is stateless, the task is to restore the binaries and configuration files in the shortest time possible. There are two alternatives in this case.

- Backup the entire software directory structure. You can restore the directory structure to the same directory path should a Management Service failure occur. At the same time, backup the Management Agent associated with this Management Service install. You will need to restore this Management Agent should a restore of the Management Service be required.
- Reinstall from the original media.

For any highly available Management Service install, it is a recommended practice that you ensure that the `/recv` directory is protected with a mirroring technology. The `/recv` directory is the directory the Management Service uses to stage files it receives from Management Agents before writing their contents to the database Management Repository.

After the Management Agent finishes transmitting its XML files to the Management Service, the Management Agent deletes its copy of the XML files. Metric data sent from the Management Agents would then be lost.

9.3.2.3 Recovering the Oracle Management Agent

The recovery of the Management Agent is similar to the Management Service recovery except that the Management Agent is not stateless. There are two strategies that can be used:

- If the host name has changed, and you are using an SLB to manage connections, you have to modify the connection pools in the SLB to drop the old host name and add the new name. If you are not using an SLB, each agent that previously pointed to the old OMS host must have its `emd.properties` file modified to point to the new OMS host name. You can use this procedure to handle a case where you need to bring up a new OMS on a new host because the former machine has crashed.

Assuming the host name has not changed, a disk backup and restore is sufficient.

- a. Delete the `agntstmp.txt` and the `lastupld.xml` files from the `/sysman/emd` directory.
 - b. Clear the `/state` and `/upload` subdirectories of all entries before restarting the Management Agent.
 - c. Start the Management Agent. This step forces a rediscovery of the targets on the host.
- Reinstall the Management Agent from the original media.

As with the Management Service, it is recommended you protect the `/state` and `/upload` directories with a mirroring technology.

9.3.2.4 Preventing Data Loss and Down Time While Switching From a Non-shared File System to a Shared File System

Data loss and down time can be prevented while switching from a Non-shared File System to a Shared File System by switching each OMS to a Shared File System in rolling fashion and ensuring that there is no backlog in the receive directory. To prevent data loss, follow these steps for each OMS:

1. Shutdown the http server on the OMS.
2. Wait for the existing backlog to get processed. To determine whether the existing backlog has been processed, continue to monitor the Loader receive directory. Wait until all the files in the receive directory are uploaded.
3. Run `emctl config oms loader -shared yes -dir <sharedfs>`. If there is any backlog, this command prompts you to clear the backlog.
4. Bounce the OMS.

9.3.3 Best Practice for Disaster Recovery (DR)

In the event of a node failure, you can restore the database using RMAN or OS commands. To speed this process, implement Data Guard to replicate the Management Repository to a different hardware node.

9.3.3.1 Management Repository

If you are restoring the Management Repository to a new host, restore a backup of the database and modify the `emoms.properties` file for each Management Service manually to point to the new Management Repository location. In addition, you must update the `targets.xml` file for each Management Service to reflect the new Management Repository location. If there is a data loss during recovery, see [Recovering the Management Repository](#) for information.

To speed Management Repository reconnection from the Management Service in the event of a single Management Service failure, configure the Management Service with a Transparent Application Failover (TAF) aware connect string. You can configure the Management Service with a TAF connect string in the `emoms.properties` file that

will automatically redirect communications to another node using the `FAILOVER` syntax. An example follows:

```
EM=
(description=
(failover=on)
(address_list=
(failover=on)
(address=(protocol=tcp) (port=1522) (host=EMPRIM1.us.oracle.com))
(address=(protocol=tcp) (port=1522) (host=EMPRIM2.us.oracle.com)))
(address_list=
(failover=on)
(address=(protocol=tcp) (port=1522) (host=EMSEC1.us.oracle.com))
(address=(protocol=tcp) (port=1522) (host=EMSEC2.us.oracle.com)))
(connect_data=(service_name=EMrep.us.oracle.com)))
```

9.3.3.2 Oracle Management Service

Preinstall the Management Service and Management Agent on the hardware that will be used for Disaster Recovery. This eliminates the step of restoring a copy of the Enterprise Manager binary files from backup and modifying the Management Service and Management Agent configuration files.

Note: In the event of a disaster, do not restore the Management Service and Management Agent binaries from an existing backup to a new host because there are host name dependencies. Always do a fresh install.

9.3.3.3 Management Agent

In the event of a true disaster recovery, it is easier to reinstall the Management Agent and allow it to do a clean discovery of all targets running on the new host.

Maintaining and Troubleshooting the Management Repository

This chapter describes maintenance and troubleshooting techniques for maintaining a well-performing Management Repository.

Specifically, this chapter contains the following sections:

- [Management Repository Deployment Guidelines](#)
- [Management Repository Data Retention Policies](#)
- [Changing the SYSMAN Password](#)
- [Dropping and Recreating the Management Repository](#)
- [Troubleshooting Management Repository Creation Errors](#)
- [Cross Platform Enterprise Manager Repository Migration](#)
- [Improving the Login Performance of the Console Home Page](#)

10.1 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 Grid 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 the Data Guard environment for no data loss.

See Also: *Oracle Database High Availability Architecture and Best Practices*

Oracle Data Guard Concepts and Administration

- Oracle strongly recommends that archive logging be turned on and that a comprehensive backup strategy be in place prior to an Enterprise Manager implementation going live in a production environment. The backup strategy should include both incremental and full backups as required.

See Also: *Oracle Enterprise Manager Grid Control Installation and Basic Configuration* for information about the database initialization parameters required for the Management Repository

10.2 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 Grid 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.

10.2.1 Management Repository Default Aggregation and Purging Policies

Enterprise Manager aggregates your management data by hour and by day to minimize the size of the Management Repository. Before the data is aggregated, each data point is stored in a raw data table. Raw data is rolled up, or aggregated, into a one-hour aggregated metric table. One-hour records are then rolled up into a one-day table.

After Enterprise Manager aggregates the data, the data is then considered eligible for purging. A certain period of time has to 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.

One-hour aggregate data records are purged 31 days after they are rolled up to the one-day data table. The highest level of aggregation, one day, is kept for 365 days.

The default data retention policies are summarized in [Table 10-1](#).

Table 10-1 *Default Repository Purging Policies*

Aggregate Level	Retention Time
Raw metric data	7 days
One-hour aggregated metric data	31 days
One-day aggregated metric data	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 10–2](#).

Table 10–2 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

10.2.2 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.

10.2.3 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 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. Modifying these default policies can affect the performance of the Management Repository and have adverse reactions on the scalability of your Enterprise Manager installation.

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.

To modify the default retention time for each level of management data in the Management Repository, you must insert additional rows into the MGMT_PARAMETERS table in the Management Repository database. [Table 10–3](#) shows the parameters you must insert into the MGMT_PARAMETERS table to modify the retention time for each of the raw data and aggregate data tables.

Table names that contain "_RT_" indicate tables used for Application Performance Monitoring response time data. In the **Table Name** column, replace *datatype* with one of the three response time data types: DOMAIN, IP, or URL.

Table 10–3 Parameters for Modifying Default Data Retention Times in the Management Repository

Table Name	Parameter in MGMT_PARAMETERS Table	Default Retention Value
MGMT_METRICS_RAW	mgmt_raw_keep_window	7 days
MGMT_METRICS_1HOUR	mgmt_hour_keep_window	31 days
MGMT_METRICS_1DAY	mgmt_day_keep_window	365 days
MGMT_RT_METRICS_RAW	mgmt_rt_keep_window	24 hours
MGMT_RT_datatype_1HOUR	mgmt_rt_hour_keep_window	7 days
MGMT_RT_datatype_1DAY	mgmt_rt_day_keep_window	31 days
MGMT_RT_datatype_DIST_1HOUR	mgmt_rt_dist_hour_keep_window	24 hours
MGMT_RT_datatype_DIST_1DAY	mgmt_rt_dist_day_keep_window	31 days

Note: If the first three tables listed in Table 8-3 are not partitioned, the Default Retention Value for each is 1, 7, and 31 days respectively, rather than the 7, 31, and 365 days listed for partitioned tables.

For example, to change the default retention time for the table MGMT_METRICS_RAW from seven days to 14 days:

1. Use SQL*Plus to connect to the Management Repository database as the Management Repository user.

The default Management Repository user is `sysman`.

2. Enter the following SQL to insert the parameter and change the default value:

```
INSERT INTO MGMT_PARAMETERS (PARAMETER_NAME, PARAMETER_VALUE)
VALUES ('mgmt_raw_keep_window', '14');
```

Similarly, to change from the default retention time for all of the MGMT_RT_datatype_1DAY tables from 31 days to 100 days:

1. Use SQL*Plus to connect to the Management Repository database as the Management Repository user.

The default Management Repository user is `sysman`.

2. Enter the following SQL to insert the parameter and change the default value:

```
INSERT INTO MGMT_PARAMETERS (PARAMETER_NAME, PARAMETER_VALUE)
VALUES ('mgmt_rt_day_keep_window', '100');
```

10.2.4 Modifying Data Retention Policies When Targets Are Deleted

By default, when you delete a target from the Grid Control Console, Enterprise Manager automatically deletes all target data from the Management Repository.

However, deleting raw and aggregated metric data for database and other data-rich targets is a resource consuming operation. Targets can have hundreds of thousands of rows of data and the act of deleting this data can degrade performance of Enterprise Manager for the duration of the deletion, especially when several targets are deleted at once.

To avoid this resource-consuming operation, you can prevent Enterprise Manager from performing this task each time you delete a target. When you prevent Enterprise Manager from performing this task, the metric data for deleted targets is not purged as part of target deletion task; instead, it is purged as part of the regular purge mechanism, which is more efficient.

In addition, Oracle strongly recommends that you do not add new targets with the same name and type as the deleted targets within 24 hours of target deletion. Adding a new target with the same name and type will result in the Grid Control Console showing data belonging to the deleted target for the first 24 hours.

To disable raw metric data deletion:

1. Use SQL*Plus to connect to the Management Repository as the Management Repository user.

The default Management Repository user is SYSMAN. For example:

```
SQL> connect sysman/sysman_password;
```

2. To disable metric deletion, run the following SQL command.

```
SQL> EXEC MGMT_ADMIN.DISABLE_METRIC_DELETION();
SQL> COMMIT;
```

To enable metric deletion at a later point, run the following SQL command:

1. Use SQL*Plus to connect to the Management Repository as the Management Repository user.

The default Management Repository user is SYSMAN. For example:

```
SQL> connect sysman/oldpassword;
```

2. To enable metric deletion, run the following SQL command.

```
SQL> EXEC MGMT_ADMIN.ENABLE_METRIC_DELETION();
SQL> COMMIT;
```

10.2.5 How to Modify the Retention Period of Job History

Enterprise Manager Grid 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 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 re-create 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 timeframe 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 job runs may vary with each Enterprise Manager installation. To determine this time in your setup follow these steps:

1. Login to the Repository database using the SYSMAN account
2. Execute the following command:

```
SQL> alter session set nls_date_format='mm/dd/yy hh:mi:ss
pm';
```

```
SQL> select what, next_date from user_jobs where what like
'%JOB_ENGINE%';
```

```
WHAT
-----
NEXT_DATE
-----
MGMT_JOB_ENGINE.apply_purge_policies();
09/23/08 10:26:17 am
```

In this example, the purge policy DBMS job will run every day at 10:26:17 AM, repository time.

10.3 Changing the SYSMAN Password

The SYSMAN account is the default super user account used to set up and administer Enterprise Manager. It is also the database account that owns the objects stored in the Oracle Management Repository. From this account, you can set up additional administrator accounts and set up Enterprise Manager for use in your organization.

The SYSMAN account is created automatically in the Management Repository database during the Enterprise Manager installation. You also provide a password for the SYSMAN account during the installation.

See Also: *Oracle Enterprise Manager Grid Control Installation and Basic Configuration* for information about installing Enterprise Manager

If you later need to change the SYSMAN database account password, use the following procedure:

1. Shut down all the Oracle Management Service instances that are associated with the Management Repository.

See Also: ["Controlling the Oracle Management Service"](#) on page 7-4

2. Stop the agent that is monitoring the target OMS and Repository.

Failure to do this will result in the agent attempting to connect to the target with a wrong password once it is changed with SQL*Plus. This may also result in the SYSMAN account being locked which can subsequently prevent logins to the Grid Control console to change the password of the target OMS and Repository.

3. Change the password of the SYSMAN database account using the following SQL*Plus commands:

```
SQL>connect sysman/oldpassword;
SQL>alter user sysman identified by newpassword;
```

4. For each Management Service associated with the Management Repository, locate the `emoms.properties` configuration file.

The `emoms.properties` file can be found in the following directory of the Oracle Application Server Home where the Oracle Management Service is installed and deployed:

`IAS_HOME/sysman/config/`

5. Locate the following entries in the `emoms.properties` file:

```
oracle.sysman.eml.mntr.emdRepPwd=ece067ffc15edc4f
oracle.sysman.eml.mntr.emdRepPwdEncrypted=TRUE
```

6. Enter your new password in the first entry and enter `FALSE` in the second entry.

For example:

```
oracle.sysman.eml.mntr.emdRepPwd=new_password
oracle.sysman.eml.mntr.emdRepPwdEncrypted=FALSE
```

7. Save and exit the `emoms.properties` file and restart each Management Service associated with the Management Repository.
8. In the Grid Control console, click the **Targets** tab and then click **All Targets** on the sub tab.
9. Select the Management Services and Repository target and click **Configure**. Enterprise Manager displays the Monitoring Configurations page.
10. Enter the new password in the **Repository password** field and click **OK**.

See Also: "Specifying New Target Monitoring Credentials" on page 7-10

11. After the Management Service has started, you can check the contents of the `emoms.properties` file to be sure the password you entered has been encrypted.

For example, the entries should appear as follows:

```
oracle.sysman.eml.mntr.emdRepPwd=ece067ffc15edc4f
oracle.sysman.eml.mntr.emdRepPwdEncrypted=TRUE
```

10.3.1 Overview of the MGMT_VIEW User

During repository creation, the MGMT_VIEW user is created. This view is used by Grid Control for the reporting framework to execute queries for Table from SQL and Chart from SQL report elements. The OMS is the only entity that uses the account so there is no need to know the password. However, you can still change the password if you choose, which requires that you bounce the OMS. To change the password, you can use either a PL/SQL call or an EMCTL command:

PL/SQL:

```
SQL> exec mgmt_view_priv.change_view_user_password('<random
pwd>');
```

EMCTL command:

```
emctl config oms -change_view_user_pwd [-sysman_pwd <pwd>]
[-user_pwd <pwd>] [-autogenerate]
```

10.4 Dropping and Recreating the Management Repository

This section provides information about dropping the Management Repository from your existing database and recreating the Management Repository after you install Enterprise Manager.

10.4.1 Dropping the Management Repository

To recreate the Management Repository, you first remove the Enterprise Manager schema from your Management Repository database. You accomplish this task using the `-action drop` argument to the `RepManager` script, which is described in the following procedure.

To remove the Management Repository from your database:

1. Locate the `RepManager` script in the following directory of the Oracle Application Server Home where you have installed and deployed the Management Service:

```
IAS_HOME/sysman/admin/emdrep/bin
```

2. At the command prompt, enter the following command:

```
$PROMPT> RepManager repository_host repository_port repository_SID
-sys_password password_for_sys_account -action drop
```

In this syntax example:

- `repository_host` is the machine name where the Management Repository database is located

- `repository_port` is the Management Repository database listener port address, usually 1521 or 1526
- `repository_SID` is the Management Repository database system identifier
- `password_for_sys_account` is the password of the SYS user for the database. For example, `change_on_install`.
- `-action drop` indicates that you want to drop the Management Repository.

Alternatively, you can use a connect descriptor to identify the database on the RepManager command line. The connect descriptor identifies the host, port, and name of the database using a standard Oracle database syntax.

For example, you can use the connect descriptor as follows to create the Management Repository:

```
$PROMPT> ./RepManager -connect "(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)
(HOST=host1)(PORT=1521)) (CONNECT_DATE=(SERVICE_NAME=service)))"
-sys_password efkl34lmm -action drop
```

See Also: "Establishing a Connection and Testing the Network" in the *Oracle Database Net Services Administrator's Guide* for more information about connecting to a database using connect descriptors

10.4.2 Recreating the Management Repository

The preferred method for creating the Management Repository is to create the Management Repository during the Enterprise Manager installation procedure, which is performed using Oracle Universal Installer.

See Also: *Oracle Enterprise Manager Grid Control Installation and Basic Configuration* for information about installing Enterprise Manager

However, if you need to recreate the Management Repository in an existing database, you can use the RepManager script, which is installed when you install the Management Service. Refer to the following sections for more information:

- [Using the RepManager Script to Create the Management Repository](#)
- [Using a Connect Descriptor to Identify the Management Repository Database](#)

10.4.2.1 Using the RepManager Script to Create the Management Repository

To create a Management Repository in an existing database:

1. Review the hardware and software requirements for the Management Repository as described in *Oracle Enterprise Manager Grid Control Installation and Basic Configuration*, and review the section "[Management Repository Deployment Guidelines](#)" on page 10-1.
2. Locate the RepManager script in the following directory of the Oracle Management Service home directory:

```
ORACLE_HOME/sysman/admin/emdrep/bin
```

3. At the command prompt, enter the following command:

```
$PROMPT> ./RepManager repository_host repository_port repository_SID
-sys_password password_for_sys_account -action create
```

In this syntax example:

- *repository_host* is the machine name where the Management Repository database is located
- *repository_port* is the Management Repository database listener port address, usually 1521 or 1526
- *repository_SID* is the Management Repository database system identifier
- *password_for_sys_account* is the password of the SYS user for the database. For example, *change_on_install*.

Enterprise Manager creates the Management Repository in the database you specified in the command line.

10.4.2.2 Using a Connect Descriptor to Identify the Management Repository Database

Alternatively, you can use a connect descriptor to identify the database on the RepManager command line. The connect descriptor identifies the host, port, and name of the database using a standard Oracle database syntax.

For example, you can use the connect descriptor as follows to create the Management Repository:

```
$PROMPT> ./RepManager -connect "(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)
(HOST=host1)(PORT=1521)) (CONNECT_DATA=(SERVICE_NAME=servicename)))"
-sys_password efkl34lmn -action create
```

See Also: "Establishing a Connection and Testing the Network" in the *Oracle Database Net Services Administrator's Guide* for more information about connecting to a database using a connect descriptor

The ability to use a connect string allows you to provide an address list as part of the connection string. The following example shows how you can provide an address list consisting of two listeners as part of the RepManager command line. If a listener on one host becomes unavailable, the second listener can still accept incoming requests:

```
$PROMPT> ./RepManager -connect "(DESCRIPTION=
(ADDRESS_LIST=
(ADDRESS=(PROTOCOL=TCP) (HOST=host1) (PORT=1521)
(ADDRESS=(PROTOCOL=TCP) (HOST=host2) (PORT=1521)
(CONNECT_DATA=(SERVICE_NAME=servicename)))"
-sys_password efkl34lmn -action create
```

See Also: *Oracle Database High Availability Architecture and Best Practices*

Oracle Enterprise Manager Grid Control Installation and Basic Configuration

10.5 Troubleshooting Management Repository Creation Errors

Oracle Universal Installer creates the Management Repository using a configuration step at the end of the installation process. If the repository configuration tool fails, note the exact error messages displayed in the configuration tools window, wait until the other configuration tools have finished, exit from Universal Installer, and then use the following sections to troubleshoot the problem.

10.5.1 Package Body Does Not Exist Error While Creating the Management Repository

If the creation of your Management Repository is interrupted, you may receive the following when you attempt to create or drop the Management Repository at a later time:

```
SQL> ERROR:
ORA-00604: error occurred at recursive SQL level 1
ORA-04068: existing state of packages has been discarded
ORA-04067: not executed, package body "SYSMAN.MGMT_USER" does not exist
ORA-06508: PL/SQL: could not find program unit being called
ORA-06512: at "SYSMAN.SETEMUSERCONTEXT", line 5
ORA-06512: at "SYSMAN.CLEAR_EMCONTEXT_ON_LOGOFF", line 4
ORA-06512: at line 4
```

To fix this problem, see ["General Troubleshooting Techniques for Creating the Management Repository"](#) on page 10-11.

10.5.2 Server Connection Hung Error While Creating the Management Repository

If you receive an error such as the following when you try to connect to the Management Repository database, you are likely using an unsupported version of the Oracle Database:

```
Server Connection Hung
```

To remedy the problem, upgrade your database to the supported version as described in *Oracle Enterprise Manager Grid Control Installation and Basic Configuration*.

10.5.3 General Troubleshooting Techniques for Creating the Management Repository

If you encounter an error while creating the Management Repository, drop the repository by running the `-drop` argument to the RepManager script.

See Also: ["Dropping the Management Repository"](#) on page 10-8

If the RepManager script drops the repository successfully, try creating the Management Repository again.

If you encounter errors while dropping the Management Repository, do the following:

1. Connect to the database as SYSDBA using SQL*Plus.
2. Check to see if the SYSMAN database user exists in the Management Repository database.

For example, use the following command to see if the SYSMAN user exists:

```
prompt> SELECT username FROM DBA_USERS WHERE username='SYSMAN';
```

3. If the SYSMAN user exists, drop the user by entering the following SQL*Plus command:

```
prompt> DROP USER SYSMAN CASCADE;
```

4. Check to see if the following triggers exist:

```
SYSMAN.EMD_USER_LOGOFF
SYSMAN.EMD_USER_LOGON
```

For example, use the following command to see if the EMD_USER_LOGOFF trigger exists in the database:

```
prompt> SELECT trigger_name FROM ALL_TRIGGERS
        WHERE trigger_name='EMD_USER_LOGOFF';
```

5. If the triggers exist, drop them from the database using the following commands:

```
prompt> DROP TRIGGER SYSMAN.EMD_USER_LOGOFF;
prompt> DROP TRIGGER SYSMAN.EMD_USER_LOGON;
```

10.6 Cross Platform Enterprise Manager Repository Migration

There are user requirements for migrating an Enterprise Manager repository across servers - same and cross platforms.

The Enterprise Manager repository migration process is not exactly the same as database migration. In case of Enterprise Manager Repository migration you must take care of Enterprise Manager specific data, options, and pre-requisites for the repository move. You should make sure data integrity is maintained from both the Enterprise Manager and Oracle database perspective.

This brings up need for defining the process that can be followed by end users for successful and reliable migration of repository in minimum time and with maximum efficiency.

The overall strategy for migration depends on:

- The source and target database version
- The amount of data/size of repository
- Actual data to migrate [selective/full migration]

Cross platform transportable tablespace along with data pump (for metadata) is the fastest and best approach for moving large Enterprise Manager Grid Control repository from one platform to another. Other option that can be considered for migration is to use Data Pump for both data and metadata moves but this would require more time than the cross platform transportable tablespace approach for the same amount of data. The advantage to using the data pump approach is that it provides granular control over options and the overall process, as in the case of selective data being migrated and not the whole of source data. If the source and target is not on version 10g 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*.

10.6.1 Common Prerequisites

The following lists the common prerequisites for a repository migration:

- Source and target database must use the same character set and should be at same version.
- Source and target database should meet all the pre-requisites mentioned for Enterprise Manager Repository software requirements mentioned in Enterprise Manager install guide.
- If source and target database are NOT on 10g - only Export/Import can be used for cross platform migration

- If Source and target database are on 10g - either of three options Cross platform transportable tablespaces migration, Data Pump or Export/Import can be used for cross platform repository migration
- You cannot transport a tablespace to a target database in which a tablespace with the same name already exists. However, you can rename either the tablespace to be transported or the destination tablespace before the transport operation.
- To plug a transportable tablespace set into an Oracle Database on a different platform, both databases must have compatibility set to at least 10.0.
- Most of the platforms (but not all) are supported for cross-platform tablespace transport. You can query the V\$TRANSPORTABLE_PLATFORM view to see the platforms that are supported, and to determine their platform IDs and their endian format (byte ordering).
- Source and Destination host should have EM agent running and configured to the instance which is to be migrated
- If target database has EM repository installed, it should be first dropped using RepManager before target database related steps are carried out.

10.6.2 Methodologies

The following sections discuss the methodologies of a repository migration.

10.6.2.1 Cross Platform Transportable Tablespaces

Oracle's transportable tablespace feature allows users to quickly move a user tablespace across Oracle databases. It is the most efficient way to move bulk data between databases. Prior to Oracle Database 10g, if you want to transport a tablespace, both source and target databases need to be on the same platform. Oracle Database 10g adds the cross platform support for transportable tablespaces. With the cross platform transportable tablespace, you can transport tablespaces across platforms.

Cross platform transportable tablespaces allows a database to be migrated from one platform to another (use with Data Pump or Import/Export).

10.6.2.1.1 Preparation for Transportable Tablespaces

Use these steps to prepare for transportable tablespaces:

1. Prepare set of user tablespaces and Check for containment violation


```
execute DBMS_TTS.TRANSPORT_SET_CHECK('MGMT_TABLESPACE,MGMT_ECM_DEPOT_TS', TRUE);
select * FROM transport_set_violations;
```
2. Shutdown OMS instances and prepare for migration


```
Shutdown OMS, set job queue_processes to 0 and run
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_remove_dbms_jobs.sql
```
3. Make the tablespaces to be transported read only


```
alter tablespace MGMT_TABLESPACE read only;
alter tablespace MGMT_ECM_DEPOT_TS read only;
```

10.6.2.1.2 Extract metadata

Extract Metadata for transportable tablespaces using Data Pump Utility:

1. Create data pump directory

```
create directory data_pump_dir as '/scratch/gachawla/EM102/ttsdata';
```
2. Extract the metadata using data pump (or export)

```
expdp DUMPFILE=ttsem102.dmp TRANSPORT_TABLESPACES=MGMT_
TABLESPACE,MGMT_ECM_DEPOT_TS TRANSPORT_FULL_CHECK=Y
```
3. Extract other objects (packages, procedures, functions, temporary tables etc - Not contained in user tablespaces)

```
expdp SCHEMAS=SYSMAN CONTENT=METADATA_ONLY
EXCLUDE=INDEX,CONSTRAINT DUMPFILE=data_pump_dir:postexp.dmp
LOGFILE=data_pump_dir:postexp.log JOB_NAME=expmet
```

10.6.2.1.3 Endian check and conversion

Run Endian check and convert the datafiles if endian is different between source and destination:

1. For Endian check, run this on both source and destination database

```
SELECT endian_format
FROM v$transportable_platform tp, v$database d
WHERE tp.platform_name = d.platform_name;
```

If the source platform and the target platform are of different endianness, then an additional step must be done on either the source or target platform to convert the tablespace being transported to the target format. If they are of the same endianness, then no conversion is necessary and tablespaces can be transported as if they were on the same platform.

Example:

```
Source Endian
Linux IA (32-bit) - Little

Destination Endian
Solaris[tm] OE (32-bit) - Big
```

2. Ship datafiles, metadata dump to target and Convert datafiles using RMAN

Ship the datafiles and the metadata dump to target and On target convert all datafiles to destination endian:

```
CONVERT DATAFILE
'/d14/em10g/oradata/em102/mgmt.dbf',
'/d14/em10g/oradata/em102/mgmt_ecm_depot1.dbf'
FROM PLATFORM 'Linux IA (32-bit)';
```

Conversion via RMAN can be done either on source or target (For more details refer RMAN doc). Parallelism can be used to speed up the process if the user tablespaces contains multiple datafiles.

10.6.2.1.4 Import metadata and plugin tablespaces

Use the following steps to import metadata and plugin tablespaces:

1. Run RepManager to drop target repository (if target database has EM repository installed)

```
RepManager repository_host repository_port repository_SID -sys_password
password_for_sys_account -action drop
```

2. Run pre import steps to create sysman user and grant privs on target database


```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_create_repos_user.sql
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_pre_import.sql
```
3. Invoke Data Pump utility to plug the set of tablespaces into the target database.


```
impdp DUMPFILE=ttsem102.dmp DIRECTORY=data_pump_dir
TRANSPORT_
DATAFILES=/d14/em10g/oradata/em102/mgmt.dbf,/d14/em10g/oradata/em102/mgmt_ecm_depot1.dbf
```
4. Import other objects (packages, procedures, functions etc)


```
impdp CONTENT=METADATA_ONLY EXCLUDE=INDEX,CONSTRAINT
DUMPFILE=data_pump_dir:postexp.dmp LOGFILE=data_pump_dir:postexp.log
```

10.6.2.1.5 Post Plug In Steps

Follow these post plug in steps:

1. Run post plugin steps to recompile any invalids, create public synonyms, create other users, enable VPD policy, repin packages


```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_create_synonyms.sql
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_post_import.sql
```

Check for invalid objects - compare source and destination schemas for any discrepancy in counts and invalids.
2. Bring user tablespaces back to read write mode


```
alter tablespace MGMT_TABLESPACE read write;
alter tablespace MGMT_ECM_DEPOT_TS read write;
```
3. Submit EM dbms jobs


```
Reset back job_queue_processes to original value and run
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_submit_dbms_jobs.sql
```
4. Update OMS properties and startup OMS

Update emoms.properties to reflect the migrated repository. Update host name - oracle.sysman.eml.mntr.emdRepServer and port with the correct value and start the OMS.
5. Relocate Management Services and Repository target

If Management Services and repository target needs to be migrated to the destination host, run em_assoc.handle_relocated_target to relocate the target or recreate the target on the target host.
6. Discover/relocate Database and database Listener targets

Discover the target database and listener in EM or relocate the targets from source agent to destination agent.

10.6.2.2 Data Pump

Oracle Data Pump technology enables high-speed, parallel movement of bulk data and metadata from one database to another. Data Pump uses APIs to load and unload data instead of usual SQL commands. Data pump operations can be run via EM interface and is very useful for cross platform database migration.

The migration of the database using the Data Pump export and Data Pump import tools comprises these steps: export the data into a dump file on the source server with the expdp command; copy or move the dump file to the target server; and import the dump file into Oracle on the target server by using the impdp command; and run post import EM specific steps.

Tuning parameters that were used in original Export and Import, such as BUFFER and RECORDLENGTH, are neither required nor supported by Data Pump Export and Import

10.6.2.2.1 Prepare for Data Pump

Use the following steps to prepare for data pump:

1. Pre-requisite for using Data pump for EM repository

Impdp fails for EM repository because of data pump bug - Bug 4386766 - IMPDP WITH COMPRESSED INDEXES FAILS WITH ORA-14071 AND ORA-39083. This bug is fixed in 10.2. Backport is available for 10.1.0.4. This RDBMS patch has to be applied to use expdp/impdp for EM repository migration or workaround is to use exp/imp for extract and import.

2. Create data pump directory

Create directory data_pump_dir as '/scratch/gachawla/EM102/ttsdata';

3. Shutdown OMS instances and prepare for migration

Shutdown OMS, set job queue_processes to 0 and run @IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_remove_dbms_jobs.sql

To improve throughput of a job, PARALLEL parameter should be used to set a degree of parallelism that takes maximum advantage of current conditions. In general, the degree of parallelism should be set to more than twice the number of CPUs on an instance.

All data pump actions are performed by multiple jobs (server processes not DBMS_JOB jobs). These jobs are controlled by a master control process which uses Advanced Queuing. At runtime an advanced queue table, named after the job name, is created and used by the master control process. The table is dropped on completion of the data pump job. The job and the advanced queue can be named using the JOB_NAME parameter.

DBMS_DATAPUMP APIs can also be used to do data pump export/import. Please refer to Data pump section in 10g administration manual for all the options.

10.6.2.2.2 Data Pump Export

Use these steps to run data pump export:

1. Run data pump export:

```
expdp FULL=y DUMPFILE=data_pump_dir:dpfull1%U.dmp, data_pump_dir:dpfull2%U.dmp
PARALLEL=4 LOGFILE=data_pump_dir:dpexpfull.log JOB_NAME=dpexpfull
Verify the logs for any errors during export
```

Data pump direct path export sometimes fails for mgmt_metrics_raw and raises ORA 600. This is due to Bug 4221775 (4233303). This bug is fixed in 10.2. Workaround: if using expdp data pump for mgmt_metrics_raw , run expdp with ACCESS_METHOD+EXTERNAL_TABLE parameter.

```
expdp directory=db_export dumpfile=exp_st2.dmp logfile=exp_st2.log
tables=sysman.mgmt_metrics_raw access_method=external_table
```

10.6.2.2.3 Data Pump Import

Use these steps to run data pump import:

1. Run RepManager to drop target repository (if target database has EM repository installed)

```
RepManager repository_host repository_port repository_SID -sys_password
password_for_sys_account -action drop
```

2. Prepare the target database

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_create_
tablespaces.sql
```

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_create_
repos_user.sql
```

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_pre_
import.sql
```

3. Run data pump import

```
Impdp FULL=y DUMPFILE=data_pump_dir:dpfull1%U.dmp, data_pump_
dir:dpfull2%U.dmp PARALLEL=4 LOGFILE=data_pump_dir:dpimpfull.log JOB_
NAME=dpimpfull
```

Verify the logs for any issues with the import.

10.6.2.2.4 Post Import EM Steps

Use the following steps for post import Enterprise Manager steps:

1. Run post plugin steps to recompile any invalids, create public synonyms, create other users, enable VPD policy, repin packages

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_create_
synonyms.sql
```

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_post_
import.sql
```

Check for invalid objects - compare source and destination schemas for any discrepancy in counts and invalids.

2. Submit EM dbms jobs

Reset back job_queue_processes to original value and run

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_submit_
dbms_jobs.sql
```

3. Update OMS properties and startup OMS

Update emoms.properties to reflect the migrated repository. Update host name - oracle.sysman.eml.mntr.emdRepServer and port with the correct value and start the OMS.

4. Relocate Management Services and Repository target

If Management Services and repository target needs to be migrated to the destination host, run em_assoc.handle_relocated_target to relocate the target or recreate the target on the target host.

5. Discover/relocate Database and database Listener targets

Discover the target database and listener in EM or relocate the targets from source agent to destination agent.

10.6.2.3 Export/Import

If the source and destination database is non-10g, then export/import is the only option for cross platform database migration.

For performance improvement of export/import, set higher value for BUFFER and RECORDLENGTH . Do not export to NFS as it will slow down the process considerably. Direct path can be used to increase performance. Note - As EM uses VPD, conventional mode will only be used by Oracle on tables where policy is defined.

Also User running export should have EXEMPT ACCESS POLICY privilege to export all rows as that user is then exempt from VPD policy enforcement. SYS is always exempted from VPD or Oracle Label Security policy enforcement, regardless of the export mode, application, or utility that is used to extract data from the database.

10.6.2.3.1 Prepare for Export/Import

Use the following steps to prepare for Export/Import:

1. Mgmt_metrics_raw partitions check

```
select table_name,partitioning_type type,
partition_count, subpartitioning_type subtype from
dba_part_tables where table_name = 'MGMT_METRICS_RAW'
```

If MGMT_METRICS_RAW has more than 3276 partitions please see Bug 4376351 - This is Fixed in 10.2 . Workaround is to export mgmt_metrics_raw in conventional mode.

2. Shutdown OMS instances and prepare for migration

```
Shutdown OMS, set job_queue_processes to 0 and run @IAS_
HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_remove_dbms_
jobs.sql
```

10.6.2.3.2 Export

Follow these steps for export:

1. Export data

```
exp full=y constraints=n indexes=n compress=y file=fullem102_1.dmp
log=fullem102exp_1.log
```

2. Export without data and with constraints

```
exp full=y constraints=y indexes=y rows=n ignore=y file=fullem102_2.dmp
log=fullem102exp_2.log
```

10.6.2.3.3 Import

Follow these steps to import:

1. Run RepManager to drop target repository (if target database has EM repository installed)

```
RepManager repository_host repository_port repository_SID -sys_password
password_for_sys_account -action drop
```

2. Pre-create the tablespaces and the users in target database

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_create_
tablespaces.sql
```

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_create_
repos_user.sql
```

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_pre_
import.sql
```

3. Import data

```
imp full=y constraints=n indexes=n file=fullem102_1.dmp log=fullem102imp_
1.log
```

4. Import without data and with constraints

```
imp full=y constraints=y indexes=y rows=n ignore=y file=fullem102_2.dmp
log=fullem102imp_2.log
```

10.6.2.3.4 Post Import EM Steps

Follow these steps for post import EM steps:

1. Run post plugin steps to recompile any invalids, create public synonyms, create other users, enable VPD policy, repin packages

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_create_
synonyms.sql
```

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_post_
import.sql
```

Check for invalid objects - compare source and destination schemas for any discrepancy in counts and invalids.

2. Submit EM dbms jobs

Reset back job_queue_processes to original value and run

```
@IAS_HOME/sysman/admin/emdrep/sql/core/latest/admin/admin_submit_
dbms_jobs.sql
```

3. Update OMS properties and startup OMS

Update emoms.properties to reflect the migrated repository. Update host name, oracle.sysman.eml.mntr.emdRepServer and port with the correct value and start the OMS.

4. Relocate Management Services and Repository target

If Management Services and repository target needs to be migrated to the destination host, run em_assoc.handle_relocated_target to relocate the target or recreate the target on the target host.

5. Discover/relocate Database and database Listener targets

Discover the target database and listener in EM or relocate the targets from source agent to destination agent.

10.6.3 Post Migration Verification

These verification steps should be carried out post migration to ensure that the migration was completely successful:

- Verify any discrepancy in objects by comparing source and target databases through EM
- Verify migrated database through EM whether database is running without any issues
- Verify repository operations, dbms jobs and whether any management system errors reported
- Verify all EM functionalities are working fine after migration
- Make sure Management Services and Repository target is properly relocated by verifying through EM

10.7 Improving the Login Performance of the Console Home Page

Oracle Enterprise Manager now provides an option that will more quickly display the Console Home page even in a scenario where the Management Repository is very large. Normally, factors such as the number of alerts, errors, policies, and critical patches can contribute to delayed displayed times. Since there is no single factor nor any simple way to scale the SQL or user interface, a simple option flag has been added that removes the following page elements for all users.

When the `emoms.properties` flag, `LargeRepository=`, is set to true (when normally the default is false), the SQL for the following items is not executed and thus the items will not be displayed on the Console page.

1. Three sections from the Overview Page segment:
 - All Target Alerts
 - Critical
 - Warning
 - Errors
 - All Target Policy Violations
 - Critical
 - Warning
 - Informational
 - All Target Jobs
 - Problem Executions (last 7 days)
 - Suspended Executions (last 7 days)
2. The page segment which includes Security Patch Violations and Critical Patch Advisories.

The Deployment Summary section would move up to fill in the vacated space.

Part II

Advanced Topics

This section of the guide discusses the application of Enterprise Manager functionality to monitor and maintain your entire ecosystem.

Part II contains the following chapters:

- [Chapter 11, "Setting Up the JVM Diagnostics Manager"](#)
- [Chapter 12, "Setting Up Request Monitoring"](#)
- [Chapter 13, "Managing Compliance"](#)
- [Chapter 14, "Configuring Services"](#)
- [Chapter 15, "Extending Enterprise Manager"](#)

Setting Up the JVM Diagnostics Manager

This chapter describes the process of installing and setting the JVM Diagnostics Manager. The `DeployAD4Jmanager.sh` script is used to deploy the JVM Diagnostics Manager. The script does the following:

- Create (clones) the managed server from the existing OMS server.
- Deploys the `jamanager.ear` file on the new managed server.

11.1 Deploying the JVM Diagnostics Manager

This section describes the following:

- [Prerequisites](#)
- [Running the DeployAD4Jmanager.sh Script](#)
- [Creating a Special Repository User for Loading JVM Diagnostics Heaps](#)

11.1.1 Prerequisites

- The `DeployAD4Jmanager.sh` script must be started from the `$ORACLE_HOME/oms11g/ad4j` directory.
- The port number specified by user to create a new managed server should not be in use.
- You must have appropriate login credentials to access and execute the `DeployAD4Jmanager.sh` script.

11.1.2 Running the DeployAD4Jmanager.sh Script

To install the JVM Diagnostics Manager, follow the steps given below:

1. Run the `DeployAD4Jmanager.sh` script from the command prompt:

```
$ ./DeployAD4JManager.sh
```

You will see the following:

```
Welcome to AD4J Manager Installation script
Please follow the guided instructions to install AD4J manager
Setting up environment variables
Environment variable used in the script:
ORACLE_HOME=/u01/atejaswy/Oracle/Middleware/oms11g
MW_HOME=/u01/atejaswy/Oracle/Middleware
WL_HOME=/u01/atejaswy/Oracle/Middleware/wlserver_10.3
MODULES_DIR=/u01/atejaswy/Oracle/Middleware/modules
Delete previous existing ad4j directory..
```

```
Done
Unzipping ad4j.zip.....
Archive:  ad4j.zip
creating: ad4j/
inflating: ad4j/ad4jjobtype.xml
inflating: ad4j/jamagent.war
inflating: ad4j/jamagent.ear
```

2. You are now prompted for the WebLogic Server host name. Press **Enter** to select the default values. The following prompts are displayed:

Enter weblogic host name:

Enter the host name or the IP address of the machine on which the WebLogic server is running.

Is secured connection (uses https)? [Default:[y]es]

Press **Enter** to indicate that it is a secure connection.

Enter weblogic ssl port number

Enter the ssl(https) port number if user selects https, otherwise its http port number.

Enter weblogic admin username

Enter the name of the admin user.

Enter weblogic admin password

Enter the password for the admin user.

3. The WebLogic server details that you have selected are displayed. You are then prompted for the details of the new managed server. Press **Enter** to accept the default values. The following prompts are displayed:

Enter managed server name:

Enter the name of the new managed server. The default name is EMAD4JMANAGER.

Enter managed server machine name:

Enter the name of machine on which managed server is to be created. The default is EMGC_MACHINE1.

Enter managed server listen address

Enter the listen address of the managed server. The default is the name of the current machine.

Enter managed server listen port

Enter the user configured port number for the managed server. The default value is 3800.

Enter managed server SSL listen port

Enter the user configured ssl port number for the managed server. The default value is 3801.

Enter existing OMS server name

Enter the name of the OMS server.

Enter existing EMGC_DOMAIN path

Enter the location of the EM domain.

Enter existing INSTANCE_HOME path

Enter the directory in which the OMS server has been installed.

Do you want secure communication between AD4J Agent and AD4J Manager (y/n)
[Default:y]

Enter **Y** if secure communication should be supported.

4. If you enter Y, you are prompted for the following:

Enter the absolute location of the wallet

Enter the complete path of the wallet including the wallet name.

Enter wallet password

Enter the password for the wallet. This will ensure that there is secure communication between the AD4J Agent and the AD4J Manager.

5. You are prompted to continue with the installation as follows:

Do you wish to continue with the installation of AD4J Manager (y/n) [Default:y]

Press **Enter** to continue. The following details are displayed:

Initializing WebLogic Scripting Tool (WLST) ...

Welcome to WebLogic Server Administration Scripting Shell

Type help() for help on available commands

Deployment Summary:

WEBLOGICHOSTNAME dadvmi0105

WEBLOGICPORTNUMBER: 7101

USERNAME: weblogic

MACHINENAME: EMGC_MACHINE1

ORACLEHOME: /u01/atejaswy/Oracle/Middleware/oms11g

WLHOME: /u01/atejaswy/Oracle/Middleware/wlserver_10.3

MODULEHOME: /u01/atejaswy/Oracle/Middleware/modules

SERVERNAME ad4jmanager

LISTENADDRESS dadvmi0105

LISTENPORT 3800

SSL LISTENPORT 3801

Connecting to t3s://dadvmi0105:7101 with userid weblogic ...

Successfully connected to Admin Server 'EMGC_ADMINSERVER' that belongs to domain 'EMGC_DOMAIN'.

Location changed to edit tree. This is a writable tree with DomainMBean as the root. To make changes you will need to start an edit session via startEdit().

For more help, use help(edit)

You already have an edit session in progress and hence WLST will continue with your edit session.

Starting an edit session ...

Started edit session, please be sure to save and activate your changes once you are done.

```
Saving all your changes ...
Saved all your changes successfully.
Activating all your changes, this may take a while ...
The edit lock associated with this edit session is released
once the activation is completed.
Activation completed

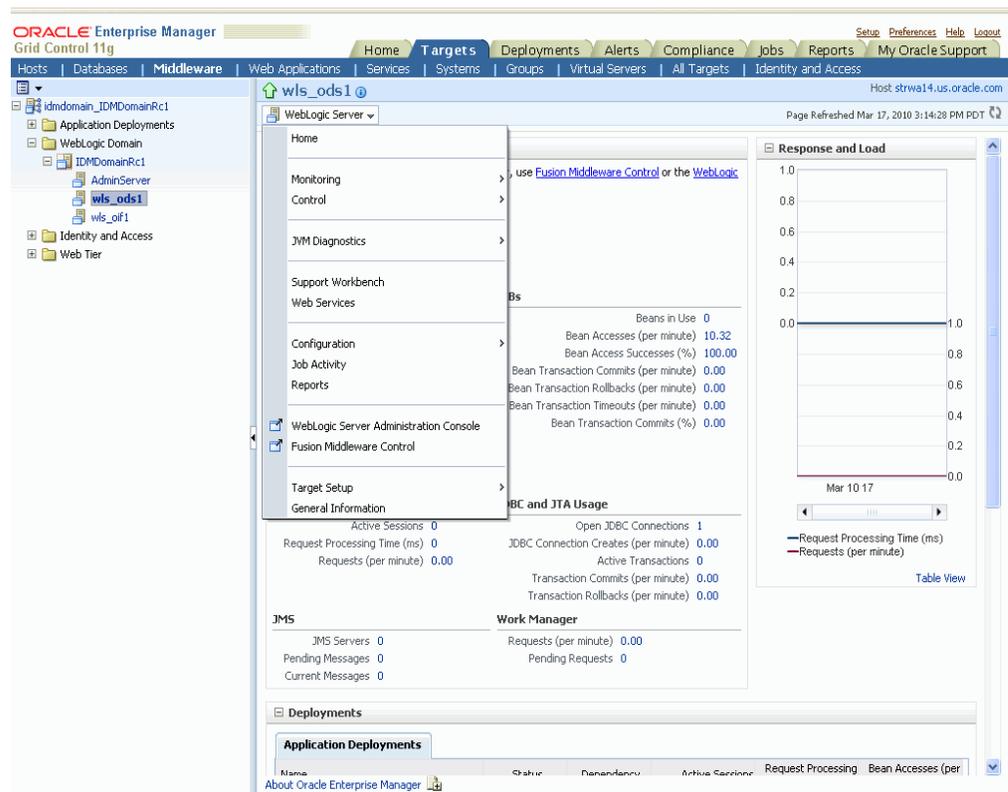
Starting server ad4jmanager
.....
.....
Server with name ad4jmanager started successfully
Deploying application from
/u01/atejaswy/Oracle/Middleware/oms11g/ad4j/./ad4j/jamanager.ear to targets
ad4jmanager (upload=false) ...
<Jan 8, 2010 6:33:16 AM PST> <Info> <J2EE Deployment SPI> <BEA-260121>
<Initiating deploy operation for application, ad4jmanagerapp [archive:
/u01/atejaswy/Oracle/Middleware/oms11g/ad4j/./ad4j/jamanager.ear], to
ad4jmanager .>
.....Completed the deployment of Application with status
completed
Current Status of your Deployment:
Deployment command type: deploy
Deployment State      : completed
Deployment Message    : no message

Exiting WebLogic Scripting Tool.
```

The JVM Diagnostics Manager has now been installed. To start using this feature, follow these steps:

1. Login to Enterprise Manager Grid Control.
2. Click **Targets** and the **Middleware** tab to navigate to the Middleware page.
3. Click on a WebLogic Server target in the list. The following screen is displayed.

Figure 11–1 WebLogic Server - JVM Diagnostics Page



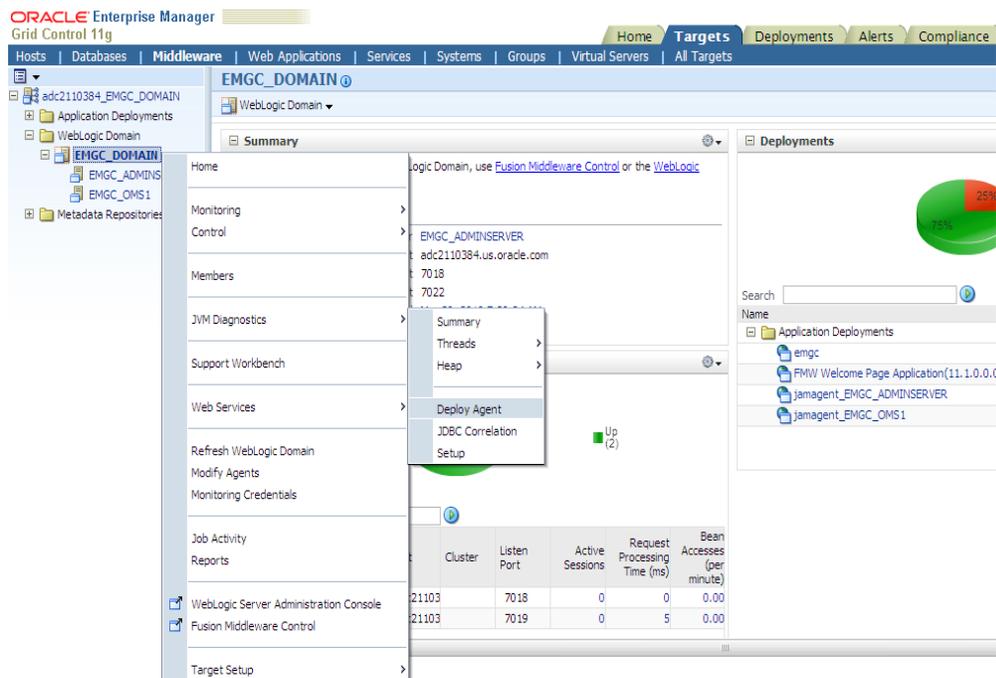
4. To start using JVM Diagnostics, select the JVM Diagnostics option from the WebLogic Server drop-down menu and choose the appropriate option.

11.1.3 Deploying the JVM Diagnostics Agent

To deploy the JVM Diagnostics Agent, follow these steps:

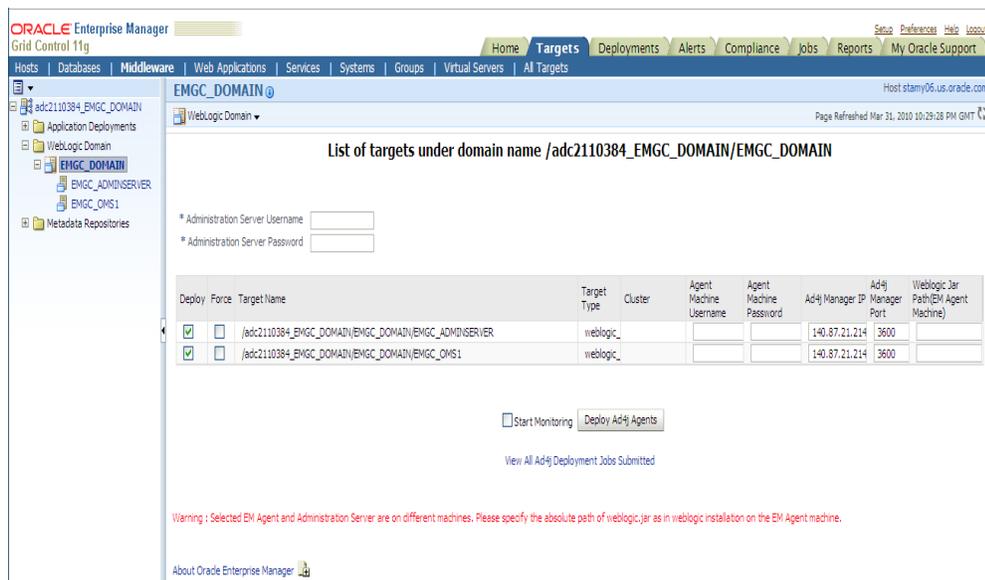
1. Select a WebLogic Server target on the Middleware page.
2. Select the JVM Diagnostics option in the drop-down list and click **Deploy Agent** as shown below:

Figure 11–2 Deploy Agent Selection Page



3. The following screen is displayed.

Figure 11–3 Deploy Agent Page



4. Enter the user name and password for the Administration Server.
5. Check the **Deploy** checkbox to select the agents that are to be deployed and click **Deploy AD4J Agents** to deploy them.

11.1.4 Creating a Special Repository User for Loading JVM Diagnostics Heaps

To create a special repository user who can load JVM Diagnostics Heaps to the repository, follow these steps:

Note: This is an optional step and can be used to create less privileged users who can load heaps using the loadHeap script.

1. Run the `create_jvm_diagnostic_db_user.sh` script from the command prompt

```
$ ./create_jvm_diagnostic_db_user.sh
```

2. You will see a message prompting you to continue with the user creation. Enter Y to continue. The following prompts are displayed:

```
Enter Database Host Machine Name or IP address:[localhost]:
```

Enter the host name or the IP address of the machine on which the database has been installed.

```
Enter Database Port on "localhost":[1521]
```

Enter the database port number as 25062.

```
Enter ORACLE_SID for database on "localhost":[s0222]
```

Enter the Oracle system identifier for the database.

```
Enter database admin user name:[SYSTEM]
```

Enter the name of the database administration user.

```
Enter "SYSTEM" user password:
```

Enter the password for the database administrator.

```
Enter JVM Diagnostics database username:[JVMDIAG]
```

Enter the user name for the JVM Diagnostics Administrator being created.

```
Enter "JVMDIAG" user password:
```

Enter the password for the JVM Diagnostics Administrator. You are then prompted to re-enter the password. After you have re-entered the password, you will see the following:

```
Creating Database user "JVMDIAG" ....
Database user "JVMDIAG" created
Public synonyms would be created for following database objects:
Tables
1) JAM_HEAPREL
2) JAM_HEAPOBJ
3) JAM_HEAPUSAGE
4) JAM_HEAPROOTS
5) JAM_HEAPROOTREL
6) JAM_HEAPOBJSUM
7) JAM_HEAPSNAP
```

3. You are prompted to continue with the setup. Press Enter to continue. You will see the following:

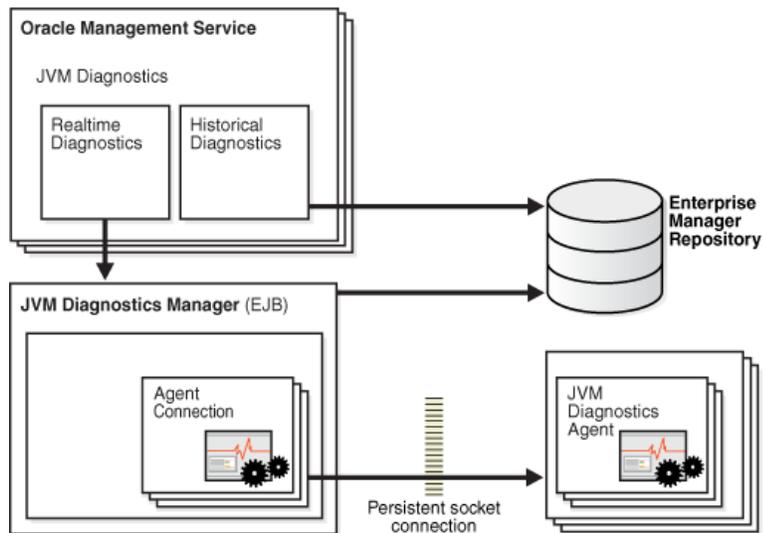
```

Required privileges granted
JVM Diagnostics Database User "JVMDIAG" created
Now onwards, to load heap, use "JVMDIAG" user
    
```

11.1.5 JVM Diagnostics Architecture

The following diagram shows the JVM Diagnostics Architecture.

Figure 11-4 JVM Diagnostics Architecture



The JVM Diagnostics Manager that is present on the WebLogic Server target retrieves historical metric data from the Management Repository (Enterprise Manager Repository). The JVM Diagnostics Agent deployed on the WebLogic Server target, collects real-time data and transmits it to the JVM Diagnostics Manager. The communication between the JVM Diagnostics Manager and Agent can be a secure (SSL) or non-secure connection.

Setting Up Request Monitoring

The Request Monitoring feature provides end-to-end visibility into requests for Java applications deployed to WebLogic Server. It establishes a clear view of a request flowing through a multi-tier java environment. This chapter describes the procedure to setup and enable the Request Monitoring feature in Enterprise Manager Grid Control. This chapter covers the following:

- [Prerequisites](#)
- [Setting Up Request Monitoring](#)
- [Request Monitoring Architecture](#)

12.1 Prerequisites

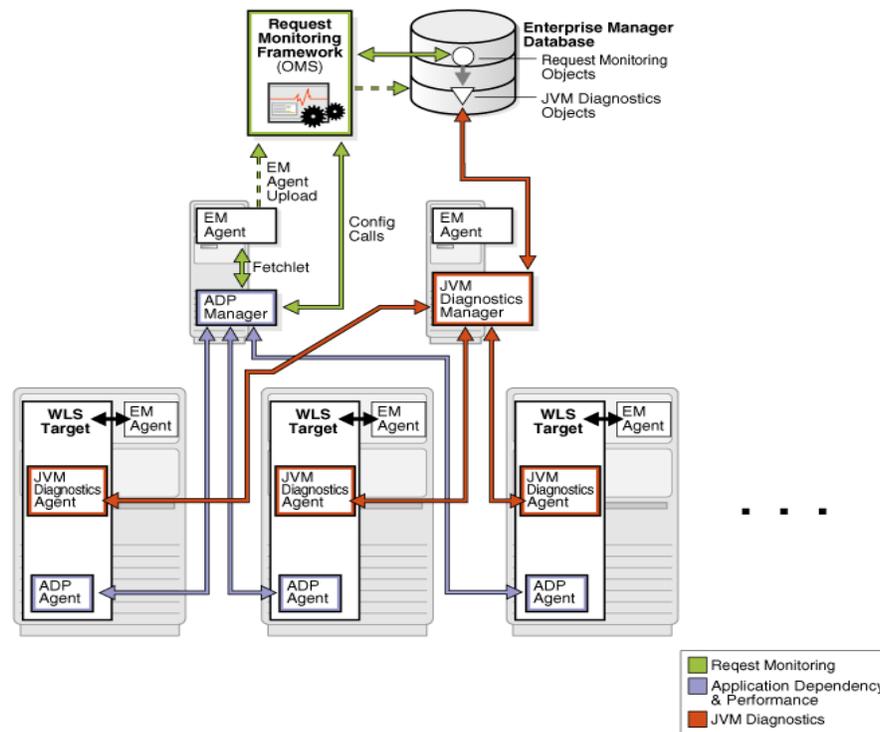
Before you enable Request Monitoring, you must do the following:

- Identify the WebLogic Server targets and J2EE applications that are to be monitored.
- Ensure that the necessary hardware for the ADP Manager, ADP database, and JVM Diagnostics Manager is available.

12.2 Request Monitoring Architecture

The following diagram shows the request monitoring architecture.

Figure 12-1 Request Monitoring Architecture



The JVM Diagnostics Agent and the ADP Agent are deployed on the WebLogic Server targets that are being monitored. These agents collect JVM Diagnostics and Application metrics and periodically report them to the JVM Diagnostics and ADP Managers. The Management Agent (EM Agent) on the ADP Manager machine, periodically runs a Fetchlet to retrieve the metrics from the ADP Manager and uploads them to the Management Repository. The Request Monitoring framework processes these metrics and uses them to monitor and analyze requests.

12.3 Setting Up Request Monitoring

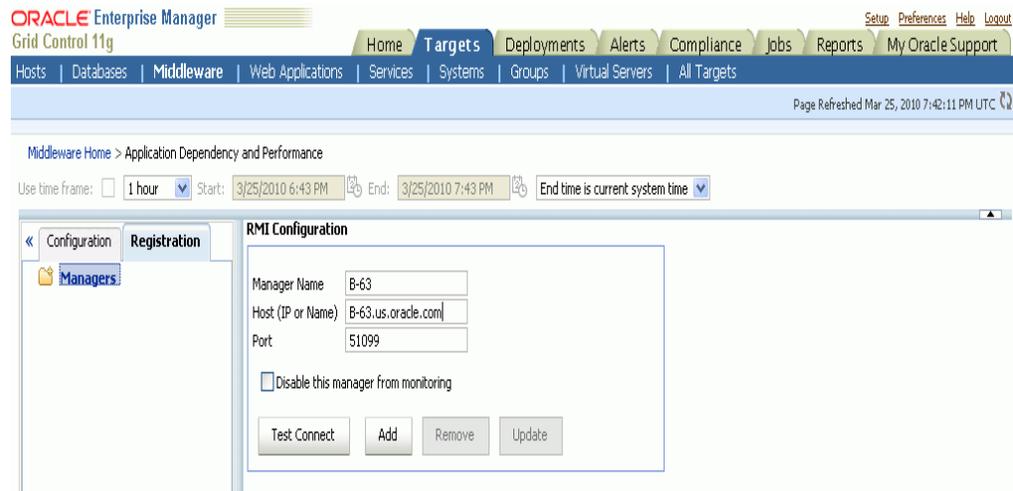
To setup end-user request tracing, follow these steps:

1. Deploy the Oracle Management Service and the Management Repository database. For details, see the *Enterprise Manager Basic Installation and Configuration Guide*.
2. Deploy the ADP Manager(s) on the designated machines. For details on deploying the ADP Manager, see the *Getting Started with SOA Management Pack Plus Guide*. The ADP Manager must be installed as the same user as the Management Agent (as described in the step below).
3. Deploy the Management Agent on all the machines on which the ADP Managers have been installed.
4. Deploy the Management Agent on the machine on which the WebLogic server target to be monitored is running. In Enterprise Manager Grid Control, click **Targets** and **Middleware**. On the Middleware page, check if the WebLogic Server target to be monitored has been discovered and registered as a target. If it has not been registered, add the WebLogic Server as a target and register it. For details on discovering and registering the WebLogic Server target, see the *Oracle Fusion*

Middleware Administration Guide. If this WebLogic server target is configured to connect to a database running on a remote machine, ensure that the Management Agent is deployed on the remote machine, and confirm that this database is listed as a target in Enterprise Manager Grid Control.

5. Register the WebLogic Server target with ADP Manager. To do so, click the Application Dependency and Performance link on the Middleware page. Click the **Registration** tab on the right panel. The following screen is displayed.

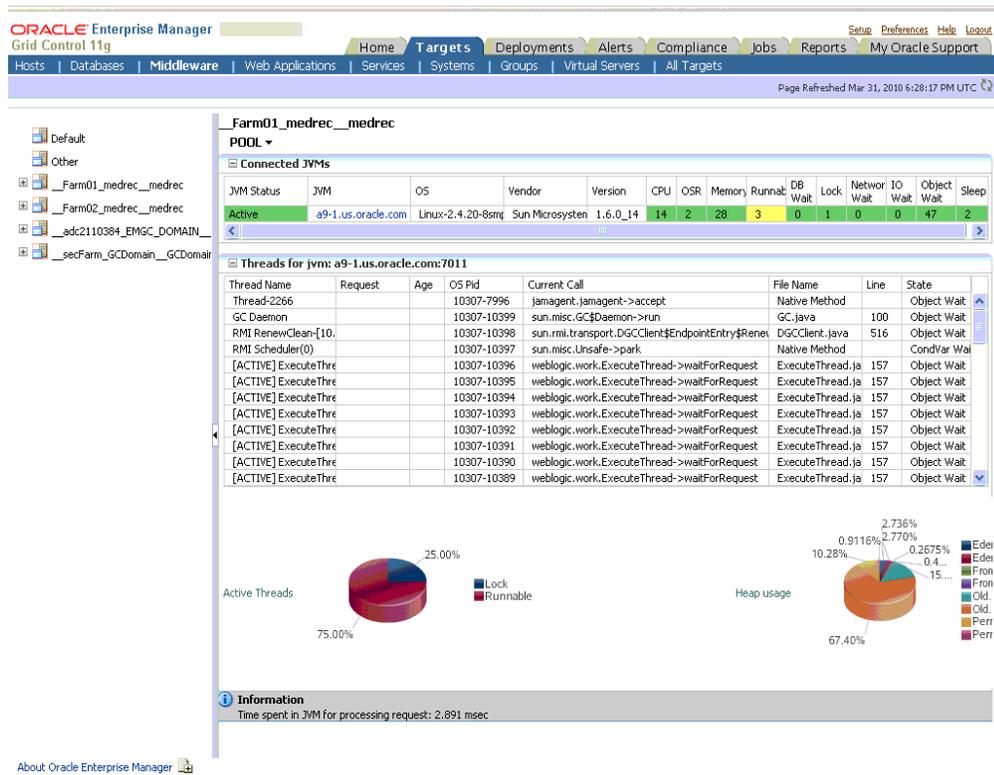
Figure 12–2 RMI Configuration Page



Enter the name of the ADP Manager to be registered and the fully qualified host name of the machine that the ADP Manager is running on. The default port 51099 is pre-populated for you. Click **Test Connect**. If you can successfully connect to the host machine, click **Add** to complete the ADP Manager registration. For more details on registering the ADP Manager, see the *Getting Started with SOA Management Pack Plus Guide*.

6. Install the JVM Diagnostics Manager and deploy the JVM Diagnostics Agent on WebLogic Server targets. For more details, see the *Setting Up the JVM Diagnostics Manager* chapter.
7. To validate that the JVM Diagnostics Manager is communicating with the JVM Diagnostics Agent, click the **JVM Diagnostics** link on the Middleware page. Click on a target in the left hand panel, for example `__Farm01_medrec__medrec`. In the right hand panel, click **Pool** > **Threads** > **Real-Time Analysis**. The following screen is displayed:

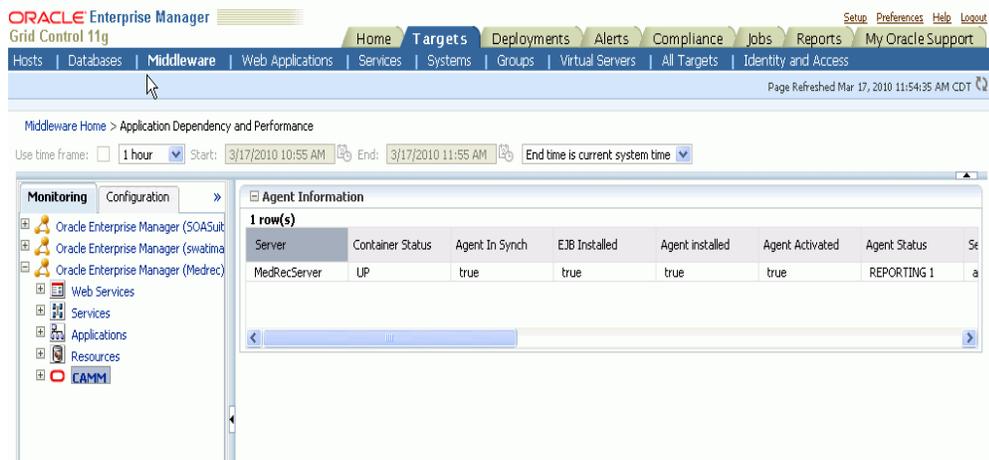
Figure 12–3 JVM Diagnostics All Threads Page



On this page, you can view the threads for each connected JVM.

8. Restart the ADP Manager and the WebLogic Server targets.
9. Click **Application Dependency and Performance** link on the Middleware page. Click the **Monitoring** tab on the right panel and select the **CAMM** option from the menu. The following screen is displayed.

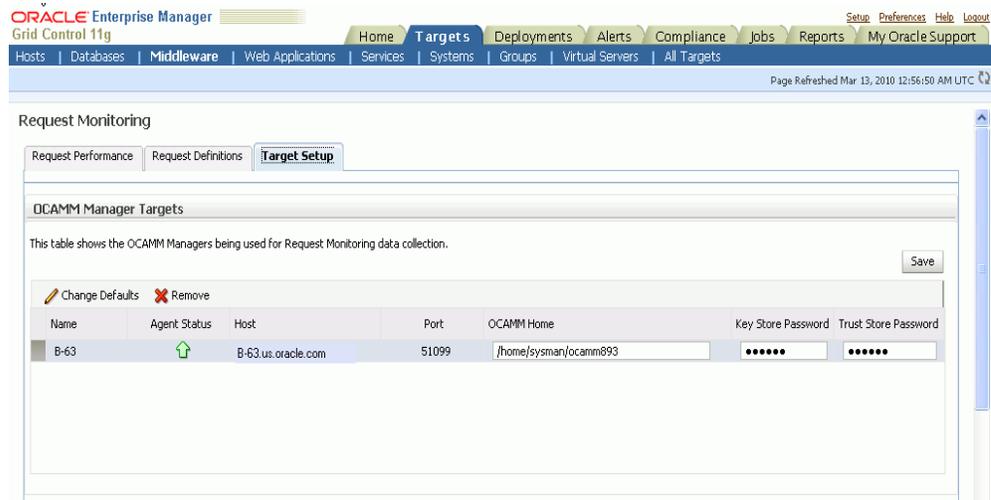
Figure 12–4 Agent Information Page



You can validate that the ADP Agent has discovered the Java EE target by checking the Agent Status on this page.

10. Click the **Request Monitoring** link on the Middleware page. Click the **Target Setup** tab. A list of ADP Managers available for data collection is displayed as shown in the following screen.

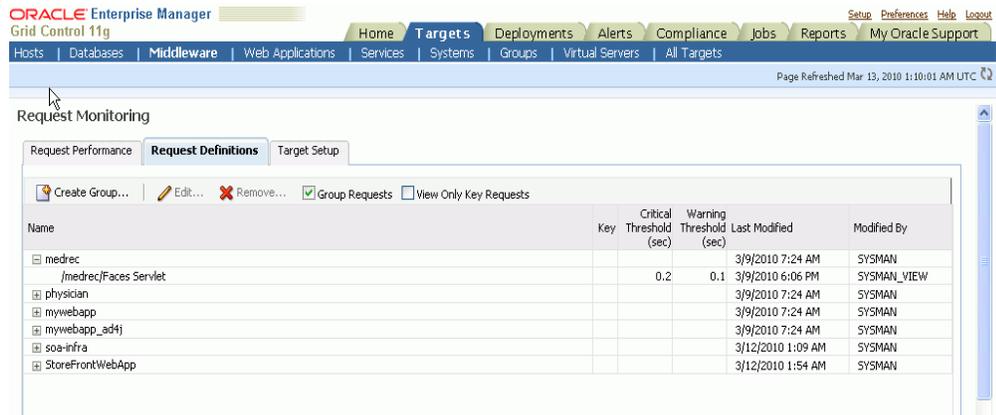
Figure 12–5 Target Setup Page



On this page, do the following:

- Select an ADP Manager from the list.
 - Select the Management Agent from the drop down list.
 - Enter the Key Store and Trust Store Password. The default password is `acseramanager`. For security reasons, it is recommended that you change this password. For more details, see the *Getting Started with SOA Management Pack Plus Guide*.
 - Click **Create Target**. The ADP Manager is now registered as a target, and starts collecting data for Request Monitoring.
11. Request Monitoring setup is now complete. To view and configure requests, grant the Request Monitoring Administrator privilege which provides full access to the Request Monitoring pages and allows the user to modify the request configuration settings. For richer Request Monitoring functionality, we recommend that you grant the following privileges:
 - View any Target privilege which allows user to drill down to the Target Home page and to view the target metrics in the Topology view.
 - JVM Diagnostics privilege which allows user to drill down to the JVM Diagnostics pages from the Topology page.
 12. After the appropriate privileges have been granted, click the **Request Monitoring** link on the Middleware page and then click the **Request Definitions** tab. The following screen is displayed.

Figure 12–6 Request Definitions Page



You can now view requests and their performance, group requests, and set request response time violation thresholds (Critical and Warning).

Managing Compliance

This chapter explains how Enterprise Manager Grid Control verifies that applications in your enterprise comply with preestablished standards. This chapter contains the following sections:

- [Compliance Overview](#)
- [Compliance Management](#)
- [Setting Up Compliance Evaluations](#)
- [Policy Groups Provided by Oracle](#)

13.1 Compliance Overview

To have your enterprise run efficiently, it must adhere to standards that promote the best practices for security, configuration, and storage. Once these standards are developed, you can apply and test for these standards throughout your organization; that is, test for compliance. Compliance is the conformance to standards, or requirements, or both.

Using Enterprise Manager Grid Control, you can test the conformance of your targets for security standards, and configuration and storage requirements.

By continually testing your systems, services, and targets, you are ensuring the best possible protection and performance your system can have.

Compliance is two-fold: evaluating the compliance of your targets, and the policy group life cycle management. The following sections describe these two concepts.

Note: To view the compliance features:

1. Navigate to the Grid Control Home page.
 2. Click the **Compliance** tab to access information regarding the policies, policy groups, and security statistics for your enterprise.
-
-

13.2 Compliance Management

Oracle provides two types of compliance management: policies and policy groups. Policies and policy groups define the optimal configurations of systems.

Policies and policy groups are similar in purpose, that is, they both provide rules against which managed entities are evaluated. However, there are differences:

- Rules within a policy group are managed as a set. They are viewed, maintained, evaluated, and reported, in the context of a policy group.

- Policy rules are not evaluated as a set. Policy rules are viewed, maintained, and evaluated as standalone entities.

Whether you use the out-of-box policies and policy groups defined by Oracle or customize policies to meet your particular system requirements, any deviations to your systems or applications are reported. Examples of deviations include inappropriate settings and incorrect system configurations.

This section contains the following subsections:

- [Accessing Compliance Management Pages in Grid Control](#)
- [Investigating Policy Violations and Policy Group Evaluation Results](#)
- [Assessing Security](#)
- [Viewing Policy Violations Results](#)
- [Policy Violations Reports](#)

13.2.1 Accessing Compliance Management Pages in Grid Control

To access compliance management pages in Grid Control:

- Click the **Compliance** tab to view violations associated with policies and policy groups.
 - Click the **Policies** tab for a roll-up view of all policy violations across all targets. From this tab, you can also access policy associations, the policy library, and errors.
 - Click the **Policy Groups** tab to view the list of policy groups against which there are violations. From this tab, you can also access the library of policy groups and policy group evaluation errors.
 - Click the **Security At a Glance** tab for a roll-up view of security statistics across the enterprise.
- Navigate to the Home page for a particular target. The links in the Policy Violations section display the number of policy violations according to severity level. Click the links to drill down to critical, warning, and informational policy violations for that target.

13.2.2 Investigating Policy Violations and Policy Group Evaluation Results

Here are a few suggestions for investigating policy violations. Attend to the most critical violations or those that have the biggest impact on your enterprise.

- Study the statistics on the Enterprise Manager Grid Control Home page. In particular, look at the statistics in the All Targets Policy Violations section. The policy violations with "Critical" severity should be dealt with first.
- Study the security-related violations reported in the Security Policy Violations section. Non-compliance with these policy rules can greatly impact the security of your enterprise.
- Address targets that have the lowest compliance scores.
- For the policy violations of a particular target, examine the home page for that target. The Policy Violations section provides overview information, but it also gives you access to the Policy Trend Overview for that target.
- To deal with policies regardless of the target, navigate to the **Compliance** tab and then click **Policies**. Using this tab, you have access to all the policy violations for

the enterprise, the policy associations, the policy rule library which lists all the policies, and policy evaluation errors.

- Navigate to the Policy Violations page and enter an appropriate value in the "Most Recent Violation within *n* days" filter.
- Suppress violations if you want to handle the violations at a later time.
- To deal with policy groups regardless of the target, click **Policy Groups**. Using this tab, you have access to all the evaluation results for the enterprise, the policy group library, and policy group evaluation errors.

Note: Only results from those targets for which you have VIEW privilege will be available for viewing.

- Navigate to the Evaluation Results page for a particular policy group. In the navigation tree, click the name of the policy group and a summary page lists all the targets along with the number of violations.
- Navigate to the Trend Overview page to see charts relating to the number of targets evaluated, the average violation count per target, number of targets by compliance score, and the average compliance score.

See Also: "About Policies" and "About Policy Groups" in the Grid Control online help for an overview of policies and policy groups and pointers to more information about viewing and managing policies and policy groups.

13.2.3 Assessing Security

Security policies are available for many targets, including Host, Database Instance, Cluster Database, Listener, OC4J, Oracle HTTP Server, and Web Cache.

Security policy groups are available for Database Instance, Cluster Database, and Listener.

Because security is crucial to the stability of your enterprise, security statistics are displayed prominently in Grid Control. On the Enterprise Manager Grid Control Home page, and many target home pages, there is a separate section displaying the security statistics for the target. This allows you to pay close attention to the security health of the enterprise.

In addition, the Security At a Glance feature provides an overview of the security health of the enterprise for all the targets or specific groups. This helps you to quickly focus on security issues by showing statistics about security policy violations and noting the critical security patches that have not been applied.

13.2.4 Viewing Policy Violations Results

To view the results of an evaluation, use any or all of the following:

- Study the statistics in the Policy Violations section of the target's home page
- Use the Policy Trend Overview page and the Trend Overview page available from the Evaluation Results page
- Access the Security At a Glance page

13.2.5 Policy Violations Reports

The Policy Violations reports and Policy Groups reports are available through the Reports feature. These reports deal with non-suppressed violations for all targets, groups, and a single target. The reports also deal with compliance summary for a group and target. In addition, suppressed violations are reported according to all targets, groups, and a single target.

See Also: [Chapter 8, "Information Publisher"](#)

13.3 Setting Up Compliance Evaluations

Compliance evaluation is the process of testing the policy and policy group rules against a target and recording any violations in the Oracle Management Repository.

13.3.1 Scheduling an Evaluation

For the evaluation to take place, you must enable the evaluation in one of two ways:

- For a policy, use the Metric Thresholds option on the Metric and Policy Settings page
- For a policy group, use the Policy Group Library page

13.3.2 Viewing Policy Group Evaluation Results

To view the results of a policy group evaluation, use the Policy Groups Evaluations Results page accessed through the Policy Groups tab.

1. From the Enterprise Manager Grid Control Home page, click the **Compliance** tab.
2. Click the **Policy Groups** tab, select the Evaluation Results page.
3. Choose the target type and policy group in which you are interested. If you are not sure what policy groups are available, click the **Library** tab and select the target type. Click **Go**. The policy group information appears.

13.3.3 Out-of-Box Policies and Policy Groups

Oracle provides a number of out-of-box policies (also known as policy rules) and policy groups for various targets.

When you add a target to Enterprise Manager, that target automatically uses predefined policy rules for that type of target. For example, Oracle provides security, configuration, and storage policy rules for the database instances and cluster databases. Security and configuration policy rules are provided for hosts.

Note: Policy Groups are *not* automatically associated with targets.

13.3.4 Customizing Policies

You can customize policies by editing the existing policy rule settings. You can enable or disable a policy evaluation, change the importance for the compliance score calculation, assign a corrective action, prevent template override, override default parameter values (when possible), and exclude objects from a policy's evaluation (when possible).

See Also: Online help for compliance scores

13.3.4.1 Defining Corrective Actions

One of the features of customizing policies is the ability to define corrective actions. Corrective Actions is a special type of job that executes automatically in response to a policy violation.

Corrective Actions utilize the Enterprise Manager Grid Control Job System and, like regular jobs, can consist of multiple steps, can be run with arbitrary host and target credentials, and reports its success or failure and its output to the Management Repository.

See Also: [Chapter 6, "Job System"](#)

13.3.4.2 Using Templates for Monitoring

A monitoring template defines all Enterprise Manager parameters you would normally set to monitor a target.

Monitoring templates simplify the task of setting up monitoring for large numbers of targets by allowing you to specify the monitoring and policy settings once and applying them as often as needed. You can save, edit, and apply these templates across one or more targets or groups.

See Also: "Monitoring Templates" in [Chapter 1, "Monitoring"](#)

13.4 Policy Groups Provided by Oracle

Policy groups serve as standards by which targets are measured. Policy groups report deviations and enable closed loop remediation by optionally taking action to bring systems back into compliance. Oracle provides the following policy groups:

- [Secure Configuration for Oracle Database](#)
- [Secure Configuration for Oracle Real Application Cluster](#)
- [Secure Configuration for Oracle Listener](#)

These standards represent best practices and allow you to maintain consistency across enterprise systems and configurations. The trend analysis feature allows fine grained tracking of compliance progress over time.

The following sections provide the highlights of each policy group.

13.4.1 Secure Configuration for Oracle Database

This policy group adheres to the security standards available for the Oracle Database. The categories deemed the most important for this policy group are:

- Post Installation
 - These rules ensure that a database is not compromised by having a default database server account left open that uses its default password.
- Oracle Directory and File Permissions
 - These rules ensure that access should be restricted, making it more difficult for an operating system user to attack the database.
- Oracle Parameters Settings
 - These rules ensure database initialization parameter settings are secure.
- Database Password Profile Settings

These rules ensure database profile settings are correctly defined. Oracle password management is controlled through the use of user profiles which are then assigned to database users, enabling greater control over database security.

- Database Access Settings

These rules ensure that access to and use of the database at the object level is restricted such that users are only given those privileges that are actually required to efficiently perform their jobs.

13.4.2 Secure Configuration for Oracle Real Application Cluster

This policy group adheres to the security standards available for the Oracle Cluster Database. The categories deemed the most important for this policy group are:

- Post Installation

These rules ensure that a database is not compromised by having a default database server account left open that uses its default password.

- Oracle Directory and File Permissions

These rules ensure that access should be restricted, making it more difficult for an operating system user to attack the database.

- Database Password Profile Settings

These rules ensure database profile settings are correctly defined. Oracle password management is controlled through the use of user profiles which are then assigned to database users, enabling greater control over database security.

- Database Access Settings

These rules ensure that access to and use of the database at the object level is restricted such that users are only given those privileges that are actually required to efficiently perform their jobs.

13.4.3 Secure Configuration for Oracle Listener

This policy group adheres to the security standards available for the Oracle Listener. The categories deemed the most important for this policy group are:

- Oracle Directory and File Permissions

These rules ensure that access should be restricted, making it more difficult for an operating system user to attack the database.

- Network Configuration Settings

These rules ensure that network configuration parameter settings are secure.

Configuring Services

This chapter describes how to configure services in Oracle Enterprise Manager 10g Grid Control Console. It contains the following sections:

- [Summary of Service Management Tasks](#)
- [Setting up the System](#)
- [Creating a Service](#)
- [Configuring a Service](#)
- [Recording Web Transactions](#)
- [Monitoring Settings](#)
- [Configuring Aggregate Services](#)
- [Configuring End-User Performance Monitoring](#)
- [Managing Forms Applications](#)
- [Configuring OC4J for Request Performance Diagnostics](#)
- [Setting Up Monitoring Templates](#)
- [Configuring Service Levels](#)
- [Configuring a Service Using the Command Line Interface](#)
- [Troubleshooting Service Tests](#)

14.1 Summary of Service Management Tasks

This table provides a summary list of all the service management features and their requirements.

Table 14–1 Summary of Service Management Tasks

Feature	Description	Requirements	Refer to
Test Performance	This feature allows you to proactively monitor services using service tests or synthetic transactions and determine their performance and availability from different user locations using beacons. For Web transactions, you can monitor the transactions at the transaction, step group and step level.	<ul style="list-style-type: none"> ■ Management Agent for enabling a beacon. ■ Microsoft Internet Explorer 5.5 or later 	Configuring a Service
End-User Performance Monitoring	Enterprise Manager allows you to gather end-user performance data and monitor the performance of the pages within your Web application. The End-User Performance Monitoring feature allows you to: <ul style="list-style-type: none"> ■ Understand real end-user page response times within your application. ■ Assess the user impact of performance problems. ■ Analyze end user response times by page, domain, region, visitors, and Web server. 	<ul style="list-style-type: none"> ■ Oracle HTTP Server Based on Apache 2.0 or Apache HTTP Server 2.0 ■ Oracle Application Server Web Cache (10.1.2, 9.0.4, 9.0.3, or 9.0.2) 	Configuring End-User Performance Monitoring
Interactive Transaction Tracing	Enterprise Manager provides a mechanism for interactively tracing Web transactions. This feature allows you to: <ul style="list-style-type: none"> ■ Diagnose performance problems at the transaction level. ■ Interactively trace transactions and analyze breakout of J2EE server activity times (servlet, JSP, EJB), and database times, including individual SQL statements. 	<ul style="list-style-type: none"> ■ Microsoft Internet Explorer 5.5 or later for creating and playing back transactions. ■ Oracle Application Server 10g (9.0.4) for playing back a transaction with trace to view J2EE server activity times. <p>Note: Recording a transaction is an optional feature. You can manually create a transaction by entering the required values.</p>	Configuring Interactive Transaction Tracing

Table 14–1 (Cont.) Summary of Service Management Tasks

Feature	Description	Requirements	Refer to
Request Performance	Enterprise Manager can gather critical request performance data about your Web application. The Request Performance feature allows you to: <ul style="list-style-type: none"> Diagnose root cause of performance problems. View historical tracing of J2EE middle tier activity. View breakouts of J2EE server processing times (servlet, JSP, EJB), and database times, including individual SQL statements. Correlate request performance to other Web application component metrics. View the full request processing call stack. 	Oracle Application Server 10g (9.0.4) and above	Configuring OC4J for Request Performance Diagnostics
Root Cause Analysis	The Root Cause Analysis (RCA) feature provides you with the ability to analyze and determine possible causes of service failure. The Topology Viewer provides a graphical representation of the service and its relationship to other services, systems and infrastructure components, with the causes identified by RCA highlighted in the display.	For the Topology Viewer <ul style="list-style-type: none"> Microsoft Internet Explorer 5.5 or higher Adobe SVG Viewer 3.0 	Root Cause Analysis Configuration
Forms Applications	A Forms Application target in Enterprise Manager can be used to model and monitor a specific Forms application. You can: <ul style="list-style-type: none"> Record and monitor a Forms transaction. Measure the End-User Performance of Forms actions such as Commit, Query, Runform, Callform, Openform, and Newform. 	<ul style="list-style-type: none"> Microsoft Internet Explorer 5.5 or later for creating and playing back Forms transactions. Oracle HTTP Server or Apache HTTP Server Oracle Application Server Web Cache (10.1.2, or 9.0.4) 	Recording and Monitoring Forms Transactions Monitoring the End-User Performance of Forms Applications

14.2 Setting up the System

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. Refer to the Enterprise Manager Online Help for details on defining systems.

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.

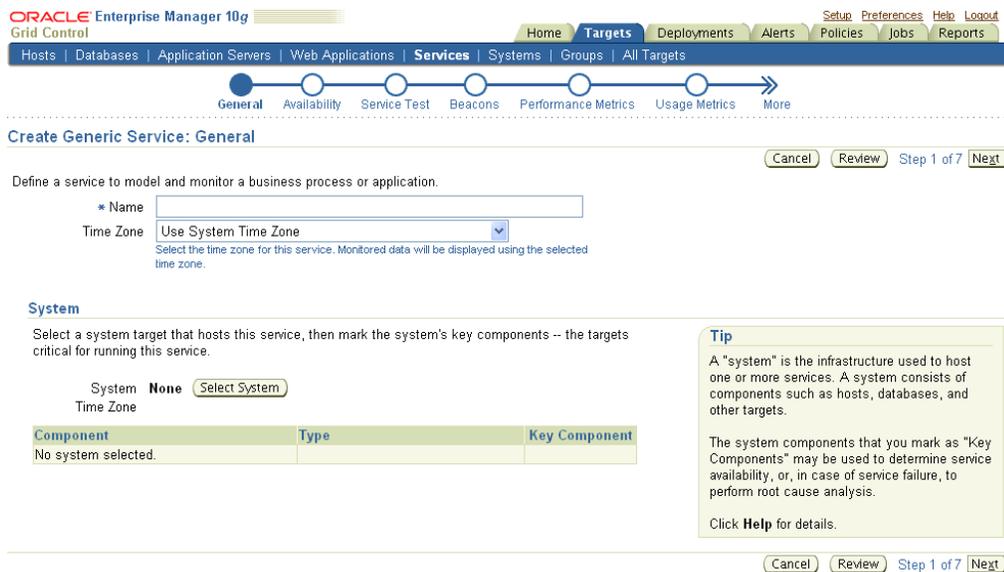
14.3 Creating a Service

Before you create a service, you must be familiar with the concepts of service management as described in the *Oracle Enterprise Manager Concepts*. You must also do the following:

- Install the Management Agent on the hosts on which the components of your service have been installed.
- Discover all the components for your service so that they can be listed as Enterprise Manager targets.
- Define systems on which the service is to be hosted.

To create a service, click the **Targets** tab and **Services** subtab. The Services main page is displayed. Select a service from the Add drop-down list and click **Go**. The following screen is displayed:

Figure 14–1 Create Service - General Page



Follow the steps in the wizard to create your service. This involves the following:

- Identifying the type of service to be created. You can define different types of services based on your requirement. Some of the services that you can define are Generic Service, Web Application, Aggregate Service, and Forms Application. A Generic Service is used to monitor a variety of different protocol based services. A Web Application is used to monitor Web transactions. Enterprise Manager provides additional monitoring and diagnostics features for Web applications. A Forms Application is used to monitor Forms transactions. Each Forms transaction can consist of one or more actions that can be monitored. You can also define other services that are specific to an application such as the OCS Service. You can combine one or more services to form an Aggregate Service.
- Specifying the name and time zone for the service.
- Selecting a system target that hosts this service and then marking the system's key components 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. For more information on defining systems and monitoring them, refer to the Service Management chapter in *Oracle Enterprise Manager Concepts*.

- 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 the one or more key beacons. If availability is based on system, availability is 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, click **Create** to create a new beacon.

Note: 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.

For Web applications, you can compare the performance of the service test execution from each remote beacon against the local beacon.

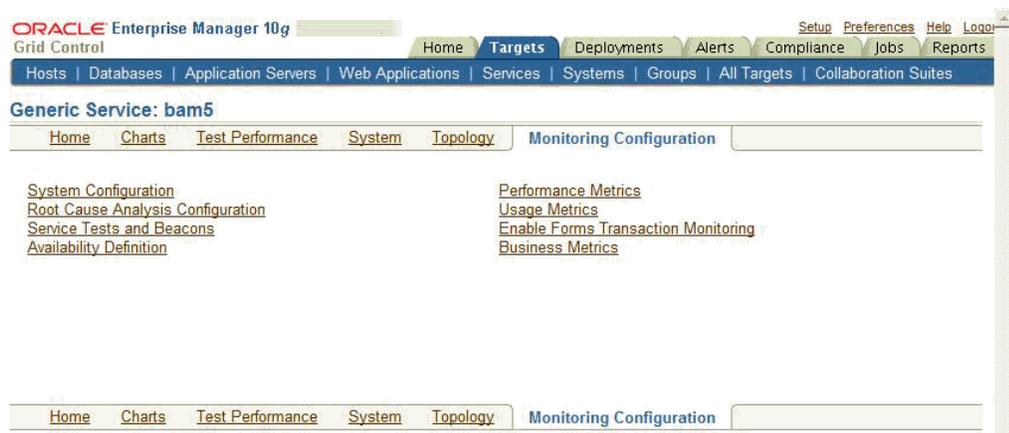
- Defining the metrics that will be used to measure the performance of the service. Performance metrics can be based on service tests or system components. After defining the metrics, you can specify the critical and warning thresholds. You can also specify the metric that is to be displayed in a graphical 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.

- 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.

14.4 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 14–2 Monitoring Configuration Page

The following options are available:

- [Availability Definition](#)
- [Performance Metrics](#)
- [Usage Metrics](#)
- [Business Metrics](#)
- [Service Tests and Beacons](#)
- [Root Cause Analysis Configuration](#)

Apart from these options, for Web applications, the end-user and request performance monitoring features can also be configured. For more information, refer to the following sections:

- [Configuring End-User Performance Monitoring](#)
- [Configuring OC4J for Request Performance Diagnostics](#)

14.4.1 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 14-13.

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.

14.4.2 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 14-2 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 14–3 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.

14.4.3 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 14–4 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.
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.

14.4.4 Business Metrics

Business metrics are used to measure the performance of business in an organization. These metrics are based on business indicators that can assess the business performance. You can define one or more system based metrics and specify critical and warning thresholds for these metrics. You can define business metrics for Generic Services and Aggregate Services.

Note: This option is available only if one of the system components is a service and has business metrics associated with it.

You can do the following:

- Add a business 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 the business metric to be based on the performance of one specific system component
 - Aggregate the metric across multiple components. Use this option if you want to measure the business performance from multiple components. Select the appropriate aggregation function from the drop down list. If you choose this option, you need select the appropriate aggregation function.

Table 14–5 Aggregation Functions - Usage Metrics

Function	Description
Maximum	The maximum value of the metric across all components will be used as the value of this business metric for the service.
Minimum	The minimum value of the metric across all components will be used as the value of this business metric for the service.
Average	The average value of the metric across all components will be used.
Sum	The sum of the values of metrics across all components will be calculated.

- Edit a business metric that has been defined.
- Delete a business metric that has been defined.

You can define system based metrics only. You can configure non-system based metrics by using the Data Exchange feature which facilitates data transfer between Enterprise Manager Grid Control and other external monitoring systems. For details, refer to the *Oracle Enterprise Manager Integration Guide*.

14.4.5 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 on the Monitoring Configuration page. The Service Tests and Beacons page appears. 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 FTP, Web Transaction, DNS, SOAP and others. The Create Service Test page is displayed. Refer to the Enterprise Manager Online Help for details on the various types of service tests.

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 the Oracle Management Service for the changes to be effective.

- 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.
- 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: 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.

For more details on creating different types of service tests, refer to the Enterprise Manager Online Help.

14.4.5.1 Configuring the Beacons

This section lists additional beacon related configuration tasks.

- **Configuring SSL Certificates for the Beacon:** To configure SSL certificates for Web transaction and Port Checker service tests, follow the steps given below:
- **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 -Xmx512m` to the `agentJavaDefines` property. This increases the Java VM heap size to 512 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.

14.4.5.2 Configuring Windows Beacons for Web Transaction (Browser) Playback

To run a Web Transaction (Browser) service test, you need beacons that are running on an 10.2.0.4 or later Management Agent on Windows. The beacon drives an Internet Explorer process. This process runs as the same user as the Management Agent service.

Verifying Web Transaction (Browser) test involves the following 3 steps:

1. Navigate to the **Service Tests and Beacons** page and select a Web Transaction (Browser) test from the list.
2. Click **Verify Test**. The Verify Service Test page is displayed.
3. Select a Windows beacon and click **Perform Test**.

One of the common problems that you may encounter is that the **Perform Test** does not respond immediately.

There may be several reasons for this delay. Complicated tests may take longer to run. However, the most probable cause for delayed response is when the Internet Explorer process from the beacon is waiting for manual confirmation, which is invisible when run as a process that does not interact with desktop.

You may need to change the browser settings on the beacon machine. These settings need to be changed for the Local Service account and are account specific. Therefore, any changes to the Internet Explorer process that was opened from the Start menu on the beacon machine, will not affect the Internet Explorer process instantiated from the beacon which runs in an invisible window. To make the Internet Explorer window instantiated from the beacon visible:

1. Login as administrator to the Windows machine on which the Management Agent is running.
2. From the Start menu, click **Run**, type `services.msc` and click **Enter**.
3. Find the Management Agent in the list of Windows services, e.g. `OracleServiceagent1`.

4. Right click the Management Agent and select **Properties**.
5. Click the **Log On** tab.
6. Click the **Select Allow service to interact with desktop** checkbox and click **OK**.
7. Right click the Management Agent and select **Stop**, and then select **Start**.

To check Internet Explorer on the Management Agent machine for any dialog confirmations. For example, SSL Certificates and security warnings.

1. Use the previous procedure to make the Internet Explorer process instantiated from the beacon visible.
2. Launch Enterprise Manager (from any machine), and navigate to the Service Tests and Beacons page of the corresponding service target.
3. Select the Web Transaction (Browser) service test, and click **Verify Service Test**.
4. From the Verify Service Test page, select the beacon running on the Windows, and click **Perform Test**.
5. If it is a SSL Certificates issue, From the Windows machine on which the Management Agent is running, you will see an Internet Explorer window open and a Security Alert with a **View Certificate** option is displayed.
6. Select the **Certificate Path** tab, click the root certificate, which should have a red cross next to the name, and click the **View Certificate** button.
7. Click **Install Certificate** and proceed with the **Certificate Import Wizard**. (Click Next and Yes for any prompts).

Note: Other security warnings may also pause the Internet Explorer automation process. Typically, these security warnings have a check box that allow you disable the display of all future warning messages for all Web sites. These warnings may have already been dismissed on the machine where the transaction was recorded.

8. Once this manual step has been performed, the Internet Explorer process should be in auto-pilot mode until the service test is completed. The warning message will not be displayed when you play back the service test next time.
9. Click **Perform Test** again to make sure the entire service test is completed automatically the second time.

To make the Internet Explorer window instantiated from the beacon invisible, you can repeat the steps 1 to 5, uncheck the **Select Allow service to interact with desktop** checkbox and continue with step 7.

To configure the proxy setting for Web Transaction (Browser) service tests:

1. Make the Internet Explorer process instantiated from the beacon visible.
2. Launch Enterprise Manager (from any machine), and navigate to the Service Tests and Beacons page of the corresponding service target.
3. Select the Web Transaction (Browser) service test, and click **Verify Service Test**.
4. From the Verify Service Test page, select the beacon running on the Windows, and click **Perform Test**.

5. From the Windows machine where the Management Agent is running, you should see two Internet Explorer windows open. From either of the windows, select the **Tools > Internet Options**.
6. Click the **Connections** tab and then click **LAN Settings**, and make all relevant changes there. These changes apply to all service tests running on this beacon.
7. Close both the Internet Explorer windows.
8. Click **Perform Test** again to make sure the entire service test is completed automatically the second time.
9. Make the Internet Explorer process instantiated from the beacon invisible.

Note: At any one time, each test run launches two Internet Explorer windows. One of the windows schedules the steps during playback. The other window actually shows the site being played back.

14.4.6 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 more 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 allows you to see a graphical representation of the service, system and component dependencies with the targets highlighted that RCA has implicated as causing the service failure.

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 component tests for the service on the Component Tests page by adding, removing, or editing tests. 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.

14.4.6.1 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.

Your services may be easier to manage in the modular fashion, and RCA will consider not only the status of a sub-service (a service that you depend on) but also on any of the system components or service that the sub-service depends on in turn and provides you with the power to encapsulate the services a key component exposes to you in the form of a managed service that your service may then depend on.

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.

14.5 Recording Web Transactions

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 Grid 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.

14.6 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, click **Monitoring Settings for Tests** link on the Service Tests and Beacons page. 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. Alerts 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. Alerts 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.

14.7 Configuring Aggregate Services

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 critical and warning alerts, 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 critical and warning alerts, 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 critical and warning alerts, 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.

14.8 Configuring End-User Performance Monitoring

Enterprise Manager allows you to monitor the response time data generated by actual end-users as they access and navigate your Web site. You can gather end-user performance data and monitor the performance of the pages within your Web application. The Web servers such as OracleAS Web Cache, Oracle HTTP Server, and Apache HTTP Server collect the end-user performance data and store it in the log file. Enterprise Manager processes this data and uploads it to the Management Repository. You can then view and analyze this data on the Page Performance page.

To gather the end-user performance data, you must configure one of the following Web servers so that Website activities are logged and stored in the correct format.

- Oracle HTTP Server Based on Apache 2.0
- OracleAS Web Cache

- Apache HTTP Server 2.0 or higher

After you have configured one of these Web servers, you can enable the collection of end-user performance data. You can then view the end-user performance data on the Page Performance page in Enterprise Manager.

Before you configure the Web server, you must do the following:

- Create a Web application target that contains one of these Web servers.
- Make this Web server as a key system component for your Web application. If this Web server is a part of the Redundancy Group, make sure that the Redundancy Group is a key system component of your Web application. To enable end-user performance monitoring, you must configure the specific Web server within the Redundancy Group.

Note: If you are using the Oracle HTTP Server Based on Apache 2.0, the Redundancy Group is referred to as the HTTP Server HA Group.

The following sections provide instructions on configuring the Web server for End-User Performance Monitoring:

- [Configuring End-User Performance Monitoring Using Oracle HTTP Server Based on Apache 2.0 or Apache HTTP Server 2.0](#)
- [Configuring End-User Performance Monitoring Using Oracle Application Server Web Cache](#)

14.8.1 Configuring End-User Performance Monitoring Using Oracle HTTP Server Based on Apache 2.0 or Apache HTTP Server 2.0

To enable End-User Performance Monitoring, you can use either of the following Apache server versions:

- Oracle HTTP Server Based on Apache 2.0
- Apache HTTP Server 2.0 or higher (This can be downloaded from <http://www.apache.org>)

Before configuring either of these Apache server versions, you must perform the following steps:

1. In the Agent Home page, select either Oracle HTTP Server or Apache HTTP Server as a target type.
2. Add the target of the corresponding type and make sure the following properties are set in the Monitoring Configuration page:
 - For Oracle HTTP Server, fill in the version number (`stdApache10.1.2`), Log file directory and Log file name.
 - For Apache HTTP Server 2.0, fill in the install home directory, Log file directory and Log file name.

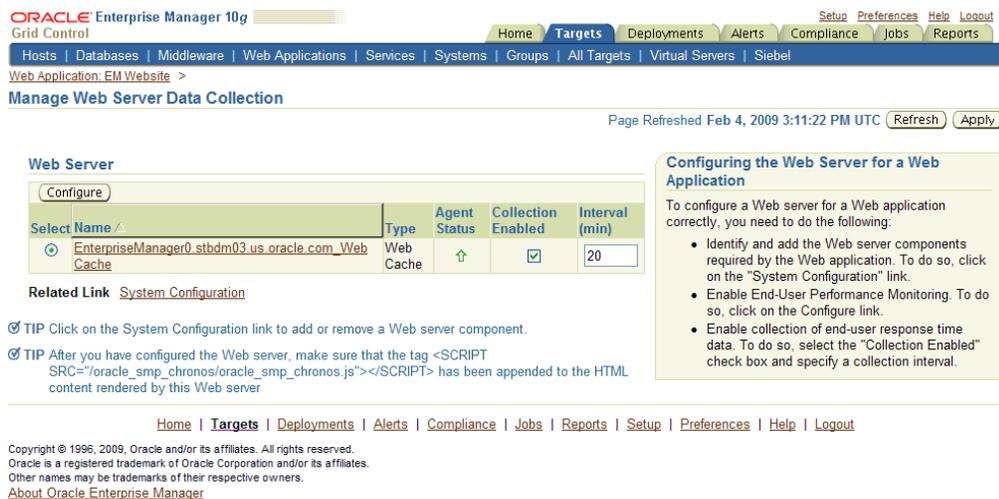
Note: If the Oracle HTTP Server is installed before the Management Agent has been installed, and is up and running during agent installation, then the target will be discovered automatically. Otherwise you need to manually create the Oracle HTTP Server target and specify the following properties: Machine name, Port number, Version of the Apache Server, Oracle home path, Log file directory (for EUM), Log file name (for EUM) where EUM refers to End-User Performance Monitoring.

3. Make sure you have created the Web application with this Web server target. For details on creating a Web application, refer to the pre-requisites in the "Configuring End-User Performance Monitoring" section on page 14-16.

To configure the Apache server and enable collection of end-user performance data, follow the steps given below:

1. Navigate to the Web Application Home page in the Grid Control Console and click **Monitoring Configuration**.
2. Click **Manage Web Server Data Collection**. You will see a table which lists the Web Servers including Oracle HTTP Server Based on Apache 2.0 or higher, Apache HTTP Server version 2.0 or higher, or OracleAS Web Cache.

Figure 14-3 Manage Web Server Data Collection



3. Select the Oracle HTTP Server or Apache HTTP Server from the table and click **Configure**. Enter the host credentials required for modifying the Apache configuration file.
4. After logging in, you will see a table containing the list of sites that are being hosted by the Apache server. These include a list of virtual hosts defined by the user in the Apache Configuration file. The up and the down arrows under the **Monitoring Status** column shows the corresponding site is currently being monitored. For each site, check or uncheck the **Enable Monitoring** checkbox to indicate whether this site is to be monitored. For the site that is to be monitored, enter the log file name in the text box to indicate the location in which the end-user performance data is to be stored. By default, the log file will be created under the

logs/directory under Apache root directory. To save the log file in a different directory, enter a file name with the absolute path.

5. Make sure that the log file name and the location you specify here match the Log file name and Log file directory in the Monitoring Configuration page of the Oracle HTTP Server or Apache HTTP Server target.
6. You can also use the one button accelerator to enable all sites or disable all sites all at once.
7. To selectively disable or enable certain URLs on a specific site, select the site, click **Set URLs**. Click **Insert Before** or **Insert After** to create a URL rule and place it in the desired place among all URL rules. A URL rule contains a **URL Pattern**, **URL Pattern Type**, and a check box indicating if this URL is to be monitored or not. For example, a URL rule with **URL Pattern** "abc" and **URL Pattern Type** "Ends With" and Monitor unchecked means that any URL ending with "abc" will not be monitored by End-User Performance Monitoring. The user can also delete a URL rule, move a URL rule up or down to increase or decrease its priority.
8. After you have made the configuration changes, click **OK** to go to the Apache Restart page. Restarting the Apache server will finalize all configuration changes, and end-user performance data will be logged by the Apache server.
9. After you have configured the Apache server, you will return to the Manage Web Server Data Collection page. You can now enable the collection of end-user performance data. For more details, refer to ["Starting and Stopping End-User Performance Monitoring"](#) on page 14-28. If you do not see data after End-User Performance Monitoring has been enabled, refer to the ["Verifying and Troubleshooting End-User Performance Monitoring"](#) on page 14-29.

14.8.1.1 Setting up the Third Party Apache Server

To set up the Third Party Apache HTTP Server 2.0, follow these steps:

1. Install the third party Application Server.
2. Install Apache HTTP Server 2.0.
3. Install the plug-in for the Apache HTTP Server 2.0 that was provided by the Application Server.
4. Ensure that the Web application works with the Apache HTTP Server 2.0 server. You can then follow the steps to configure the Apache server and enable collection of end-user performance data.

14.8.2 Configuring End-User Performance Monitoring Using Oracle Application Server Web Cache

Enterprise Manager uses data from Oracle Application Server Web Cache to gather statistics about the performance of pages within your Web applications. As a result, you must configure Oracle Application Server Web Cache to ensure that it logs your Web site activity and that the data is in the correct format.

When Oracle Application Server Web Cache is properly configured, Enterprise Manager can begin collecting the end-user performance data and load it into the Oracle Management Repository.

See Also: "Configuring End-User Performance Monitoring" in the *Oracle Application Server Web Cache Administrator's Guide*.

The following sections describe how to configure and collect end-user performance data if you are using the OracleAS Web Cache:

- [Configuring Oracle Application Server Web Cache 10.1.2](#)
- [Configuring Oracle Application Server Web Cache 9.0.4](#)
- [Configuring End-User Performance Monitoring Using Earlier Versions of Oracle Application Server Web Cache](#)
- [Configuring End-User Performance Monitoring Using Standalone Oracle Application Server Web Cache](#)

14.8.2.1 Configuring Oracle Application Server Web Cache 10.1.2

To configure the OracleAS Web Cache for End-User Performance Monitoring, follow the instructions in the following sections:

1. Navigate to the Web Application Home page in the Grid Control Console and click **Monitoring Configuration**.
2. Click **Manage Web Server Data Collection**. Enterprise Manager displays the Manage Web Server Data Collection page.
3. Select the Web Cache target and click **Configure**. Enterprise Manager displays the login dialog box for the Oracle Application Server Control.

Tip: If the login dialog box does not appear or if you see an error message in your browser window, navigate to the Web Cache Home page. Click **Administer** in the Related Links section. You will be prompted for the user name and password for the Application Server Control. Click **Administration** and scroll down and click **End-User Performance Monitoring**.

4. Enter the username and password for the Application Server Control user or the `ias_admin` account. The password for the `ias_admin` account is defined during the installation of Oracle Application Server.
5. After you have logged into Oracle Application Server Control, you can then configure the Oracle Application Server Web Cache using the Set Up End-User Performance Monitoring page. Check the **Enable End-User Performance Monitoring** checkbox and click **OK** to enable End-User Performance Monitoring at the Web Cache level.
6. At the site-level configuration section, select a site and check **Enable Monitoring** for that site.

Tip: Disabling End-User Performance Monitoring at the Web Cache level will override site-level settings.

7. From the drop-down list, select the Access Log Format as **access log:WCLF** for each site you want to monitor. If this format is not in the list, click **Use Required Log Format**. This automatically picks up the End-User Performance Monitoring log format.
8. Click the link under the **URLs to Monitor** column. The URLs To Monitor page is displayed. Click **Add Another Row** to create a URL rule and place it in the desired place among all URL rules. A URL rule contains a **URL Pattern**, **URL Pattern Type**, and a check box indicating if this URL is to be monitored or not. For example, a URL rule with **URL Pattern** "abc" and **URL Pattern Type** "Ends With" and **Monitor** unchecked means that any URL ending with "abc" will not be

monitored by End-User Performance Monitoring. The user can also change the priority of the URL rule by clicking **Reorder**. Click **OK** to save the changes and return to the Set Up End-User Performance Monitoring page.

9. After you have configured the Web Cache in the Set Up End-User Performance Monitoring page, click **OK** to save the changes. You will then return to the Web Cache Administration page in Oracle Application Server Control. Click **Restart** to restart the Web Cache. For more detailed information about configuring these options, click **Help** on the Set Up End-User Performance Monitoring page.
10. Close the Application Server Control browser window and return to the Manage Web Server Data Collection page in the Grid Control console. You can now enable the collection of end-user performance data. For more details, refer to "[Starting and Stopping End-User Performance Monitoring](#)" on page 14-28. If you do not see data after end-user performance has been enabled, refer to "[Verifying and Troubleshooting End-User Performance Monitoring](#)" on page 14-29.

14.8.2.2 Configuring Oracle Application Server Web Cache 9.0.4

To configure the Oracle Application Server Web Cache Manager 9.0.4, follow the instructions given in these sections:

1. Navigate to the Web Application home page in the Grid Control console and click **Monitoring Configuration**.
2. Click **Manage Web Server Data Collection**. Enterprise Manager displays the Manage Web Server Data Collection page.
3. Select the Web Cache target and click **Configure**. Enterprise Manager displays the login dialog box for the Web Cache Manager.

Tip: If the login dialog box does not appear or if you receive an error message in your browser window, you may have to start the Oracle Application Server Web Cache Manager. For more information about starting and using Oracle Application Server Web Cache Manager, refer to the *Oracle Application Server Web Cache Administrator's Guide*.

4. Enter the username and password for the Web Cache administrator account. The first time you log in to the Oracle Application Server Web Cache administrator account, the password is `administrator`. The password for the `ias_admin` account is defined during the installation of Oracle Application Server.
5. Enable OracleAS Web Cache logging for End-User Performance Monitoring:
 - a. Select **Logging and Diagnostics** and then select End-User Performance Monitoring in the OracleAS Web Cache Manager navigator frame.
You can enable monitoring for a particular Web cache or for an entire site.
 - b. To enable monitoring for a particular Web cache, select the Web cache from the **Cache-Specific End-User Performance Monitoring** section and click **Enable**.
Be sure to enable the Web cache that you are using as a front-end to your Web application.
 - c. To enable monitoring for the entire site, select the site from the **Site-Specific End-User Performance Monitoring** section and click **Enable**.
6. Configure Oracle Application Server Web Cache to use the Web Cache Log Format (WCLF):

- a. Select **Logging and Diagnostics** and then select Access Logs in the OracleAS Web Cache Manager navigator frame.
 - b. In the Cache-Specific Access Log Configuration table, click **Edit Selected** and enable the access log for your selected cache.
 - c. In the Site-Specific Access Log Configuration table, make sure that the Format style of the selected Site Name is `WCLF` and that it is enabled.
7. Click **Apply Changes** at the top of the Web Cache Manager window and restart OracleAS Web Cache by clicking **Restart** on the Web Cache Manager Cache Operations page.
 8. Close the Web Cache Manager browser window and return to the Manage Web Server Data Collection page in the Grid Control console. You can now enable the collection of end-user performance data. For more details, refer to "[Starting and Stopping End-User Performance Monitoring](#)" on page 14-28. If you do not see data after end-user performance has been enabled, refer to "[Verifying and Troubleshooting End-User Performance Monitoring](#)" on page 14-29.

14.8.2.3 Configuring End-User Performance Monitoring Using Earlier Versions of Oracle Application Server Web Cache

If you are managing an earlier version of the Oracle Application Server using the Oracle Enterprise Manager 10g Grid Control Console, you can monitor your Web applications with End-User Performance Monitoring, but you cannot configure your Oracle Application Server Web Cache instance from within the Grid Control console.

Instead, you configure End-User Performance Monitoring for Oracle Application Server Web Cache 9.0.2 and 9.0.3 by running the `chronos_setup.pl` script on the computer that hosts your Oracle HTTP Server.

14.8.2.3.1 Using the `chronos_setup.pl` Configuration Script

Before you begin, consider the following:

- The `chronos_setup.pl` script is installed in the `bin` directory of your Management Agent home when you install the Management Agent using the instructions in *Oracle Enterprise Manager Grid Control Installation and Basic Configuration*.
- You must run the `chronos_setup.pl` script as an operating system user with the privilege to write to the document root of your Oracle HTTP Server.
- If you have trouble running the script, run it with no arguments to display the help text.

To enable End-User Performance Monitoring for Oracle Application Server Web Cache 9.0.2 or Oracle Application Server Web Cache 9.0.3, you must run the `chronos_setup.pl` script three times, each time with a different argument:

- Once to configure the document root for each Web server in your Web site
- Once to configure Oracle Application Server Web Cache
- Once to start collecting response time data

The following sections describe each step of enabling End-User Performance Monitoring for Oracle Application Server Web Cache 9.0.2 or Oracle Application Server Web Cache 9.0.3.

14.8.2.3.2 Configuring the Document Root For Each Web Server When you run the `chronos_setup.pl` script with the `webserver` argument, the script:

- Creates a new directory inside the document root. The directory is called:

```
oracle_smp_chronos
```

- Installs two files into the `oracle_smp_chronos` directory:

```
oracle_smp_chronos.js
oracle_smp_chronos.gif
oracle_smp_eum_init.js
oracle_smp_eum_main.js
```

The `oracle_smp_chronos.js` must be installed in the document root of each Web server that serves content for your Website.

Note: If you have more than one document root, you must run the `chronos_setup.pl` script on each document root.

For example, if Oracle Application Server Web Cache and your Web server are on different machines and an Oracle Management Agent is present on the Web server machine, you must run the `chronos_setup.pl` script with the `webserver` option on the Web Server host to configure the document root for the remote Web server.

If Oracle Application Server Web Cache and your Web server are installed on different machines and you have no plans to install a Management Agent or to monitor the Web server, you will need to create a directory called `oracle_smp_chronos` under the Web server document root directory, and using FTP, place the `oracle_smp_chronos.js` file in the `oracle_smp_chronos` directory.

To configure the document root for each Web server:

1. Change directory to the `/bin` directory in the Management Agent home directory.

For example:

```
$PROMPT> cd AGENT_HOME/bin
```

2. Make sure you have write access to the Web server document root directory and then run the script as follows:

```
$PROMPT> perl chronos_setup.pl webserver location_of_the_webserver_DocumentRoot
```

An example of a Document Root is as follows:

```
$ORACLE_HOME/Apache/Apache/htdocs
```

To find the location of the document root, you can perform either of these steps:

- Log in to the Oracle Application Server Release 2 (9.0.2) Enterprise Manager Web site and navigate to the Oracle HTTP Server Home Page. The document root is displayed in the General section of the HTTP Server Home Page.
- Use a text editor or a command-line search utility to search for the term `DocumentRoot` in the following Oracle HTTP Server configuration file:

```
$ORACLE_HOME/Apache/Apache/conf/httpd.conf
```

14.8.2.3.3 Configuring Oracle Application Server Web Cache for End-User Performance Monitoring

To configure Oracle Application Server Web Cache for End-User Performance Monitoring, you run the `chronos_setup.pl` script with the `webcache` argument. The script sets up Oracle Application Server Web Cache for End-User Performance

Monitoring, and stops and restarts Oracle Application Server Web Cache automatically.

To configure Oracle Application Server Web Cache for End-User Performance Monitoring:

1. Make sure you have write access to the Oracle Application Server Web Cache directory.

For example, if Web Cache is installed in an Oracle Application Server home directory, you will need access to the `IAS_HOME/webcache` directory.

2. Change directory to the `/bin` directory in the Management Agent home directory.

For example:

```
$PROMPT> cd /private/agent_home/bin
```

3. Run the script as follows:

```
$PROMPT> perl chronos_setup.pl webcache webcache_installation_directory
```

Note: After running `chronos_setup.pl`, if you cannot restart Oracle Application Server Web Cache, back out of the configuration process by copying the following files back to their original name and location:

- `internal.xml<timestamp>`
 - `webcache.xml<timestamp>`
-
-

14.8.2.3.4 Starting End-User Performance Monitoring To start End-User Performance Monitoring, you run the `chronos_setup.pl` script with the `collection` argument. The script creates a collection file for the specified target and restarts the agent.

To start End-User Performance Monitoring:

1. Log in as the user who installed the Management Agent so you have write access to the following directory:

```
AGENT_HOME/sysman/emd/collection
```

2. Change directory to the `/bin` directory in the Management Agent home directory.

For example:

```
$PROMPT> cd AGENT_HOME/bin
```

3. Locate the name of the Oracle Application Server Web Cache target.

You can locate the name of the target in one of three ways:

- From the Oracle Enterprise Manager 10g Grid Control Console, locate the Oracle Application Server Web Cache target on the Targets tab. The name listed in the first column of the Target table is the name you must enter as an argument to the `chronos_setup.pl` script. Note the use of spaces and underscores.
- Search the contents of the `targets.xml` configuration file, which lists all the targets managed by the Management Agent. Locate the Oracle Application Server Web Cache entry in the file and use the `NAME` attribute for the Web Cache target. The `targets.xml` file is located in the following directory of the Management Agent home:

```
AGENT_HOME/sysman/emd/targets.xml
```

- Use the `emctl config agent listtargets` command to list the target names and target types currently being monitored by the Management Agent.

See Also: ["Listing the Targets on a Managed Host"](#) on page 7-11.

4. Start the collection for the Oracle Application Server Web Cache target by running the script as follows:

```
$PROMPT> perl chronos_setup.pl collection webcache_targetname
```

Note: If the name of the Oracle Application Server Web Cache target includes spaces, you must use quotation marks around the name

14.8.2.4 Configuring End-User Performance Monitoring Using Standalone Oracle Application Server Web Cache

Oracle Application Server Web Cache is available as a standalone download from the Oracle Technology Network (OTN). The standalone version of Oracle Application Server Web Cache allows you to improve the performance and reliability of your Web server even if you are not using Oracle Application Server.

If you are using standalone Oracle Application Server Web Cache with a third-party Web server, you can still manage Oracle Application Server Web Cache using the Oracle Enterprise Manager 10g Grid Control Console. As a result, you can also use End-User Performance Monitoring to monitor the Web applications that your users access through Oracle Application Server Web Cache.

Configuring End-User Performance Monitoring for standalone Oracle Application Server Web Cache involves the following steps, which are described in the following sections:

- [Installing Standalone Oracle Application Server Web Cache](#)
- [Configuring Standalone Oracle Application Server Web Cache](#)
- [Enabling End-User Performance Monitoring for Standalone Oracle Application Server Web Cache](#)

14.8.2.4.1 Installing Standalone Oracle Application Server Web Cache

To install the standalone version of Oracle Application Server Web Cache:

1. Navigate to the Oracle Technology Network (OTN):


```
http://otn.oracle.com/software/content.html
```
2. Locate and select the Oracle Application Server Web Cache download option and follow the links for your operating system.
3. Use the instructions on the OTN Web site to download Oracle Application Server Web Cache.
4. Use the instructions in the Web Cache readme file to install Oracle Application Server Web Cache in its own Oracle Home.

14.8.2.4.2 Configuring Standalone Oracle Application Server Web Cache

End-User Performance Monitoring uses data from Oracle Application Server Web Cache to gather statistics about the performance of pages within your Web

applications. As a result, Enterprise Manager obtains End-User Performance Monitoring data only when Oracle Application Server Web Cache is configured to improve the performance and reliability of your Web server.

See Also: *Oracle Application Server Web Cache Administrator's Guide* for complete instructions for configuring Oracle Application Server Web Cache

Specifically, you must perform the following Oracle Application Server Web Cache configuration tasks:

1. Change the default listening port of your HTTP Server (for example, 7777) to a new port number (for example, 7778) and restart the HTTP Server.

See Also: "Specifying Listening Addresses and Ports" in the Enterprise Manager Online Help if you are using Oracle HTTP Server and managing the server with Enterprise Manager.

Oracle HTTP Server Administrator's Guide for information about modifying the `httpd.conf` file if you are not managing the server with Enterprise Manager.

2. Start Oracle Application Server Web Cache and its administration tools.
3. Configure Oracle Application Server Web Cache so it receives requests on the default port previously assigned to your Web server (for example, 7777).
4. Configure Oracle Application Server Web Cache so it so it sends cache misses to your newly defined Web server default port number (for example, 7778), which is also referred to as the origin server.
5. Create an Oracle Application Server Web Cache *site* and map the site to your origin server.
6. Apply the changes and restart Oracle Application Server Web Cache.
7. Test the installation to be sure Oracle Application Server Web Cache and your Web server are working properly.

14.8.2.4.3 Enabling End-User Performance Monitoring for Standalone Oracle Application Server Web Cache

After you have installed and configured Oracle Application Server Web Cache and tested the configuration to be sure your Web site data is being cached, you can then enable End-User Performance Monitoring.

The procedure for enabling End-User Performance Monitoring is similar to the procedures documented earlier in this chapter. Use the Oracle Application Server Control for Web Cache 10.1.2 or Oracle Application Server Web Cache Manager for Web Cache 9.0.4 to configure End-User Performance Monitoring, and use Grid Control to start End-User Performance Monitoring, as described in ["Starting and Stopping End-User Performance Monitoring"](#) on page 14-39.

14.8.3 Configuring End-User Performance Monitoring for Web Page Extensions

End User Performance Monitoring feature automatically recognizes all pages with extensions `htm`, `txt`, `jhtml`, `shtml`, `jsp`, and `asp`. However, additional configuration is required if a Web Application has pages with extensions that are not recognized automatically. For example, for Web Applications that have pages with `.do` extension, you will have to make additional configuration so that they get recognized.

To configure end-user performance monitoring for Web page extensions that are not recognized automatically, do these:

1. Access the Web Cache or HTTP Server Home Page.
2. From the Related Links section, click **Monitoring Configuration**.
3. To specify single page extensions, provide the following value in the property **Additional Optional Properties (for EUM)**

```
pageext <appropriate page extension>
```

For example, if the Web page has the extension .do, then provide the following:

```
pageext do
```

To specify multiple page extensions, provide the following value:

```
pageext <appropriate page extension>/<appropriate page extension>
```

For example, if the Web pages have the extensions .do and .html, then provide the following:

```
pageext do/html
```

14.8.4 Configuring End-User Performance Monitoring for Web Pages Having the Same URI

By default, Page Performance reports the performance data for the pages identified by URIs without any query parameters. For example, if the complete URL for petstore search page is /petstore/search?cat=cats, then Page Performance reports data only for /petstore/search.

This works fine if the Web Application pages can be identified by URI uniquely without any query parameters. However, it is not possible to identify the pages if a Web Application has the same URI that is used for all pages. For example, the petstore search page URL /petsore?pageid=search and the petstore cart page URL /petsore?pageid=cart.

To configure end-user performance monitoring for Web pages that have the same URI, do these:

1. Access the Web Cache or HTTP Server Home Page.
2. From the Related Links section, click **Monitoring Configuration**.
3. Provide the following value in the property **Page Identifying Parameters (for EUM)**

```
<query parameter name>
```

For example, if the URI for the petstore search page is /petsore?pageid=search, the specify the following:

```
pageid
```

The query parameters specified can be applicable to all URI paths, or specific to particular URI paths.

For example, if you want all URIs that have a query parameter called 'target' or 'event' to be reported with those query parameters, then specify the following:

```
target,event
```

For example, if you want the URIs that have '/em' as the path and have 'target' or 'event' to be reported with those query parameters, then specify the following:

```
/em:target,event
```

For example, if you want the URIs that have '/em' as the path to be reported, then specify the following:

```
/em/console:event
```

To show how the reported data will look like, here is an example. Consider that the following are the URIs for the application:

```
/portal/page?tab=home&event=login&id=12312312
```

```
/portal/page?tab=home&event=submit&id=553634
```

```
/portal/page?tab=admin&event=update&id=23423234
```

```
/portal/page?tab=admin&event=cancel&id=6784532
```

If you do not specify anything, then you will see one URI, that is, "/portal/page".

If you specify 'tab', then you will see two URIs, that is, "/portal/page?tab=home" and "/portal/page?tab=admin".

If you specify 'tab,event', then you will see four URIs (and EUM data for each), that is, the following:

```
"/portal/page?tab=home&event=login"
```

```
"/portal/page?tab=home&event=submit"
```

```
"/portal/page?tab=admin&event=update"
```

```
"/portal/page?tab=admin&event=cancel"
```

14.8.5 Starting and Stopping End-User Performance Monitoring

After you have configured the Web server to enable collection, you can then start collecting end-user performance data.

1. Navigate to the Web Application home page in the Grid Control console and click **Monitoring Configuration**.
2. Click **Manage Web Server Data Collection**. Enterprise Manager displays the Manage Web Server Data Collection page.
3. In the **Interval (minutes)** column, enter the interval at which Enterprise Manager will collect performance data.
4. Check the **Collection Enabled** checkbox.
5. Click **Apply**, review the changes and confirm by clicking **Apply** again. End-User Performance Monitoring collection is enabled and data will soon be uploaded to the database and shown under the Page Performance page.

To stop collecting end-user performance data:

1. Navigate to the Manage Web Server Data Collection page.
2. Clear the check box in the **Collection Enabled** column of the table and click **Apply**.
3. Click **Apply** again to confirm the changes.

14.8.6 Verifying and Troubleshooting End-User Performance Monitoring

To verify that End-User Performance Monitoring is working properly:

1. Wait a period of time to allow Enterprise Manager to begin collecting end-user performance data and to start loading the data into the Management Repository. Specifically, you should wait until the next upload of data from the Management Agent to the Management Service. The Management Service then loads the data into the Management Repository. For more information about how Enterprise Manager gathers and uploads to the repository, see Oracle Enterprise Manager Concepts.
2. Navigate to the Web Application home page, select a Web application and navigate to the Page Performance tab. Verify that there is data in the **Slowest Response Times** table.
3. Another way to verify the existence of end-user performance data, is to note the value of the **Number of Unprocessed Samples**. Samples for an hour that has not ended are referred to as **Unprocessed Samples**. For example, data is processed for the time period between 10 am to 11 am, 11 am to 12 pm and so on. Therefore, data from 10 am to 11 am will be considered as **Unprocessed Samples** if the 11 am boundary has not been crossed or if there is no incoming end-user traffic after 11 am. If this is a non-zero value, click **Process Samples**. End-user performance data is displayed in the **Slowest Response Times** table.
4. If you still do not see any data on the Page Performance page, consider the following troubleshooting tips:
 - a. Be sure you have completed all the steps required to configure End-User Performance Monitoring. Make sure that the Web server you are using to collect end-user performance data, is either OracleAS Web Cache or Oracle HTTP Server Based on Apache 2.0 (stdApache10.1.2), or Apache HTTP Server (2.0 or higher). You can see the Web server version in the Monitoring Configuration page.
 - b. To monitor Web pages from a third party Application Server, follow the instructions for installing an Apache 2.0 server with the Application Server.
 - c. Install End-User Performance Monitoring after installing the plug-in for the Application Server.
 - When using the Apache Configuration page, log in using the same account used to install Apache.
 - If the Apache server is running on a port less than 1024, the server must be started as root. Apache can be started as root with a lower privileged account by changing ownership of `bin/httpd` to `root` and setting its `setuid` flag. When Apache is started as root, the 'User' and 'Group' directives in `httpd.conf` need to be set to the user who installed the Apache server.

Note: Only pages with a Content-Type header of text or HTML will be monitored. Pages that pass through the Apache Server with a Content-Encoding header (like gzip) will not be monitored because the JavaScript tag cannot be added to these pages.

- If your Web site uses IFrames and End-User Monitoring is not working on those pages, you will need to switch to the newer JavaScript version with

IFrame support. In the `<apache root>/conf/eum.conf` file, follow the directions for enabling IFrame support.

- d. Be sure there is enough activity on your site. If no user is visiting and using your Web application, there may be no end-user performance data to collect or to upload to the Management Repository.
- e. Be sure you have waited long enough for the Management Agent on the Web server host to upload data to the repository. Check the Management Agent home page to determine the last time the Management Agent successfully uploaded data to the Management Repository.
- f. Check the html source of the URLs that you wanted to monitor: make sure the tag `<SCRIPT SRC="/oracle_smp_chronos/oracle_smp_chronos.js"></SCRIPT>` has been appended to the HTML source of these URLs.
 - If it is present, proceed to the next step.
 - If it is not present, check the configuration of your OracleAS Web Cache, Oracle HTTP Server, or Apache HTTP Server. Make sure that all configurations are correct, the site has been enabled, and the Web server has been successfully restarted after saving any configuration changes.
- g. Go to the OracleAS Web Cache or Apache server target home page, click **Monitoring Configuration**, and check if the log file in the defined Log file directory contains any recent data.
 - If it does not have data, go to the next step.
 - If the log file does contain data and the Web server is OracleAS Web Cache, login to Oracle Application Server Control or Web Cache Manager and make sure that the access log is in WCLF or End-User Performance Monitoring format.
- h. Verify that the OracleAS Web Cache / Apache server Monitoring Configuration properties that specify the location and name of the log file are accurate.
- i. Check the Web Server target Home page for any collection errors. Often, the collection error will provide information describing why performance data cannot be collected.
- j. Navigate to the All Metrics page for the Web server target and check to be sure the APM Mining Performance Details metrics are being collected successfully.

14.8.7 Enabling End-User Performance Monitoring for Third-Party Application Servers

For enabling End-User Performance Monitoring for third-party application servers like IBM WebSphere Application Server, BEA WebLogic Managed Server, and JBoss Application Server, after you configure one of the Web servers as explained in this chapter, you have to enable the Application Server Diagnostics Pack for the Web applications hosted on these servers.

To do so, perform the following steps:

1. Click **Setup** on the top-right corner of the Grid Control console and navigate to the Overview of Setup page.
2. Click **Management Pack Access** from the panel to the left.
3. On the Management Pack Access page, select the **All Targets** option in the View Options section of this page.

4. Select **Web Application** from the Search menu, and click **Go**. The table lists all the Web applications monitored.
5. For the Web application for which you want to enable End-User Performance Monitoring, check **Application Server Diagnostics Pack** and **Pack Access Agreed**, and then click **Apply**.
6. Now return to the Web Application Home Page and click **Page Performance** to see the end-user performance monitoring data that has been collected.

Note: End-User Performance Monitoring for a Web application is not supported if the J2EE container hosting that application is SSL enabled. This applies to Oracle J2EE containers, that is OC4J, and any non-Oracle J2EE containers for third-party application servers like BEA WebLogic Managed Server, IBM WebSphere Application Server, or JBoss Application Server. To activate End-User Performance Monitoring for such a Web application, disable SSL for that J2EE container.

For information about configuring SSL for Oracle Application Servers, refer to the Security Guide for your Oracle Application Server release. Documentation for all the Oracle Application releases is available from the Oracle Technology Network:
<http://www.oracle.com/technology/documentation/index.html>.

For information about configuring SSL for third-party servers, refer to your third-party documentation.

14.9 Managing Forms Applications

A Forms Application target in Enterprise Manager can be used to model and monitor a specific Forms application. To use a Forms Application target, you must ensure that the following prerequisites are met:

- Install the Management Agent on the hosts on which the components of your Forms Application have been installed.
- Verify that all the components for your Forms Application has been discovered so that they can be listed as Enterprise Manager targets.
- Create a system that contains all the components that are required for the Forms Application that is to be monitored. The system can contain an Oracle HTTP Server, Apache HTTP Server or an OracleAS Web Cache. For more details on creating a system, refer to [Setting up the System](#).
- After you have created a system for the Forms Application, you can create a Forms Application target using the Create Service Wizard. See [Creating a Service](#) for details. Before you create a service, you must be familiar with the concepts of service management as described in the *Oracle Enterprise Manager Concepts*.

After you have set up the Forms Application target, you can use it to do the following:

- Record and monitor a Forms transaction. See [Recording and Monitoring Forms Transactions](#) for details.
- Measure the End-User Performance of Forms actions such as Commit, Query, Runform, Callform, Openform, and Newform. See [Monitoring the End-User Performance of Forms Applications](#) for details.

14.9.1 Recording and Monitoring Forms Transactions

A Forms transaction consists of a set of user actions within a single application when using Forms. For example, an Update Employee Salary transaction may consist of several user actions like open salary form, update salary form, and save salary form. You can record multiple Forms transactions by using the intuitive playback recorder that automatically records a series of Forms actions.

Before recording a Forms transaction, you must do the following:

- Set the permissions of the `.java.policy` file on each Windows client. See [Setting the Permissions of the .java.policy File](#) on page 14-32.
- Ensure that a trusted Enterprise Manager certificate is used. See [Using a Trusted Enterprise Manager Certificate](#) on page 14-33.
- Add a certificate to the Enterprise Manager Agent to play back secure Forms transactions. See [Adding a Forms Certificate to the Enterprise Manager Agent](#) on page 14-34.
- Configure the Forms server so that Forms transactions can be recorded. See [Configuring the Forms Server](#) on page 14-34.

After you have performed these steps, you can install the transaction recorder to record and play back the Forms transaction. See [Installing the Transaction Recorder to Record and Play Back Forms Transactions](#) on page 14-35.

14.9.1.1 Setting the Permissions of the .java.policy File

You must set the permissions of the `.java.policy` file on each Windows client on which the Forms transaction is being recorded. To set the permissions, follow these steps:

- Ensure that the `.java.policy` file is present under the user home directory. If the `.java.policy` file does not exist, you must create one as follows:
 - Create a `java.policy` (without the ".") file
 - Click **Start** and **Run** from your Windows desktop.
 - Type **cmd** and click **OK**.
 - At the DOS prompt, rename the file as follows:

```
move java.policy .java.policy
```

- After you have created the `.java.policy` file, set the permissions for each Forms server or Oracle Applications server as follows:

```
grant codeBase "URL" {  
    permission java.security.SecurityPermission "putProviderProperty.SunJSSE";  
};
```

where `URL` needs to be replaced with the code source location of the Forms applet. By specifying the `codeBase`, you grant permissions to the code present in that location. For example, for an out-of-box Forms installation, you must specify the `codeBase` as follows:

```
http://formsServerHost:port/forms/java/*
```

where `formsServerHost` and `port` must be replaced with the host name and port number of the Forms server.

For Oracle Applications, you must specify the `codeBase` as follows:

`http://appsHost:appsPort/OA_JAVA/oracle/apps/fnd/jar/*`

where `appsHost` and `appsPort` must be replaced with the host name and port number of the Oracle Applications.

14.9.1.2 Using a Trusted Enterprise Manager Certificate

If you are using secure Enterprise Manager to record a Forms transaction running on Oracle Jinitiator or a Java plug-in, you must ensure that the Enterprise Manager certificate is trusted by Oracle Jinitiator and JPI. For Oracle Jinitiator, you must append the Enterprise Manager certificate to Jinitiator's `certdb.txt` file. For the Java Plug-in, you must set the certificate as trusted by JPI.

To ensure that the Enterprise Manager certificate is trusted by Jinitiator and JPI, follow these steps:

1. Export the Enterprise Manager certificate to a file.
 - When you launch secure Enterprise Manager, if Enterprise Manager is using a self generated certificate, you may see a "Certificate Error". Double click on the error and click **View Certificates**. The Certificate window is displayed.
 - Click the **Details** tab and then click **Copy to file...** to export the certificate to a file. The Certificate Export Wizard is displayed.
 - Click **Next** in the Welcome page.
 - In the Export File Format page, select Base-64 encoded X.509 (.CER) and click **Next**.
 - Click **Browse** to select the name and the location of the file to which the certificate is to be saved.
 - Click **Finish**. The certificate has now been exported to a file.
2. After the certificate has been exported, you must set the certificate as trusted by Jinitiator or JPI.

For Forms applications running on Oracle Jinitiator:

- Open `certdb.txt` under `[Jinitiator InstallRoot]\lib\security\` directory. Usually Jinitiator is installed under `C:\ProgramFiles\Oracle\Jinitiator [version]`.
- Use a text editor to open the file to which the certificate has been exported. Copy the contents and append it to `certdb.txt`.

For Forms applications running on Java plug-in:

- In the Control Panel, double click the Java program that is used to run the Forms application.
 - Click the **Security** tab and then click **Certificates**.
 - From the **Certificate Type** drop down list, select **Secure Site**.
 - Click **Import** to import the file to the location in which the Enterprise Manager certificate has been saved.
 - Close the certificate windows and the Java Control Panel.
3. Close the browser window. When the Forms application is accessed again, Jinitiator or JPI is restarted. This ensures that the changes to the security settings have been saved.

14.9.1.3 Adding a Forms Certificate to the Enterprise Manager Agent

To play back a secure Forms transaction, you must add a Forms certificate to the Enterprise Manager Agent by following these steps:

1. Stop the Management Agent by entering the `emctl stop agent` command.
2. Create an importable certificate file from the forms server certificate (Base64 encoded X.509 format) and name this file as `forms.cer`.
3. Copy the `forms.cer` to `%AGENT_HOME%/jdk/jre/lib/security/` directory.
4. Run `keytool` with the following parameters (the `keytool` executable can be found under the `jdk/jre/bin` directory)

```
keytool -import -alias forms -file %AGENT_HOME%/jdk/jre/lib/security/forms.cer
-keystore AGENT_HOME%/jdk/jre/lib/security/cacerts
```

5. You will be prompted for the `cacerts` password. Enter `changeit` as the password.
6. Start the Management Agent by entering the `emctl start agent` command.

For Forms6i, you need to follow these steps:

1. Stop the Management Agent by entering the `emctl stop agent` command.
2. Obtain forms server certificate in Base64 encoded X.509 format and append to `AGENT_HOME/sysman/config/b64InternetCertificate.txt` file.
3. Start the agent by entering the `emctl start agent` command.

14.9.1.4 Configuring the Forms Server

Before recording a Forms transaction, you must configure the Forms server by following these steps:

1. Create a system based Forms Application target that contains Forms, OracleAS Web Cache or Oracle HTTP Server / Apache HTTP Server targets. These targets must be a part of the system of the Forms Application. They must also be key components of your Forms Application or part of a key Redundancy Group. If you are using the Oracle HTTP Server, the Redundancy Group is referred to as the HTTP Server HA Group.

2. Set up the Forms server for recording transactions:

1. Navigate to the Forms Application Home page in the Grid Control console and click **Monitoring Configuration**.

2. Click **Enable Forms Transaction Monitoring**.

The Enable Forms Transaction Monitoring page is displayed.

3. Select a Forms server from the list and click **Configure**.

The Configure Forms Server: Login page is displayed.

4. Enter the login credentials of the host on which Forms server is installed and click **Continue**.

The jar files required for Forms Transaction Monitoring (`formsRecorder.jar`, `jsse.jar`, `jnet.jar`, and `jcet.jar`) are copied into the Forms applet's archive directory (`ORACLE_HOME/forms/java`) and a confirmation message is displayed.

For Oracle Applications, the archive directory is located at `$JAVA_TOP/oracle/apps/fnd/jar`.

5. Click **Yes** to configure the Forms server and return to the Enable Forms Transaction Monitoring page.

After you have configured the system-based Forms Application target, you can record and play back Forms transactions to monitor the availability of the Forms application. To do so, navigate to the Monitoring Configuration page and click Availability Definition. In this page, change the Availability Definition to Service Test.

14.9.1.5 Installing the Transaction Recorder to Record and Play Back Forms Transactions

After you have configured the Forms server, you can install the transaction recorder on your computer. The transaction recorder is downloaded from the Enterprise Manager Grid Control server the first time you access the Record Forms Transaction page. The transaction recorder requires some Microsoft libraries to be installed in your computer. Make sure that your computer has access to the Internet to download these files. If these libraries are not present during installation, they are automatically downloaded and installed from the Microsoft site. After you have recorded a Forms transaction, if you need to record another one in the same browser, you must use the same JVM version for the new transaction.

You can record multiple Forms transactions on the Forms Application target and monitor these transactions periodically. Before recording a Forms transaction, ensure that all other Forms applications are closed. When you record a Forms transaction, make sure that the following parameters are specified correctly:

- **Login URL:** If you selected the Login Type as **Single Sign-On (SSO)** or **Oracle Applications Login**, the Login URL must be explicitly specified.
- **Connection Type:** This can be:
 - **Socket:** Ensure that the Forms server host name and port number are specified correctly.
 - **HTTP / HTTPS:** If the Connection Type is HTTPS and a non-standard certificate is being used, you must import the certificate into the Agent Home directory.
- **Forms Path:** This is an optional parameter and points to the absolute path of the forms files (. fmx) on the Forms server. To find the absolute path, launch the Forms Application and view the source HTML file of the Forms launcher window. The path is stored in a variable called **xmodule**. Example: The path may be stored as /myvol/oracle01/apps/apps_st/appl/fnd/12.0.0/forms/US/.

Note: This parameter is required only if the Forms transaction has been recorded on one Forms server and played back against a different Forms server with a different installation path.

For more details on recording a Forms transaction and metrics collected, refer to the Enterprise Manager Online Help.

14.9.2 Monitoring the End-User Performance of Forms Applications

The End-User Performance Monitoring utility allows you to measure the response time of your applications by viewing information about how quickly the responses are delivered to the end users. When you access a Forms application, the End-User

Performance Monitoring utility measures the response time of Forms actions such as Commit, Query, Runform, Callform, Newform, and Openform.

You can monitor the Forms actions and view reports based on the response times experienced by the user. You can also define a Watch List of the most important Forms actions to monitor and view the response metrics of these critical operations at a glance.

Note: End-User Performance Monitoring is supported with Forms server version 6i Patch 16, 10g R2. For version 6i Patch 16, only the Commit operation can be monitored.

Before you can begin monitoring the End-User Performance of a Forms Application, you must configure the Forms and Web server to enable data collection for End-User Performance Monitoring. To configure the Forms Application for End-User Performance Monitoring, follow these steps:

- Configure the Forms server to enable End-User Performance Monitoring.
- Configure the Web server (OracleAS Web Cache or Oracle HTTP Server / Apache HTTP Server) so that it can be used for End-User Performance Monitoring.
- Enable the collection of end-user performance data.

14.9.2.1 Configuring the Forms Server for End-User Performance Monitoring

Before you can enable the collection of end-user performance data, you must first configure the Forms server. To configure the Forms server, follow these steps:

1. Navigate to the Forms Application Home page in Enterprise Manager Grid Control.
2. Click **Monitoring Configuration**.
3. Click **Manage Web Server Data Collection**.
4. On the Manage Web Server Data Collection page, select the Forms server and click **Configure**. The Configure Forms Server for End-User Performance Monitoring: Login page is displayed.
5. Enter the host login credentials and click **Continue**. The Configure Forms for End-User Performance Monitoring: Configuration Sections page is displayed.
6. Select a section and check the Enable Monitoring checkbox to enable End-User Performance Monitoring on that section. Click **Enable All** or **Disable All** to enable or disable all the sections. You can also click **Add New Section** to add a section without affecting existing sections. After adding the section, you can enable End-User Performance Monitoring by selecting the checkbox. You can also delete a section that you have added.

Tip: A **section** is a parameter defined in the `formsweb.cfg`. It specifies which section of Forms configuration the user wants to run. The section usually includes the application name and other relevant parameters which are required for successful execution of the application.

7. Set the value of the End-User Performance Monitoring URL column to `http://<hostname:portnumber>/oracle_smp_chronos/oracle_smp_`

`chronos_sdk.gif`. The hostname and port number are for the Web Server that is serving the Forms application.

8. After you have configured the Forms server, click **OK** to save the changes and return to the Manage Web Server Data Collection page.

14.9.2.2 Configuring the OracleAS Web Cache

You can use the 10.1.2 or 9.0.4 versions of OracleAS Web Cache to collect end-user performance data.

- **OracleAS Web Cache 10.1.2:** To configure OracleAS Web Cache 10.1.2, follow these instructions:
 1. You can configure OracleAS Web Cache by using the Oracle Application Server Control. Navigate to the Forms Application home page in the Enterprise Manager Grid Control.
 2. Click **Monitoring Configuration**.
 3. Click **Manage Web Server Data Collection**.
 4. On the Manage Web Server Data Collection page, select the Web Cache target and click **Configure**. The Application Server Control login dialog box is displayed.

Tip: If the login dialog box does not appear or if you receive an error message in your browser window, navigate to the Web Cache Home page and click **Administer** under the Related Links. You will be prompted for the user name and password for Application Server Control. Click **Administration**, scroll down and click **End-User Performance Monitoring**.

If Application Server Control is not available, you can also use the Oracle Application Server Web Cache Manager to configure the OracleAS Web Cache for End-User Performance Monitoring. For more information about starting and using Oracle Application Server Web Cache Manager, refer to the Oracle Application Server Web Cache Administrator's Guide.

5. Enter the username and password for the Web Cache administrator account or the `ias_admin` account. The password for the `ias_admin` account is defined during the installation of Oracle Application Server.

After you have logged into Oracle Application Server Control, you can configure OracleAS Web Cache from the Set Up End-User Performance Monitoring page.
 6. Select the Access Log Format as `access_log:WCLF` for each site from the drop down list. If this format is not in the list, click **Use Required Log Format**.
 7. You will return to the Web Cache Administration page in Oracle Application Server Control. Click **Restart** to restart the Web Cache. For more detailed information about configuring these options, refer to the Enterprise Manager Online Help.
 8. Close the Oracle Application Server Control browser window and return to the Manage Web Server Data Collection page in the Enterprise Manager Grid Control.
- **OracleAS Web Cache 9.0.4:** To configure OracleAS Web Cache 9.0.4, follow these instructions:

1. You can configure OracleAS Web Cache by using the Oracle Application Server Web Cache Manager. Navigate to the Forms Application home page in the Enterprise Manager Grid Control.
2. Click **Monitoring Configuration**.
3. Click **Manage Web Server Data Collection**.
4. On the Manage Web Server Data Collection page, select the Web Cache target and click **Configure**. A login dialog box is displayed.

Tip: If the login dialog box does not appear or if you receive an error message in your browser window, navigate to the Web Cache Home page and click **Administer** under the Related Links. You will be prompted for the user name and password for Application Server Control. Click **Administration**, scroll down and click **End-User Performance Monitoring**.

5. Enter the username and password for the Web Cache administrator account. The first time you log in to the Oracle Application Server Web Cache administrator account, the password is `administrator`.
6. Configure Oracle Application Server Web Cache to use the Web Cache Log Format (WCLF):
 - Select **Logging and Diagnostics** and then select **Access Logs** in the OracleASWeb Cache Manager navigator frame.
 - In the Cache-Specific Access Log Configuration table, click **Edit Selected** and enable the access log for your selected cache.
 - In the Site-Specific Access Log Configuration table, make sure that the Format style of the selected Site Name is `WCLF` and that it is enabled.

For more details on changing the `access_log` format, refer to the Enterprise Manager Online Help.

7. Click **Apply Changes** at the top of the Oracle Application Server Web Cache Manager window and restart Oracle Application Server Web Cache by clicking **Restart** on the Oracle Application Server Web Cache Manager Cache Operations page.
8. Close the Oracle Application Server Web Cache Manager window and return to the Manage Web Server Data Collection page in the Grid Control console. You can now enable the collection of end-user performance data.

14.9.2.3 Configuring the Oracle HTTP Server / Apache HTTP Server

You can collect end-user performance data by using Oracle HTTP Server or Apache HTTP Server. Before you use these server, follow these steps:

1. On the Agent Home page, select the Oracle HTTP Server or Apache HTTP Server target type. If you are using a generic third party Apache server, select a Apache HTTP Server target.
2. Add the target of the corresponding type and make sure that the Log file directory and Log file name properties are set in the Monitoring Configuration page.

The Log file directory and Log file name you specify here will be used by the End-User Performance Mining Engine to upload end-user performance data.

Note: If the Oracle HTTP Server is installed before the Management Agent has been installed, and is up and running during agent installation, then the target will be discovered automatically. Otherwise you need to manually create the Oracle HTTP Server target and specify the following properties: Machine name, Port number, Version of the Apache Server, Oracle home path, Log file directory (for EUM), Log file name (for EUM) where EUM refers to End-User Performance Monitoring.

3. Create a system target and a Forms Application target. Add the Oracle HTTP Server or Apache HTTP Server target to the system target, and make it a key component of the Forms Application target or a part of a key Redundancy Group target. If you are using Oracle HTTP Server, the Redundancy Group is referred to as HTTP Server HA Group.
4. Navigate to the Monitoring Configuration page for the Forms Application target that contains the Oracle HTTP Server or Apache HTTP Server target. Click **Manage Web Server Data Collection**. You will see a table which lists the Web Servers including Oracle HTTP Server, Apache HTTP Server, or OracleAS Web Cache.
5. Select the Oracle HTTP Server or Apache HTTP Server from the table and click **Configure**. Enter the username and password for the host on which the Oracle HTTP Server or Apache HTTP server is installed.
6. After logging in, you will see a table containing the list of sites that are being hosted by the Apache server. These include a list of virtual hosts defined by the user in the Apache Configuration file. The up and the down arrows under the **Monitoring Status** column shows the corresponding site is currently being monitored. For each site, check or uncheck the **Enable Monitoring** checkbox to indicate whether this site is to be monitored. For the site that is to be monitored, enter the log file name in the text box to indicate the location in which the end-user performance data is to be stored. By default, the log file will be created under the `logs/directory` under Apache root directory. To save the log file in a different directory, enter a file name with the absolute path.
7. Make sure that the log file name you specify here matches the Log file directory and Log file name in Monitoring Configuration page of the Oracle HTTP Server or Apache HTTP Server target.
8. You can also use the one button accelerator to enable all sites or disable all sites all at once.
9. After you have made the configuration changes, click **OK** to go to the Apache Restart page. Restarting the Apache server will finalize all configuration changes, and end-user performance data will be logged by the Apache server.
10. After you have configured the Web server, you must configure the Forms server and enable collection of the End-User Performance data from the Manage Web Server Data Collection Page. For details on configuring the Forms server, refer to the Enterprise Manager Online Help.

14.9.2.4 Starting and Stopping End-User Performance Monitoring

After you have configured the Forms and Web server to enable collection, you can then start collecting end-user performance data.

1. Navigate to the Web Application home page in the Grid Control console and click **Monitoring Configuration**.
2. Click **Manage Web Server Data Collection**. Enterprise Manager displays the Manage Web Server Data Collection page.
3. In the **Interval (minutes)** column, enter the interval at which Enterprise Manager will collect performance data.
4. Check the **Collection Enabled** checkbox.
5. Click **Apply**, review the changes and confirm by clicking **Apply** again. End-User Performance Monitoring collection is enabled and data will soon be uploaded to the database and shown under the Page Performance page.

To stop collecting end-user performance data:

1. Navigate to the Manage Web Server Data Collection page.
2. Clear the check box in the **Collection Enabled** column of the table and click **Apply**.
3. Click **Apply** again to confirm the changes.

14.10 Configuring OC4J for Request Performance Diagnostics

Enterprise Manager can gather critical request performance data about your Web application and display this performance data. This feature can be instrumental when you are diagnosing application server and back-end performance issues.

Before you can begin collecting request performance data, you must do the following:

- Create a Web application target and associate it with a system that contains the OC4J instances to be monitored.
- Make these OC4J instances as key system components for your Web application and enable the logging and tracing capabilities. If these OC4J instances are a part of an OC4J Cluster, make sure that this OC4J Cluster is a key system component of your Web application. To enable request performance monitoring, you must configure the specific OC4J instance within the OC4J cluster.

For more information, see the following:

- [Selecting OC4J Targets for Request Performance Diagnostics](#)
- [Configuring Interactive Transaction Tracing](#)
- [Configuring OC4J Tracing for Request Performance Data](#)
- [Additional Configuration for Monitoring UIX Applications](#)

14.10.1 Selecting OC4J Targets for Request Performance Diagnostics

Before you configure the OC4J target to collect request performance data, follow the steps given below to add the target to the Web application.

1. Configure the system where the OC4J targets are defined for the Web application target.
2. Navigate to the Web application Home page and click **Monitoring Configuration**.
3. Click **System Configuration**. From the list of system components displayed on this page, select one or more OC4J targets and select the checkbox in the **Key**

Components column. The OC4J targets can now be configured and used to collect request performance data.

14.10.2 Configuring Interactive Transaction Tracing

When you use transactions to monitor your Web application, some of the transactions you create often involve application components such as servlets, Java Server Pages (JSPs), Enterprise Java Beans (EJBs), and database connections. Often, the best way to solve a performance problem is to trace these more complex transactions and analyze the time spent processing each application component.

Enterprise Manager provides a mechanism for tracing these transactions. Use the **Service Tests and Beacons** link on the **Monitoring Configuration** page of the Web application target to create your transactions and to trace the transactions as they are processed by the servlets, JSPs, EJBs, or database connections of your application.

However, before you can take advantage of transaction tracing, you must first enable tracing for the OC4J instance used to deploy the application. Each OC4J instance of an OC4J cluster must be configured independently. The OC4J instances of the OC4J clusters selected as key components of the Web application target are displayed on the Manage Web Server Data Collection page.

To enable tracing for an OC4J instance:

1. Navigate to the Web Application Home page and click **Monitoring Configuration**.
2. Click **Manage OC4J Data Collection**.

Enterprise Manager displays the Manage OC4J Data Collection page.

3. Select the OC4J to configure and click **Enable Logging**.

Enterprise Manager opens another browser window and displays the Tracing Properties page for the OC4J instance in the Application Server Control.

If you are prompted to log in to the Application Server Control Console, enter the credentials for the `ias_admin` administrator's account.

4. Select the following options on the Tracing Properties page:
 - **Enable JDBC/SQL Performance Details**
 - **Enable Interactive Trace**

You can use the default values for most of the tracing properties.

Note: Turning on the **Enable JDBC/SQL Performance Details** option allows to you drilldown to actual SQL statements but this may require more resources.

5. Click **Apply**.

If this is the first time you are enabling OC4J tracing for this application server, Enterprise Manager displays a message stating that the `transtrace` application is being deployed. The Application Server Control then prompts you to restart the OC4J instance.

6. Click **Yes** to restart the instance and enable the tracing properties.
7. Return to the Grid Control console.

Tracing is now enabled for the selected OC4J instance.

14.10.3 Configuring OC4J Tracing for Request Performance Data

You must configure OC4J instances to enable tracing so that request performance data can be collected. Each OC4J instance of an OC4J cluster must be configured independently. The OC4J instances of the OC4J clusters selected as key components of the Web application target are displayed on the Manage Web Server Data Collection page. To configure the OC4 instances, follow these steps:

1. Navigate to the Web Application home page and click Monitoring Configuration.
2. Click **Manage OC4J Data Collection**.

Enterprise Manager displays the Manage OC4J Data Collection page.

3. For the OC4J instance that you used to deploy your application, select the check box in the **Collection Enabled** column.
4. In the Interval (minutes) column, enter the interval at which to collect OC4J tracing data.

The recommended interval setting is 60 minutes.

5. Select the OC4Js to configure and click **Enable Logging**.

Enterprise Manager opens another browser window and displays the Tracing Properties page for the OC4J instances in the Application Server Control.

If you are prompted to log in to the Application Server Control Console, enter the credentials for the `ias_admin` administrator's account.

6. Select the following options on the Tracing Properties page:

- **Enable JDBC/SQL Performance Details**
- **Enable Historical Trace**

You can use the default values for most of the tracing properties. However, Oracle recommends that you set the **Frequency to Generate Trace File (seconds)** field to 3600 seconds (equivalent to 60 minutes).

Note: Modifying the value in the **Trace File Directory** field is not supported.

7. Click **Apply**.

If this is the first time you are enabling OC4J tracing for this application server, Enterprise Manager displays a message stating that the `transtrace` application is being deployed. The Application Server Control then prompts you to restart the OC4J instance.

8. Click **Yes** to restart the instance and enable the tracing properties.
9. Return to the Grid Control console.

Request Performance data should begin to appear on the Request Performance page as soon as data for the OC4J instance is collected and uploaded into the Management Repository.

14.10.4 Additional Configuration for Monitoring UIX Applications

If you used Oracle User Interface XML (UIX) to build your application, there is an additional configuration step you must perform before you can monitor the requests of your application.

See Also: Your JDeveloper documentation for information on using UIX to develop Web applications

Before you can monitor the requests of your UIX application, you must do the following:

1. Enable tracing for the OC4J instance you used to deploy your application, as described in "[Configuring OC4J Tracing for Request Performance Data](#)" on page 14-42.
2. Locate the following configuration file in the Application Server home directory where you deployed your UIX application:

```
$ORACLE_HOME/j2ee/OC4J_instance_name/config/oc4j.properties
```

For example, if you deployed your application in the OC4J instance called "home," locate the following configuration file:

```
$ORACLE_HOME/j2ee/home/config/oc4j.properties
```

3. Open the `oc4j.properties` file using your favorite text editor and add the following line to the end of the file:

```
oracle.dms.transtrace.dollarstrippingenabled=true
```

4. Save your changes and close the `oc4j.properties` file.
5. Restart the OC4J instance.

14.11 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.

14.11.1 Configuring Service Tests and Beacons

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.

14.12 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**.

14.12.1 Defining Service Level Rules

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.

14.12.2 Viewing Service Level Details

You can view service level information directly from the either of the following:

- **Enterprise Manager Grid 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.

14.13 Configuring a Service Using the Command Line Interface

Using the Command Line Interface, you can define service targets, templates and set up alerts. EM CLI 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. EM CLI can also be used by administrators interactively, and directly from an operating system console. Refer to *Enterprise Manager Command Line Interface Guide* for details.

14.14 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](#)

14.14.1 Verifying and Troubleshooting Forms Transactions

The section covers the following:

- [Troubleshooting Forms Transaction Playback](#)
- [Troubleshooting Forms Transaction Recording](#)
- [Troubleshooting End-User Performance of Forms Transactions](#)

14.14.1.1 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.

14.14.1.2 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, follow the steps in the [Configuring the Forms Server](#) section 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. See [Setting the Permissions of the .java.policy File](#) on page 14-32.
5. If the recording still fails, check if the Enterprise Manager Certificate has been imported to the secure site as described in [Using a Trusted Enterprise Manager Certificate](#). If the certificate has not been imported, you must import it and try recording again. See [Using a Trusted Enterprise Manager Certificate](#) on page 14-33.

14.14.1.3 Troubleshooting End-User Performance of Forms Transactions

This section lists troubleshooting steps that you can use when the Forms transaction End-user Performance Monitoring (EUM) data is not being displayed.

1. Ensure that the Forms server is configured with EUM.

From the **Manage Web Server Data Collection** page, select the Forms server and click **Configure**. Log in using the credentials of the host where the Forms server is installed. Ensure that the correct Forms configuration section has been configured with EUM enabled and that the correct EUM URL is specified. Go to the Forms application URL (with the correct configuration section) and perform "Save" or "Query" actions to generate EUM traffic.

2. Ensure that the Web server is configured to log End-User Performance Monitoring data.

From the **Manage Web Server Data Collection** page, select the Web server and click **Configure**.

If you are using a Web Cache to log EUM data, login to the **Web Cache Administration** page or Web Cache Manager and check if the `access_log` file is set to either End-User Performance Monitoring or WCLF format. End-User Performance Monitoring data is logged into Web Cache's `access_log`.

If you are using HTTP Server or Apache HTTP Server, log in using the credentials of the host where the HTTP Server is installed. Then check if EUM has been enabled and note the path of the log file in the configuration page.

3. Ensure that the EUM log file is being generated.

Go to the location of the End-User Performance Monitoring log file, open the log file and search for word "sdk".

"sdk" entries indicate that there is EUM traffic and that the monitoring configuration is correct. In this situation, more time is required to collect end-user performance data. If the log file exists and "sdk" entries are found, go to step 4.

4. Check the Monitoring Configuration page of the Web Cache or HTTP Server target to ensure the parameters "Log File Directory (for EUM)" and "Log File Name (for EUM)" match that of the log file path shown on the configuration page.

5. Another way to verify the existence of end-user performance data, is to note the value of the **Number of Unprocessed Samples** on the **Page Performance** page of the Forms application. Samples for an hour that has not ended are referred to as **Unprocessed Samples**. For example, data is processed for the time period between 10 am to 11 am, 11 am to 12 pm and so on. Therefore, data from 10 am to 11 am will be considered **Unprocessed Samples** if the 11 am boundary has not been crossed or if there is no incoming end-user traffic after 11 am. If this is a non-zero value, click **Process Samples**. End-user performance data is displayed in the **Slowest Response Times** table.

14.14.2 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) OracleEMAgentURLTiming/3.0.

Note: For Grid Control 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.

Extending Enterprise Manager

Enterprise environments consist of a wide variety of components: OS platforms, hardware, software, network, and storage devices. All of these components work in concert to deliver critical information and functionality required to keep enterprise operations performing optimally and providing information to make important business decisions. While Oracle Enterprise Manager Grid Control allows you to monitor and manage a variety of components out-of-box, you may want to monitor third party components or custom applications specific to your environment. For example, with this release, you can seamlessly monitor WebLogic and WebSphere application servers. Additional plug-ins are being developed and will be announced as they become available.

In addition, you can use the same mechanism used by Oracle and partners to extend Enterprise Manager to monitor custom components via modular Management Plug-ins. Once a plug-in is defined, you use the Enterprise Manager Grid Control console to deploy the new plug-in throughout your enterprise environment.

15.1 Benefits of Extending Enterprise Manager

Extending Enterprise Manager for monitoring additional components provides the following benefits:

- **Centralize management information in a single console:** When deployed, components defined by the Management Plug-in automatically appear in the Grid Control console. Being able to monitor all targets in your environment provides a consolidated view of your entire enterprise, thus allowing you to monitor and manage all components from a central point.
- **Extend Enterprise Manager's monitoring and management features to non-Oracle components:** Newly added components automatically inherit Enterprise Manager's powerful monitoring and management features, such as: alerts, policies, blackouts, monitoring templates, groups/systems, configuration management, and enterprise reporting.
- **Comprehensive Service-Level Management:** By managing all of your enterprise components with Enterprise Manager Grid Control, you can more fully utilize the Service-Level Management features offered in Enterprise Manager Grid Control.

15.2 More Extensibility Information

Complete Management Plug-in development information can be found in the *Oracle Enterprise Manager Extensibility Guide*.

The latest information on all aspects of Oracle Enterprise Manager extensibility can be found at the Oracle Enterprise Manager Grid Control Extensions Exchange located on the Oracle Technology Network web site.

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

Part III

Enterprise Manager High Availability

This section covers high availability best practices and strategies that allow you to safeguard your Oracle Enterprise Manager installation.

Part III contains the following chapters:

- [Chapter 16, "High Availability Solutions"](#)
- [Chapter 17, "High Availability: Single Resource Configurations"](#)
- [Chapter 18, "High Availability: Multiple Resource Configurations"](#)
- [Chapter 19, "Management Agent and Management Services"](#)

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 Grid 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 Grid 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 Services
Results in targets no longer monitored by Enterprise Manager, though the Enterprise Manager console is still available and one can view historical data from the Management Repository.
- Management Service failure
Results in the unavailability of Enterprise Manager console, as well as unavailability of almost all Enterprise Manager services.
- Management Repository failure
Results in failure on the part of Enterprise Manager to save the uploaded data by the Management Agents as well as unavailability of almost all Enterprise Manager services.

Overall, failure in any component of Enterprise Manager can result in substantial service disruption. Therefore it is essential that each component be hardened using a highly available architecture.

16.1 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 Website
<http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>
- Support Note 330072.1: "How To Configure Grid Control Components for High Availability "

16.2 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.

16.2.1 Levels of High Availability

Each high availability solution level is driven by your business requirements and available IT resources. The following table summarizes the four high availability levels for Oracle Enterprise Manager installations.

Table 16–1 Enterprise Manager Availability levels

High Availability Level	Business Need	Hardware Requirement
Level 1	Responsiveness to business application events.	Well-tuned, single instance (one host) with redundant storage. 11.1.0.7 RDBMS + Protected Storage
Level 2	Ability to ensure business application service quality.	Cold Failover Cluster configuration using Data Guard on a single site. Level 1 + Data Guard
Level 3	Operational overhead and heavy costs of manual processes.	<i>n</i> -Instance RAC with local site Data Guard (protection against limited site failure). Level 2 + Primary site on RAC
Level 4	Revenue impact on loss of key business services and applications.	<i>n</i> -Instance RAC with secondary side Data Guard (protection against site loss) Level 3 + Data Guard on a remote site.

Note: Levels 3 and 4 are not covered in this manual. For more information, refer to Real Application Cluster and Dataguard documentation.

16.3 Determining Your High Availability Needs

As previously mentioned, the availability level you choose depends on factors such as the hardware resources available and the business need of your organization. However, developing your high availability plan in a way that objectively encompasses all aspects of your high availability needs (hardware, business processes, effort, cost) can be problematic. The solution is to define high availability needs in terms of Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

- **Recovery Time Objective** - The period of time within which your business process or technological resources must be restored after failure. Key Question: How fast do your business processes/resources need to be running again before the bottom line is impacted?
- **Recovery Point Objective** - The period of time between the time of failure and the last backup. Key Question: How much data are you willing to lose?

Defining your high availability needs in terms of RTO and RPO allows you to effectively meet the demands of users. Both values should be determined using the worst-case scenarios.

16.4 RTO, RPO, and Availability Levels

Given the broad range of factors that must be taken into consideration when implementing a highly available Enterprise Manager environment, your ultimate decision will be based on the interrelationship between RTO, RPO and the cost involved with implementing one of the availability levels. The following table shows the interrelationship between these factors.

Table 16–2 Comparison of High Availability Levels

Level	RTO	RPO	Build Time	Cost
1	98.0%	Hours	Hours to Days	\$
2	98.8%	Minutes	Hours to Days	\$\$
3	99.9%	Minutes to Seconds	Days	\$\$\$
4	99.9%	Minutes to Seconds	Days	\$\$\$\$

The table is not a prescriptive recommendation for choosing a high availability level, but instead should be used to aid your decision making process based on your business needs. For example, you have an uptime requirement of 95% and a desired mean time to recovery of seconds, the you should select level four.

What is not reflected in the table are such factors as survivability and scalability. Hence, although the differences between level three and level four seem outwardly insignificant, there are differences. If you need survivability in the event of a primary site loss you need to go with a Level 4 architecture. If you need equalized performance in the event of site loss it's essential. A level three architecture with DG that's asymmetrically scaled will mean degradation in performance when activated.

If you need to maintain performance levels you will need for level 4 with a symmetrically sized architecture on both sites. 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.

High Availability: Single Resource Configurations

Single resource configurations consist of a single instance Enterprise Manager configuration utilizing some form of protected storage to protect both the repository database and the software installation. As one of the most common installation types, implementing high availability for single resource configurations is cost effective from both time and resource standpoints as the objective is to leverage the technology that is already available, such as Recovery Manager, Flashback, Ultrasafe, and Automated Storage Management.

This chapter covers the following topics:

- [About Single Resource Configurations](#)
- [Deploying Grid Control Components on a Single Host](#)
- [Backup and Recovery](#)

17.1 About Single Resource Configurations

The configurations described in this chapter are provided as examples only. The actual Grid Control configurations that you deploy in your own environment will vary depending upon the needs of your organization.

For example, in a production environment you will likely want to implement firewalls and other security considerations. For specific information on implementing firewalls and security protocols, you should refer to the following Enterprise Manager documentation:

- [Chapter 2, "Enterprise Manager Security"](#) for information about securing the connections between Grid Control components
- For information about configuring firewalls between Grid Control components, see the *Oracle® Enterprise Manager Grid Control Advanced Installation and Configuration Guide*.

Besides providing a description of common configurations, this chapter can also help you understand the architecture and flow of data among the Grid Control components. Based on this knowledge, you can make better decisions about how to configure Grid Control for your specific management requirements.

The Grid Control architecture consists of the following software components:

- Oracle Management Agent
- Oracle Management Service

- Oracle Management Repository
- Oracle Enterprise Manager 11g Grid Control Console

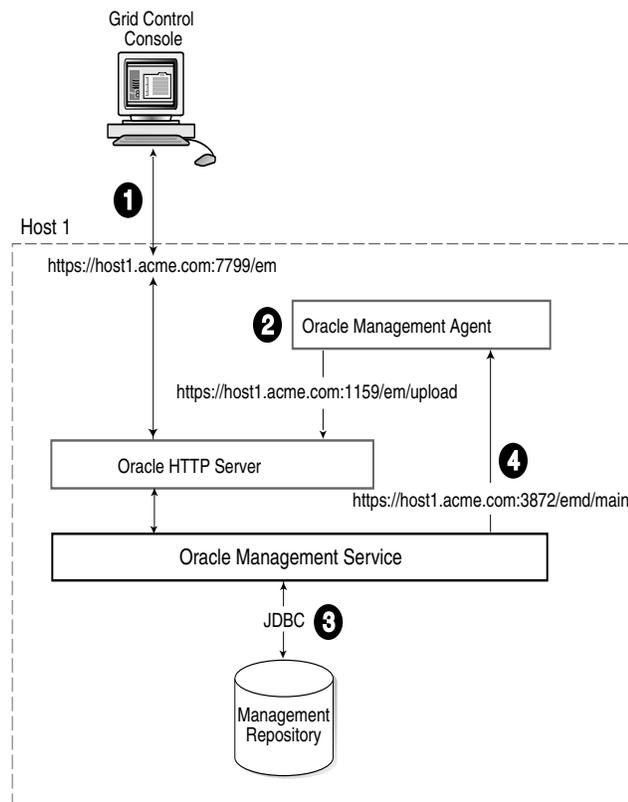
See Also: *Oracle Enterprise Manager Concepts* for more information about each of the Grid Control components

17.2 Deploying Grid Control Components on a Single Host

Figure 17–1 shows how each of the Grid Control components are configured to interact when you install Grid Control on a single host. This is the default configuration that results when you use the **Installing a New Grid Control** installation type.

In Enterprise Manager release 11g, the installation does not create a new database. You must install a database which should be on the same host as Enterprise Manager.

Figure 17–1 *Grid Control Components Installed on a Single Host*



When you install all the Grid Control components on a single host, the management data travels along the following paths:

1. Administrators use the Grid Control console to monitor and administer the managed targets that are discovered by the Management Agents on each host. The Grid Control console uses the following URL to connect to the Oracle HTTP Server:

`https://host1.acme.com:7799/em`

The Management Service retrieves data from the Management Repository as it is requested by the administrator using the Grid Control console.

2. The Management Agent loads its data (which includes management data about all the managed targets on the host, including the Management Service and the Management Repository database) by way of the Oracle HTTP Server upload URL. The Management Agent uploads data directly to Oracle HTTP Server. The default port for the upload URL is 1159 (if it is available during the installation procedure). The upload URL is defined by the `REPOSITORY_URL` property in the following configuration file in the Management Agent home directory:

```
AGENT_HOME/sysman/config/emd.properties (UNIX)
AGENT_HOME\sysman\config\emd.properties (Windows)
```

See Also: For more information about the Oracle Enterprise Manager directory structure (`AGENT_HOME` directory in particular), see the *Oracle® Enterprise Manager Grid Control Advanced Installation and Configuration Guide*.

3. The Management Service uses JDBC connections to load data into the Management Repository database and to retrieve information from the Management Repository so it can be displayed in the Grid Control console. The Management Repository connection details can be listed and changed by using the following `emctl` commands:

```
emctl config oms -list_repos_details
emctl config oms -store_repos_details
```

See Also: ["Reconfiguring the Oracle Management Service"](#) on page 19-9 for more information on modifying the Management Repository connection information.

4. The Management Service sends data to the Management Agent by way of HTTP. The Management Agent software includes a built-in HTTP listener that listens on the Management Agent URL for messages from the Management Service. As a result, the Management Service can bypass the Oracle HTTP Server and communicate directly with the Management Agent. If the Management Agent is on a remote system, no Oracle HTTP Server is required on the Management Agent host.

The Management Service uses the Management Agent URL to monitor the availability of the Management Agent, submit Enterprise Manager jobs, and other management functions.

The Management Agent URL can be identified by the `EMD_URL` property in the following configuration file in the Management Agent home directory:

```
AGENT_HOME/sysman/config/emd.properties (UNIX)
AGENT_HOME\sysman\config\emd.properties (Windows)
```

For example:

```
EMD_URL=https://host1.acme.com:3872/emd/main
```

In addition, the name of the Management Agent as it appears in the Grid Control console consists of the Management Agent host name and the port used by the Management Agent URL.

17.3 Backup and Recovery

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)
- Oracle Management Agent (Agent)
- Oracle Management Repository (Repository)

Each component, being uniquely different in composition and function, requires different approaches to backup and recovery. For this reason, the backup and recovery strategies are discussed on a per-tier basis in this chapter. For an overview of Enterprise Manager architecture, refer to the Oracle® Enterprise Manager Grid Control Basic Installation Guide.

Oracle Configuration Manager

Oracle Configuration Manager (OCM) is used to collect client configuration information and upload it to the Oracle repository. When the client configuration data is uploaded on a regular basis, customer support representatives can analyze this data and provide better service to the customers.

When installing Oracle software, the installer provides an option to setup and configure OCM. In recovery scenarios where software is installed in software-only mode, OCM can be configured manually by running the following from OMS and agent Oracle Homes:

```
<OracleHome>/ccr/bin/setupCCR
```

The Oracle Configuration Manager client is installed into the ORACLE_HOME directory. Once installed, OCM collects configuration data related to the ORACLE_HOME directory and the host on which it is installed. In addition to collecting and uploading configuration data, it also checks if any software updates to the Oracle Configuration Manager client are available. If updates are available, it downloads them and updates the Oracle Configuration Manager software installed on the customer's system.

17.3.1 Repository Backup and Recovery

The Repository is the storage location where all the information collected by the 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 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 Enterprise Manager console.

17.3.1.1 Repository Backup

Oracle recommends using High Availability Best Practices for protecting the 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 Enterprise Manager 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.

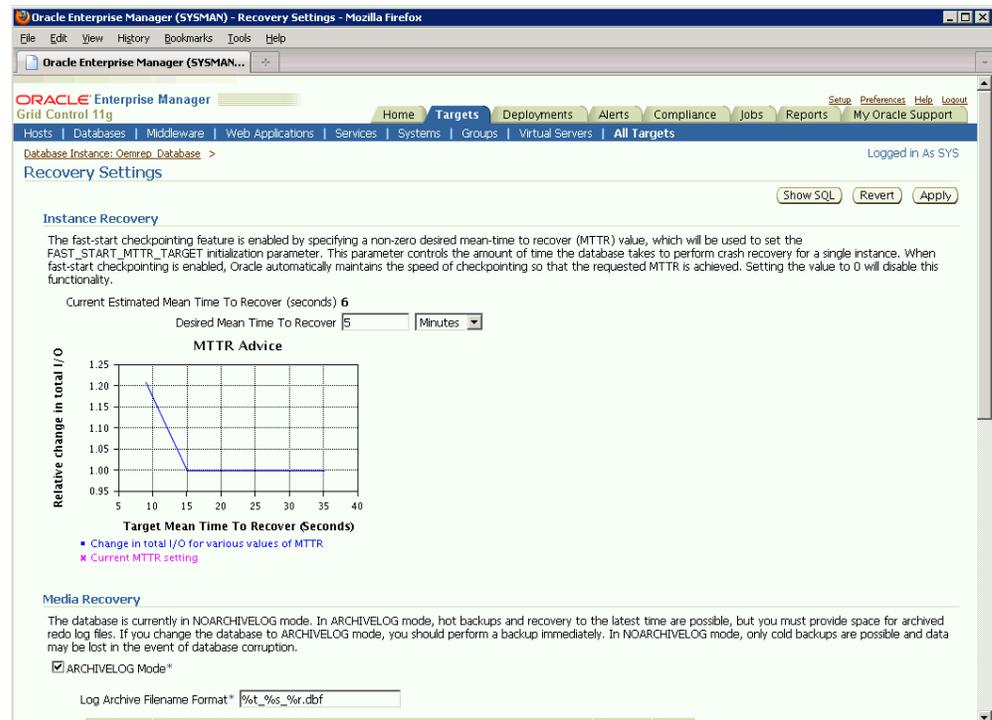
Using the *Recommended Backup Strategy* also takes advantage of the capabilities of Enterprise Manager to execute the backups: Jobs will be automatically scheduled through the Job sub-system of Enterprise Manager. The history of the backups will then be available for review and the status of the backup 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.

Setting Up the Backup

First, navigate to the Enterprise Manager Recovery Settings page (Target-->Database--><Repository Database Target>-->Availability-->Recovery Settings) and enable *Archive Logging* then *Flashback Database* as shown in Figure 17-2.

Figure 17-2 Recovery Settings Page



Next, navigate to the Backup Policies page (Target-->Database--><Repository Database Target>-->Availability-->Backup Settings-->Policy) and enable *Block Change Tracking* to speed up backup operations as shown in Figure 17-3.

Figure 17-3 Backup Policy Page

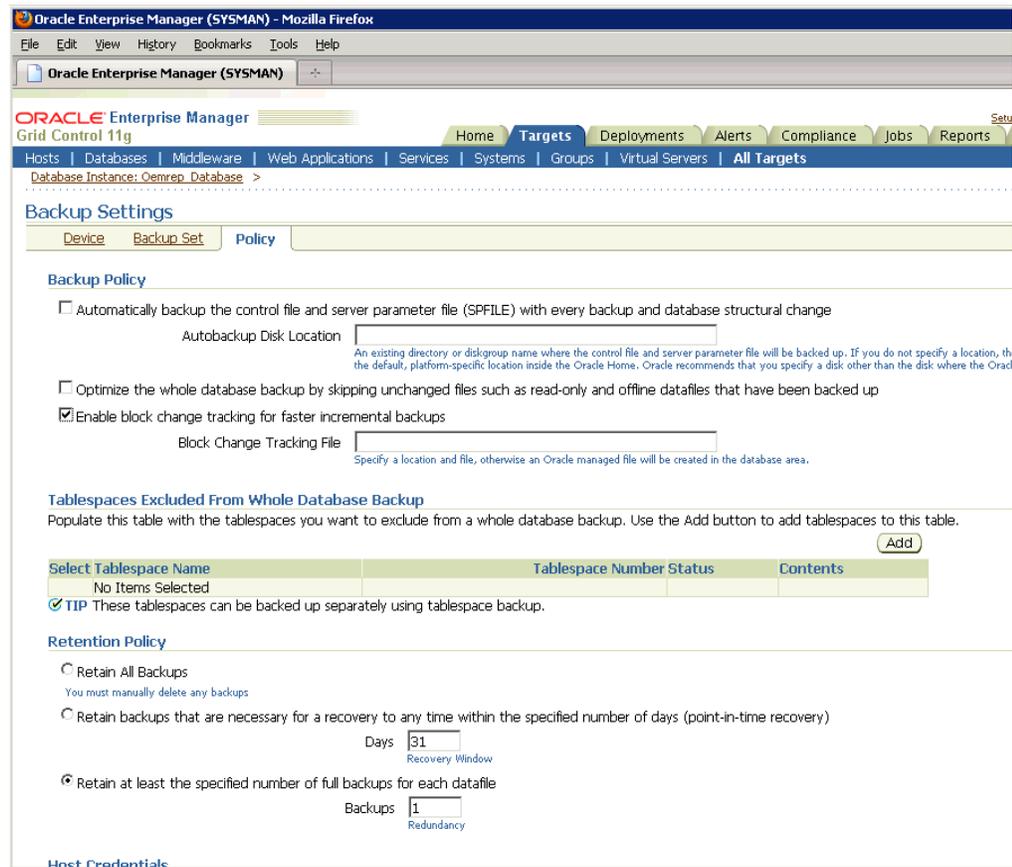
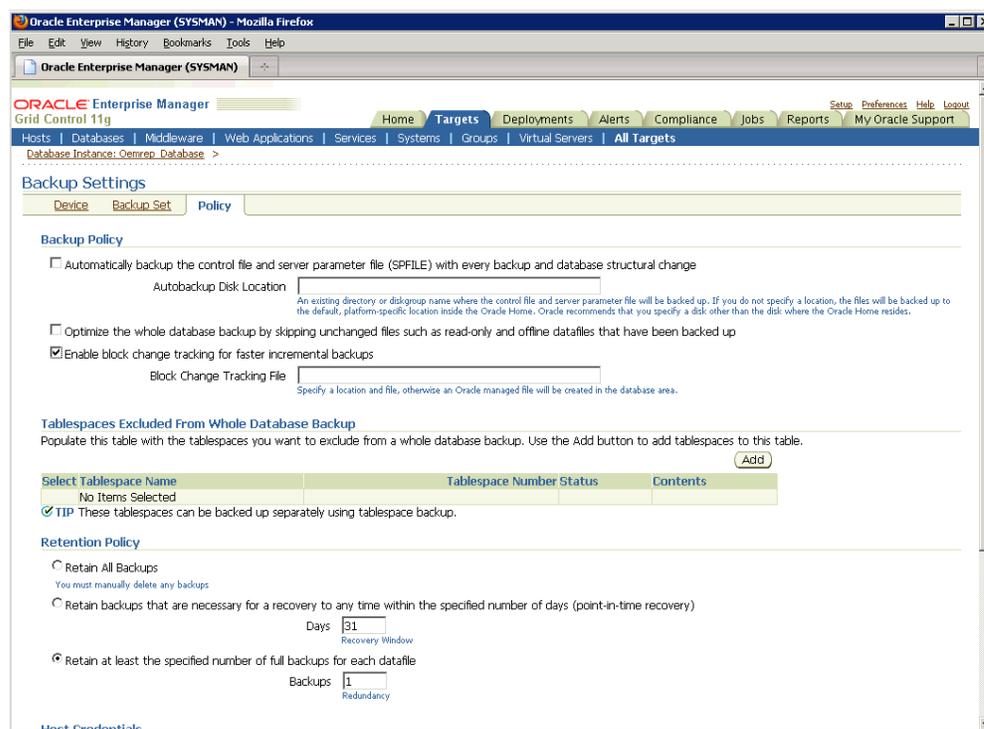


Figure 17-4 Backup Policy Page



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.

17.3.1.2 Repository Recovery

Recovery of the Repository database must be performed using RMAN since Grid 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 Grid Control console unless the repository is synchronized with the latest state available on the Agents.

A repository resync feature (Enterprise Manager version 10.2.0.5 and later) allows you to automate the process of synchronizing the Enterprise Manager repository with the latest state available on the Agents.

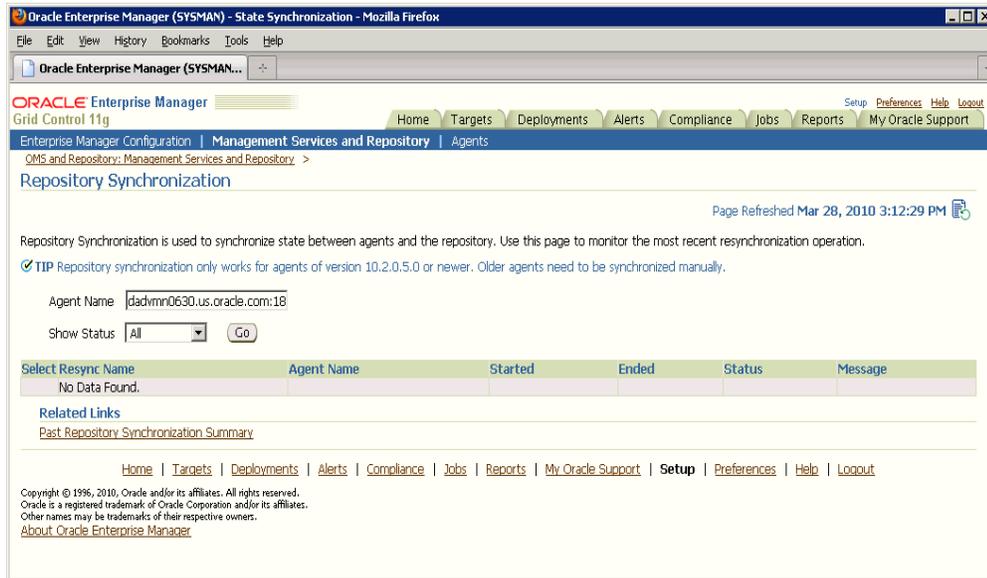
Note: resync requires Agents version 10.2.0.5 or later. Older Agents must be synchronized manually. See "[Manually Resynchronizing Agents](#)" on page 17-8.

To resynchronize the repository with the 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 repository but BEFORE starting the OMS. After submitting the command, start up all OMS's and monitor the progress of repository resynchronization from the Enterprise Manager console's Repository Resynchronization page, as shown in [Figure 17-5](#).

Figure 17-5 Repository Synchronization Page



Repository recovery is complete when the resynchronization jobs complete on all Agents.

Oracle strongly recommends that the 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:

- Notification Rules
- Preferred Credentials
- Groups, Services, Systems
- Jobs/Deployment Procedures
- Custom Reports
- New Agents

Manually Resynchronizing Agents

The Enterprise Manager Repository Synchronization feature can only be used for Agents 10.2.0.5 or later. Older Agents must be resynchronized manually by shutting down the Agent using the following procedure:

1. Shut down the Agent.

2. Delete the `agentstmp.txt`, `lastupld.xml`, `state/*` and `upload/*` files from the `<AGENT_HOME>/sysman/emd` directory.
3. Restart the Agent.

17.3.1.3 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.

17.3.1.3.1 Full Recovery on the Same Host Repository database is running in *archive* 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 the OMS(s) using `emctl stop oms`.
2. Recover the database using RMAN
3. Bring the site up using the command `emctl start oms` on all OMS(s).
4. Verify that the site is fully operational.

17.3.1.3.2 Incomplete Recovery on the Same Host Repository database is running in *noarchive* mode. Full offline backup is available. The repository database disk crashes. All datafiles and control files are lost.

Resolution:

1. Stop the OMS(s) using `emctl stop oms`.
2. Recover the database using RMAN.
3. Initiate Repository Resync using `emctl resync repos -full -name "<resync name>"` from one of the OMS Oracle Home.
4. Start the OMS(s) using `emctl start oms`.
5. Manually fix any pre-10.2.0.5 Agent by shutting down the Agent, deleting the `agentstmp.txt`, `lastupld.xml`, `state/*` and `upload/*` files under the `<AGENT_HOME>/sysman/emd` directory, and then restarting the Agents.
6. Log into Grid Control. Navigate to **Management Services and Repository Overview** page. Click on **Repository Synchronization** under **Related Links**. Monitor the status of resync jobs. Resubmit failed jobs, if any, after fixing the error.

7. Verify that the site is fully operational.

17.3.1.3.3 Full Recovery on a Different Host The 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(s) 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 in credential store by running

```
$emctl config oms -store_repos_details -repos_conn_desc <connect descriptor>
-repos_user sysman
```
4. Start the OMS(s) using the command `emctl start oms`.
5. Relocate the 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.

If a new Agent had been installed on host "B", you must ensure there are NO previously discovered database targets.

6. Change the monitoring configuration for the OMS and Repository target: by running the following command from the OMS:

```
$emctl config emrep -conn_desc "<TNS connect descriptor>"
```
7. Verify that the site is fully operational.

17.3.1.3.4 Incomplete Recovery on a Different Host The 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(s) 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_conn_desc <connect descriptor>
-repos_user sysman
```
4. Initiate Repository Resync:

```
emctl resync repos -full -name "<resync name>"
```

from one of the OMS Oracle Homes.

5. Start the OMS(s) using the command `emctl start oms`.
6. Run the command to relocate the repository database target to the 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. Manually fix all pre-10.2.0.5 Agents by shutting down the Agents, deleting the `agentstmp.txt`, `lastupld.xml`, `state/*` and `upload/*` files under the `<AGENT_HOME>/sysman/emd` directory and then restarting the Agents.
9. Log in to Grid Control. Navigate to **Management Services and Repository Overview** page. Click on **Repository Synchronization** under **Related Links**. Monitor the status of resync jobs. Resubmit failed jobs, if any, after fixing the error mentioned.
10. Verify that the site is fully operational.

17.3.2 Oracle Management Service Backup and Recovery

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. For Enterprise Manager version 11.1, the OMS architecture has changed.

17.3.2.1 Backing Up the OMS

The OMS is generally stateless. Some transient and configuration data is stored on the OMS file system. The shared loader "recv" directory stores metric data uploaded from Agents temporarily before the data is loaded into the repository. If an OMS goes down, other surviving OMS(s) upload the data stored in the shared loader location. In a High Availability (HA) configuration, the shared loader receive directory should be protected using an HA storage technology, such as a redundant disk.

A snapshot of OMS configuration can be taken using the `emctl exportconfig oms` command.

```
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.
                          For example: ORACLE_HOSTNAME
```

Note: The `exportconfig oms` command is only available with Enterprise Manager version 10.2.0.5 or newer.

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. 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 three WebLogic components – Middleware Home, the OMS Oracle Home and the WebTier (OHS) Oracle Home. Software Homes only change when patches or patchsets are applied. For this reason, filesystem-level backups should be taken after each patch/patchset application. You should back up the Oracle inventory files along with the Software Homes. .

Important: Beginning with Enterprise Manager version 11.1, the location of the OMS Oracle Home must be the same for all OMS's in your monitored environment.

- **Instance Home**

Composed of WebLogic, OMS and WebTier configuration files. The Instance Home can be backed up using the `emctl exportconfig oms` command.

- **Software Library**

Composed of components used by Enterprise Manager patching and provisioning functions. Oracle Database Filesystem (DBFS) is recommended for software library backup. DBFS technology allows an Oracle database tablespace to be exposed to applications as a mounted filesystem. Internally, all the files are stored as secure files in the Oracle database. Storing the software library in the Enterprise Manager repository database using DBFS lets you leverage the comprehensive capabilities of the Oracle database to take consistent backups of the software library along with the Enterprise Manager repository. For more information about DBFS, see the *Oracle® Database SecureFiles and Large Objects Developer's Guide*.

- **Shared Loader RECV Directory**

The shared loader receive (RECV) directory temporarily stores metric data uploaded from Agents before the data is loaded into the repository. Use a high availability storage technology to protect the receive directory.

- **AdminServer**

Beginning with Enterprise Manager version 11.1, the OMS's WebLogic architecture introduces the concept of an AdminServer. The AdminServer operates as the central control entity for the configuration of the entire OMS(s) domain. The AdminServer is an integral part of the first OMS installed in your Grid Control deployment and shares the Software Homes and Instance Home.

17.3.2.2 Recovering the OMS

If an OMS is lost, it should be reinstalled using "Installing Software Only and Configuring Later". This is an additional Management Service option documented in the *Oracle Enterprise Manager Grid Control Advanced Installation and Configuration* guide. The OMS configuration can then be restored with the OMS Configuration Assistant using the following command:

```
omsca recovery -BACKUP_FILE <file>
```

Use the export file generated by the `emctl exportconfig` command shown in the previous section.

Recovering an OMS essentially consists of two steps, recovering the Software Homes and then configuring the Instance Home. When restoring on the same host, the

software homes can be restored from filesystem backup. In case a backup does not exist, the software homes can be reconstructed using the software-only installation of WebLogic and OMS, software-only installation of add-ons (if any) and all patches that were applied before the crash. As stated earlier, the location of the OMS Oracle Home is fixed and cannot be changed. Hence, ensure that the OMS Oracle Home is restored in the same location that was used previously.

Once the Software Homes are recovered, the instance home can be reconstructed using the `omsca` command in recovery mode.

17.3.2.3 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.

17.3.2.3.1 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. Ensure that loader receive directory and software library locations are still accessible.
2. Restore the software homes from filesystem backup taken earlier. Alternately, if a backup does not exist, use the software-only install method to reconstruct the WebLogic and OMS Oracle Home, add-ons that were installed earlier need to be reinstalled in software-only mode and all patches that were applied earlier need to be reapplied. Remember that the location of OMS Oracle Home needs to be the same as one used before.
3. Run `omsca` in recovery mode specifying the export file taken earlier to configure the OMS:

```
omsca recover -as -ms -backup_file <file>
```

Note: The `-backup_file` to be passed must be the latest file generated from `emctl exportconfig oms` command.

4. Configure agent:

```
>agentca -f
>emctl secure agent -emdWalletSrcUrl <oms url>
```

5. At this point, two possibilities exist depending upon the port used by the reinstalled agent that comes along with the OMS:

Option A: Agent uses the same port as the previous installation.

- OMS automatically blocks the Agent. Resync the Agent from Agent homepage

Option B: Agent uses a different port.

- Run the command to relocate the OMS and Repository target to reinstalled Agent:

```
emctl config emrep -agent <reinstalled agent>
```

Example: `emctl config emrep -agent foo.us.oracle.com:3872`

6. Locate duplicate targets from the Management Services and Repository Overview page. Relocate duplicate targets from the old agent to the reinstalled Agent. Delete the old Agent.
7. Verify that the site is fully operational.

17.3.2.3.2 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 loader receive directory and software library locations are accessible from Host "B".
2. Usually filesystem restore does not work across hosts. Use the software only install method to reconstruct the WebLogic and OMS Oracle home, add-ons that were installed earlier need to re-installed in software only mode and all patches that were applied earlier need to be reapplied. Remember that the location of OMS Oracle Home must be the same as one used before.
3. Run `omsca` in recovery mode specifying the OMS configuration backup file generated earlier to configure the new OMS:

```
omsca recover -as -ms -backup_file <file>
```

4. Configure agent:

```
agentca -f
emctl secure agent -emdWalletSrcUrl <oms url>
```

5. Change the OMS to which all Agents point and then resecure all Agents
 - Make all Agents in the deployment point to new OMS running on Host "B". On each Agent, run the following command

```
emctl secure agent -emdWalletUrlSrc
"http://hostB:<httpport>/em"
```

- Run the command to relocate OMS and Repository target to Agent "B":

```
emctl config emrep -agent <agent on host "B">.
```

Note: 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.

6. 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".

7. Verify that the site is fully operational.

17.3.2.3.3 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.

Resolution:

1. Ensure that loader receive directory and software library locations are accessible from Host "B".
2. Usually filesystem restore does not work across hosts. Use the software-only install method to reconstruct the WebLogic and OMS Oracle home, add-ons that were installed earlier need to be re-installed in software-only mode and all patches that were applied earlier need to be reapplied. Remember that the location of OMS Oracle home needs to be the same as one used before.
3. Modify the network configuration such that HostB also responds to hostname 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 Host "B" and Host "A" network addresses resolve to the physical IP of Host "B".
- Multi-home Host "B". Configure an additional IP of Host "A" on Host "B". For example, on Host "B" run the following commands:

```
> ifconfig eth0:1 <IP of HostA> netmask <netmask>
> /sbin/arping -q -U -c 3 -I eth0 <IP of HostA>
```

4. Run `omscs` in recovery mode specifying the export file taken earlier to configure the OMS:

```
omscs recover -as -ms -backup_file <file>
```

5. Resecure the OMS:

```
emctl secure oms host <Host A>
```

6. Configure agent:

```
> agentca -f followed by
> emctl secure agent -emdWalletSrcUrl <oms url>
```

7. Run the command to relocate Management Services and Repository target to Agent "B":

```
emctl config emrep -agent <agent on host B>
```

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. Verify that the site is fully operational.

17.3.2.3.4 Multiple OMS, Server Load Balancer, Primary OMS Recovered on the Same Host Site hosts multiple OMSs. All OMSs 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. Ensure that shared loader receive directory and shared software library locations are still accessible.
2. Restore the software homes from filesystem backup taken earlier. Alternately if backup does not exist, use the software only install method to reconstruct the WebLogic and OMS Oracle home, add-ons that were installed earlier need to re-installed in software only mode and all patches that were applied earlier need to be reapplied. Remember that the location of OMS Oracle home needs to be the same as one used before.
3. Run `omsca` in recovery mode specifying the export file taken earlier to configure the OMS:

```
omsca recover -as -ms -backup_file <file>
```

4. Resecure the Agent that gets installed along with OMS.

```
emctl secure agent -emdWalletSrcUrl  
"http://slb:<httpport>/em"
```

5. At this point, two possibilities exist depending upon the port used by the reinstalled agent that comes along with the OMS:
 - Option A: Agent gets the same port as earlier--OMS automatically blocks the agent. Resync the agent from agent homepage.
 - Option B: Agent gets a different port--Run the command to relocate Management Services and Repository target to reinstalled agent:

```
emctl config emrep -agent <reinstalled agent>
```

Locate duplicate targets from the Management Services and Repository Overview page. Relocate duplicate targets from old agent to reinstalled agent. Delete the old agent.

6. Re-enroll the additional OMS, if any, with the recovered Administration Server by running `emctl enroll oms` on each additional OMS.
7. Verify that the site is fully operational.

17.3.2.3.5 Multiple OMS, Server Load Balancer configured, Primary OMS Recovered on a Different Host Site hosts multiple OMSs. OMSs 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. Ensure that shared loader receive directory and shared software library locations are accessible from the new OMS host (host "B")
2. Filesystem restore typically does not work across hosts. Use the software-only install method to reconstruct the WebLogic and OMS Oracle home. Add-ons that were installed earlier need to re-installed in software-only mode and all patches that were applied earlier need to be reapplied. Remember that the location of OMS Oracle home needs to be the same as one used before.
3. Run `omsca` in recovery mode specifying the export file taken earlier to configure the OMS:

```
omsca recover -as -ms -backup_file <file>
```

4. Configure agent:

```
agentca -f followed by
emctl secure agent -emdWalletSrcUrl
"http://slb:<httpport>/em"
```

5. Add the new OMS to the SLB
6. Relocate the OMS and Repository target to reinstalled Agent:

```
emctl config emrep -agent <agent on Host B>
```
7. Locate duplicate targets from the Management Services and Repository Overview page. Relocate duplicate targets from the Agent on Host "A" to the Agent on Host "B". Delete the Agent on Host "A".
8. Re-enroll the additional OMS, if any, with the recovered Administration Server

```
emctl enroll oms -as_host <HostB> -as_port <admin secure port>
```

Run this command on each additional OMS.
9. Verify that the site is fully operational.

17.3.2.3.6 Multiple OMS, SLB configured, additional OMS recovered on same or different host

Multi OMS site. OMSs fronted by SLB. OMS configuration backed up using `emctl exportconfig oms` command on the first OMS. Additional OMS is lost and needs to be recovered on same or different host.

1. Ensure that shared loader receive directory and shared software library locations are accessible.
2. If recovering on same host, restore the Software Homes from a filesystem backup. Alternatively, if a backup does not exist, or when recovering on a different host, use the software-only install method to reconstruct the WebLogic and OMS Oracle Home. Add-ons that were installed earlier need to be reinstalled in software-only mode and all patches that were applied earlier need to be reapplied. The location of the restored OMS Oracle home needs to be the same as the previous.
3. Run `omsca` in recovery mode specifying the export file taken earlier to configure the OMS:

```
omsca recover -ms -backup_file <file>
```
4. Configure agent:

```
agentca -f followed by emctl secure agent -emdWalletSrcUrl
"http://slb:<httpport>/em"
```
5. Add the new OMS (if recovered on a different host) to the SLB
6. At this point, three possibilities exist depending upon the port used by the reinstalled agent that comes along with the OMS:
 - Option A: OMS installed on same host and agent gets the same port as earlier. OMS automatically blocks the agent. Resync the agent from agent homepage.
 - Option B: OMS installed on same host and agent gets a different port. Locate duplicate targets from the *Management Services and Repository Overview* page. Relocate duplicate targets from old agent to reinstalled agent. Delete the old agent.
 - Option C: OMS installed on different host

Locate duplicate targets from the *Management Services and Repository Overview* page. Relocate duplicate targets from old agent to reinstalled agent.

7. Verify that the site is fully operational.

17.3.3 Agent Backup and Recovery

The 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.

17.3.3.1 Backing Up Agents

There are no special considerations for backing up Agents. As a best practice, reference 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 Grid Control console to install and maintain reference Agent installs.

17.3.3.2 Recovering Agents

If an Agent is lost, it should be reinstalled by cloning from a reference install. Cloning from a reference install is often the fastest way to recover an Agent install as it is not necessary to track and reapply customizations and patches. Care should be taken to reinstall the Agent using the same port. Using the Enterprise Manager's Agent Resynchronization feature, a reinstalled Agent can be reconfigured using target information present in the repository. When the 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 Agent from overwriting previous customizations.

Blocked Agents: A Blocked Agent is a condition where the OMS rejects all heartbeat or upload requests from the blocked Agent. Hence, a blocked Agent will not be able to upload any alerts or metric data to the OMS. However, blocked Agents continue to collect monitoring data.

The Agent can be resynchronized and unblocked from the Agent homepage by clicking on the **Resynchronize Agent** button. Resynchronization pushes all targets from the repository to the Agent and then unblocks the Agent.

17.3.3.3 Agent Recovery Scenarios

The following scenarios illustrate various Agent recovery situations along with the recovery steps. Agent recovery is supported for Agent versions 10.2.0.5 and later. The Agent resynchronization feature requires that a reinstalled Agent use the same port as the previous Agent that crashed.

17.3.3.3.1 Agent Reinstall Using the Same Port An Agent is monitoring multiple targets. The Agent installation is lost.

1. Deinstall Agent OracleHome using the Oracle Universal Installer.
2. Install a new Agent or use the Agent clone option to reinstall the Agent through Enterprise Manager. Specify the same port as used by the crashed Agent. The location of install need not be same as previous install.

The OMS detects that Agent has been re-installed and blocks the Agent.

3. Initiate Agent Resynchronization from the Agent homepage.

All targets in the repository are pushed to the new Agent. The Agent is instructed to clear backlogged files and then do a clearstate. Agent is unblocked.

4. Reconfigure User-defined Metrics if the location of User-defined Metric scripts have changed.
5. Verify that the Agent is operational and all target configurations have been restored.

17.3.3.3.2 Agent Restore from Filesystem Backup An Agent is monitoring multiple targets. File system backup for the Agent OracleHome exists. The Agent install is lost.

1. Deinstall Agent OracleHome using OUI.

2. Restore the Agent from file system backup. Start the Agent.

OMS detects that Agent has been restored from backup and blocks the Agent.

3. Initiate Agent Resynchronization from the Agent homepage.

All targets in the repository are pushed to the new Agent. The Agent is instructed to clear backlogged files and performs a clearstate. The Agent is unblocked.

4. Verify that the 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.

17.3.4 Recovering from a Simultaneous OMS-Repository Failure

When both OMS and repository fail simultaneously, the recovery situation becomes more complex depending upon factors such as whether the OMS and repository recovery has to be performed on the same or different host, or whether there are multiple OMSs fronted by an SLB. In general, the order of recovery for this type of compound failure should be repository first, followed by OMS(s) following the steps outlined in the appropriate recovery scenarios discussed earlier. The following scenarios illustrate two OMS-Repository failures and the requisite recovery steps.

17.3.4.1 Collapsed Configuration: Incomplete Repository Recovery, Primary OMS on the Same Host

Repository and the primary OMS are installed on same host (host "A"). The repository database is running in *noarchivelog* mode. Full cold backup is available. A recent OMS backup file exists (`emctl exportconfig oms`). The repository, OMS and the Agent crash.

1. Follow the repository recovery procedure shown in [Section 17.3.1.3.2, "Incomplete Recovery on the Same Host"](#) with the following exception:

Since the OMS OracleHome is not available and repository resynchronization has to be initiated before starting an OMS against the restored repository, submit "resync" via the following PL/SQL block. Log into the 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 17.3.2.3.1, "Single OMS, No Server Load Balancer \(SLB\), OMS Restored on the same Host"](#)
3. Verify that the site is fully operational.

17.3.4.2 Distributed Configuration: Incomplete Repository Recovery, Primary OMS and additional OMS on Different Hosts, SLB Configured

The Repository, primary OMS, and additional OMS all reside on the different hosts. Repository database was running in *noarchive* 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 repository recovery procedure shown in [Section 17.3.1.3.2, "Incomplete Recovery on the Same Host"](#) with the following exception:

Since OMS OracleHome is not yet available and Repository resync has to be initiated before starting an OMS against the restored repository, submit resync via the following PL/SQL block. Log into the 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 17.3.2.3.5, "Multiple OMS, Server Load Balancer configured, Primary OMS Recovered on a Different Host"](#) with the following exception:

Override the 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 17.3.2.3.5](#).

3. Follow the OMS recovery procedure shown in [Section 17.3.2.3.6, "Multiple OMS, SLB configured, additional OMS recovered on same or different host"](#) with the following exception:

Override the repository connect description and AdminServer details present in the backup file by passing the additional omsca parameters:

```
"-REPOS_CONN_STR <restored repos descriptor>" -AS_HOST <recovered admin host>
-AS_HTTPS_PORT <recovered admin port>
```

This must be added along with other parameters listed in [Section 17.3.2.3.6, "Multiple OMS, SLB configured, additional OMS recovered on same or different host"](#).

4. Verify that the site is fully operational.

17.3.5 EMCTL High Availability Commands

The Enterprise Manager command-line utility (emctl) adds many new commands that allow you to perform requisite backup and recovery operations for all major components.

exportconfig oms

Exports a snapshot of the OMS configuration to the specified directory.

Usage:

```
emctl exportconfig oms [-sysman_pwd <sysman password>]
[-dir <backup dir>] Specify the directory used to store the backup file
```

`[-keep_host]` Specify to back up hostname if no SLB is defined
(Use this option only if recovery will be performed
on the machine that responds to this hostname)

importconfig oms

Imports the OMS configuration from the specified backup file.

Usage:

```
emctl importconfig oms [-sysman_pwd <sysman password>] [-reg_pwd <registration
password>]
    -file <backup file>      Required backup file to import from
    [-key_only]              Specify to import emkey only
    [-no_resecure]          Specify not to resecure the oms after import
                           (default is to resecure after import)
```

config emrep

Configures the OMS and repository target. The command is used to change the monitoring Agent for the target and/or the connection string used to monitor this target.

Usage:

```
emctl config emrep [-sysman_pwd <sysman password>]
    [-agent <new agent>]     Specify a new destination Agent for emrep target
    [-conn_desc [<jdbc connect descriptor>]]
                           Update the Connect Descriptor with value if specified,
                           else from value stored in the emoms.properties file.
```

config repos

Configures the repository database target. The command is used to change the monitoring Agent for the target and/or the monitoring properties (hostname, Oracle Home and connection string used to monitor this target).

Usage:

```
emctl config repos [-sysman_pwd <sysman password>]
    [-agent <new agent>]     Specify new destination agent for repository target
    [-host <new host>]       Specify new hostname for repository target
    [-oh <new oracle home>] Specify new OracleHome for repository target
    [-conn_desc [<jdbc connect descriptor>]]
                           Update the Connect Descriptor with the specified value,
                           else from the value stored in emoms.properties
```

resync repos

Submits a repository resynchronization operation. When the `-full` option is specified, all agents are instructed to upload the latest state to the repository. A list of agents can be specified using the `-agentlist` option to resync with a given list of agents.

Usage:

```
emctl resync repos (-full|-agentlist "agent names") [-name "resync name"]
[-sysman_pwd "sysman password"]
```

abortresync repos

Aborts the currently running repository resynchronization operation. Use the `-full` option to stop a full repository resynchronization. Use the `-agentlist` option to stop resync on a list of agents.

Usage:

```
emctl abortresync repos (-full|-agentlist "agent names") -name "resync name"
[-sysman_pwd "sysman password"]
```

statusresync repos

Lists the status of given repository resynchronization operation.

Usage:

```
emctl statusresync repos -name "resync name"
```

create service

Valid on Windows only. The command creates a service for the Oracle Management Services on Windows. You use this command to manage the Windows service for the OMS on a failover host in a Cold Failover Cluster setup.

Usage:

```
emctl create service [-user <username>] [-pwd <password>]
                    -name <servicename>      Name of service to be created
```

delete service

Valid on Windows only. Deletes the service for the Oracle Management Services on Windows.

Usage:

```
emctl delete service
                    -name <servicename>      Name of service to be deleted
```

resyncAgent

Resynchronizes a restored or reinstalled Agent by pushing all target configuration from the repository.

Usage:

```
emcli resyncAgent -agent="Agent Name"
                  [-keep_blocked]
```

High Availability: Multiple Resource Configurations

Multiple resource configurations add redundancy to your Enterprise Manager installation, thus creating a higher level of availability. Using multiple hosts, Enterprise Manager components can be replicated and configured for failover. This not only increases survivability but also greatly reduces the time required to restore Enterprise Manager after a failure has occurred. If you have a high RTO, installing Enterprise Manager using a multi-host configuration is essential.

This chapter covers the following topics:

- [Managing Multiple Hosts and Deploying a Remote Management Repository](#)
- [Using Multiple Management Service Installations](#)
- [High Availability Configurations](#)
- [Installation Best Practices for Enterprise Manager High Availability](#)
- [Configuration With Grid Control](#)
- [Configuring Oracle Enterprise Manager for Active and Passive Environments](#)
- [Using Virtual Host Names for Active and Passive High Availability Environments in Enterprise Manager Database Control](#)
- [Configuring Grid Control Repository in Active/Passive High Availability Environments](#)
- [How to Configure Grid Control OMS in Active/Passive Environment for High Availability Failover Using Virtual Host Names](#)
- [Configuring Targets for Failover in Active/Passive Environments](#)

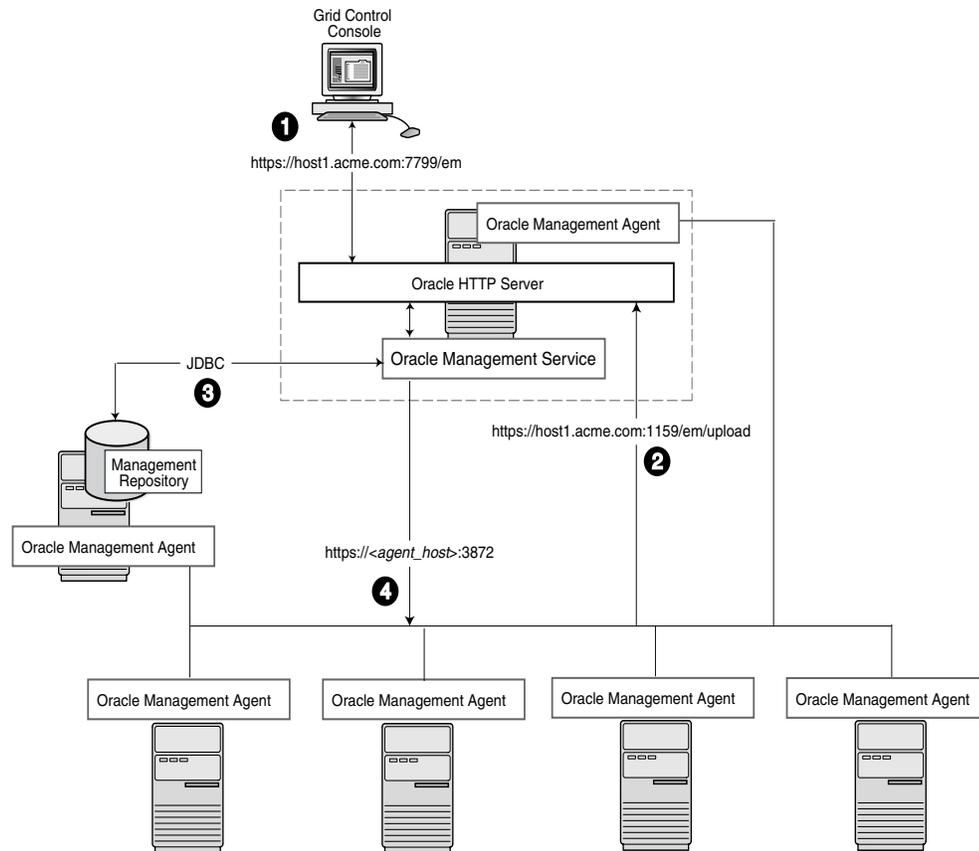
18.1 Managing Multiple Hosts and Deploying a Remote Management Repository

Installing all the Grid Control components on a single host is an effective way to initially explore the capabilities and features available to you when you centrally manage your Oracle environment.

A logical progression from the single-host environment is to a more distributed approach, where the Management Repository database is on a separate host and does not compete for resources with the Management Service. The benefit in such a configuration is scalability; the workload for the Management Repository and Management Service can now be split. This configuration also provides the flexibility to adjust the resources allocated to each tier, depending on the system load. (Such a

configuration is shown in [Figure 18–2, "Grid Control Architecture with Multiple Management Service Installations".](#))

Figure 18–1 Grid Control Components Distributed on Multiple Hosts with One Management Service



In this more distributed configuration, data about your managed targets travels along the following paths so it can be gathered, stored, and made available to administrators by way of the Grid Control console:

1. Administrators use the Grid Control console to monitor and administer the targets just as they do in the single-host scenario described in ["Deploying Grid Control Components on a Single Host"](#) on page 17-2.
2. Management Agents are installed on each host on the network, including the Management Repository host and the Management Service host. The Management Agents upload their data to the Management Service by way of the Management Service upload URL, which is defined in the `emd.properties` file in each Management Agent home directory. The upload URL uploads the data directly through the Oracle HTTP Server.
3. The Management Repository is installed on a separate machine that is dedicated to hosting the Management Repository database. The Management Service uses JDBC connections to load data into the Management Repository database and to retrieve information from the Management Repository so it can be displayed in the Grid Control console. This remote connection is defined in the Administration Server and can be accessed and changed through `emctl` commands.
4. The Management Service communicates directly with each remote Management Agent over HTTP by way of the Management Agent URL. The Management

Agent URL is defined by the `EMD_URL` property in the `emd.properties` file of each Management Agent home directory. As described in [Section 17.2, "Deploying Grid Control Components on a Single Host"](#), the Management Agent includes a built-in HTTP listener so no Oracle HTTP Server is required on the Management Agent host.

18.2 Using Multiple Management Service Installations

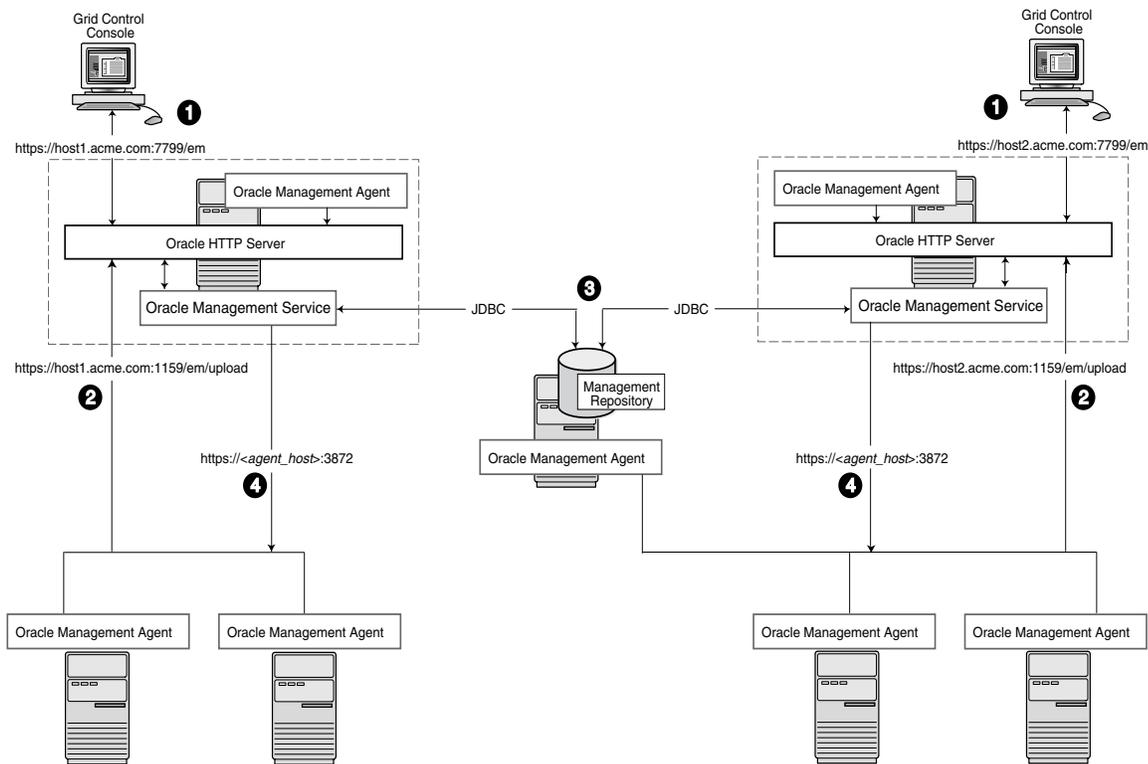
In larger production environments, you may find it necessary to add additional Management Service installations to help reduce the load on the Management Service and improve the efficiency of the data flow.

Note: When you add additional Management Service installations to your Grid Control configuration, be sure to adjust the parameters of your Management Repository database. For example, you will likely need to increase the number of processes allowed to connect to the database at one time. Although the number of required processes will vary depending on the overall environment and the specific load on the database, as a general guideline, you should increase the number of processes by 40 for each additional Management Service.

For more information, see the description of the `PROCESSES` initialization parameter in the *Oracle Database Reference*.

Understanding the Flow of Management Data When Using Multiple Management Services

[Figure 18–2, "Grid Control Architecture with Multiple Management Service Installations"](#) shows a typical environment where an additional Management Service has been added to improve the scalability of the Grid Control environment.

Figure 18–2 Grid Control Architecture with Multiple Management Service Installations

In a multiple Management Service configuration, the management data moves along the following paths:

1. Administrators can use one of two URLs to access the Grid Control console. Each URL refers to a different Management Service installation, but displays the same set of targets, all of which are loaded in the common Management Repository. Depending upon the host name and port in the URL, the Grid Control console obtains data from the Management Service (by way of the Oracle HTTP Server) on one of the Management Service hosts.
2. Each Management Agent uploads its data to a specific Management Service, based on the upload URL in its `emd.properties` file. That data is uploaded directly to the Management Service by way of Oracle HTTP Server.

Whenever more than one Management Service is installed, it is a best practice to have the Management Service upload to a shared directory. This allows all Management Service processes to manage files that have been uploaded from any Management Agent. This protects from the loss of any one Management Server causing a disruption in upload data from Management Agents.

Configure this functionality from the command line of each Management Service process as follows:

```
emctl config oms loader -shared <yes|no> -dir <load
directory>
```

Important: By adding a load balancer, you can avoid the following problems:

- Should the Management Service fail, any Management Agent connected to it cannot upload data.
- Because user accounts only know about one Management Service, users lose connectivity should the Management Service go down even if the other Management Service is up.

See [Section 18.3, "High Availability Configurations"](#) for information regarding load balancers.

Note: If the software library is being used in this environment, it should be configured to use shared storage in the same way as the shared Management Service loader. To modify the location for the software library:

1. Click the **Deployments** tab on the Enterprise Manager Home page.
 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.
-
-

3. Each Management Service communicates by way of JDBC with a common Management Repository, which is installed in a database on a dedicated Management Repository host. Each Management Service uses the same database connection information, defined in the `emgc.properties` file, to load data from its Management Agents into the Management Repository. The Management Service uses the same connection information to pull data from the Management Repository as it is requested by the Grid Control console.
4. Any Management Service in the system can communicate directly with any of the remote Management Agents defined in the common Management Repository. The Management Services communicate with the Management Agents over HTTP by way of the unique Management Agent URL assigned to each Management Agent.

As described in [Section 17.2, "Deploying Grid Control Components on a Single Host"](#), the Management Agent URL is defined by the `EMD_URL` property in the `emd.properties` file of each Management Agent home directory. Each Management Agent includes a built-in HTTP listener so no Oracle HTTP Server is required on the Management Agent host.

18.3 High Availability Configurations

You can configure Enterprise Manager to run in either active-active or active-passive mode using a single instance database as the Management Repository. The following text summarizes the active-active mode.

Refer to the following sections for more information about common Grid Control configurations that take advantage of high availability hardware and software solutions. These configurations are part of the Maximum Availability Architecture (MAA).

- [Configuring the Management Repository](#)
- [Configuring the Management Services](#)
- [Configuring Additional Management Services](#)
- [Configuring the Management Agent](#)
- [Disaster Recovery](#)

18.3.1 Configuring the Management Repository

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.

- **Configure Database**
 - For both high availability and scalability, you should configure the Management Repository in the latest certified database version, with the RAC option enabled. Check for the latest version of database certified for Enterprise Manager from the Certify tab on the My Oracle Support website.
 - Choose Automatic Storage Management (ASM) as the underlying storage technology.
 - When the database installation is complete:
 - Go to `$ORACLE_HOME/rbdms/admin` directory of the database home and execute the `'dbmspool.sql'`
 - This installs the `DBMS_SHARED_POOL` package, which will help in improving throughput of the Management Repository.
- **Install Enterprise Manager**

While installing Enterprise Manager using Oracle Universal Installer (OUI), you will be presented with the option for configuring the Management Repository using an existing database.

18.3.1.1 Post Management Service - Install Management Repository Configuration

There are some parameters that should be configured during the Management Repository database install (as previously mentioned) and some parameters that should be set after the Management Service has been installed.

Start by installing Management Agents on each Management Repository RAC node. Once the Management Agents are installed and the Management Repository database is discovered as a target, the Enterprise Manager console can be used to configure these best practices in the Management Repository.

These best practices fall in the area of:

- **Configuring Storage**
- **Configuring Oracle Database 11g with RAC for High Availability and Fast Recoverability**
 - 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.

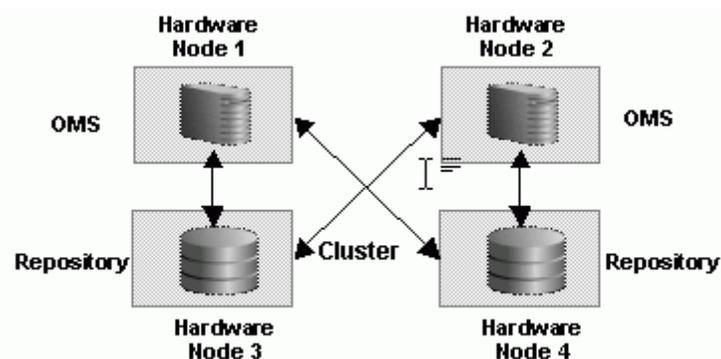
18.3.2 Configuring the Management Services

Once you configure the Management Repository, the next step is to install and configure the Enterprise Manager Grid Control mid-tier, the Management Services, for greater availability. Before discussing steps that add mid-tier redundancy and scalability, note that the Management Service itself has a built-in restart mechanism based on the Oracle Weblogic Node Manager and the Oracle Process Management and Notification Service (OPMN). These services will attempt to restart a Management Service that is down. It is advisable to run OPMN and Node Manager as operating system services, so that they restart automatically if their host machine is restarted.

18.3.2.1 Management Service Install Location

If you are managing a large environment with multiple Management Services and Management Repository nodes, then consider installing the Management Services on hardware nodes that are different from Management Repository nodes (Figure 18-3). This allows you to scale out the Management Services in the future.

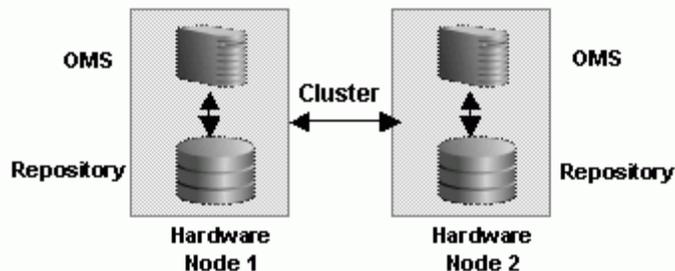
Figure 18-3 Management Service and Management Repository on Separate Hardware



Also consider the network latency between the Management Service and the Management Repository while determining the Management Service install location. The distance between the Management Service and the Management Repository may be one of the factors that affect network latency and hence determine Enterprise Manager performance.

If the network latency between the Management Service and Management Repository tiers is high or the hardware available for running Enterprise Manager is limited, then the Management Service can be installed on the same hardware as the Management Repository (Figure 18-4). This allows for Enterprise Manager high availability, as well as keep the costs down.

Figure 18-4 Management Service and Management Repository on Same Hardware



18.3.2.2 Configuring the First Management Service for High Availability

If you plan ahead, you can configure your Enterprise Manager deployment for high availability by choosing the correct options during the first Management Service install. You can also retrofit the MAA best practices into an existing Enterprise Manager deployment configured initially using the default install options. Both paths will be addressed in the following sections.

18.3.2.3 Configuring Management Service to Management Repository Communication

Management Service processes need to be configured to communicate with each node of the RAC Management Repository in a redundant fashion.

Note that Real Application Cluster (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 *Oracle Database Net Services Administrator's Guide* for details.

Configure the repository connect descriptor by running the `emctl` command from any Management Service:

```
emctl config oms -store_repos_details -repos_conndesc '(DESCRIPTION=
(AADDRESS_LIST=(FAILOVER=ON)
(AADDRESS=(PROTOCOL=TCP) (HOST=node1-vip.example.com) (PORT=1521))
(AADDRESS=(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 to make the same change to the monitoring configuration used for the Management Services and Repository target: `emctl config emrep -conn_desc`

18.3.2.4 Configuring Shared File System Loader

The Management Service for Grid Control has a high availability feature called the Shared Filesystem Loader. In the Shared Filesystem Loader, management data files received from Management Agents are stored temporarily on a common shared location called the shared receive directory. All Management Services are configured to

use the same storage location for the shared receive directory. The Management Services coordinate internally and distribute among themselves the workload of uploading files into the Management Repository. Should a Management Service go down, its workload is taken up by surviving Management Services. You must choose a shared receive directory that is accessible by all the Management Services using redundant file storage.

During the first Management Service installation, the shared receive directory can be configured out of the box by passing SHARED_RECEIVE_DIRECTORY_LOCATION=<shared recv directory> option to runInstaller (setup.exe on Windows). Oracle recommends that this location be outside the Instance Home and Middleware Home locations.

If not configured during installation, the Shared Filesystem Loader can also be configured post-installation by running the following emctl command on every Management Service:

```
emctl config oms loader -shared yes -dir <shared recv directory>
```

Note: If shared filesystem Loader is configured on the first Management Service, any additional management service that is installed later will inherit the shared filesystem loader configuration. Therefore, ensure that the shared recv directory is available on the additional Management Service prior to running install.

Consider the following while configuring Shared Filesystem Loader on Windows.

- On Windows platforms, the Enterprise Manager install may configure the Management Service to run as a service using 'LocalSystem' account. This is a local account and will typically not have access to the network drive for the shared filesystem that requires domain authentication. To resolve this issue, configure the Management Service to run as a domain user as follows:
 1. Go to the Control Panel and then open the Services panel.
 2. Double click the appropriate service (Oracleoms11gProcessManager).
 3. Change the 'Log on as' user from the 'Local System Account' to **This Account**.
 4. Specify the domain user that has access to shared network drive.
 5. Click **OK**.
- Do not use local drive letters for mapped network shares while configuring the shared filesystem on Windows. Use UNC locations instead.

```
emctl config oms loader -shared yes -dir \\<host>\<share-name>\<recv-dir>
for example
emctl config oms loader -shared yes -dir \\hostA\vol1\recv
```

Note the use of double backslashes while specifying the directory location.

Note: User equivalence should be set up properly across OMS so that files created by one OMS on the loader directory are modifiable by other OMS.

18.3.2.5 Configuring Software Library

Since software library location has to be accessed by all Management Services, considerations similar to shared filesystem loader directory apply here too. The configuration of software library is not covered during install. It needs to be configured post-install using the Enterprise Manager Console:

1. On the Enterprise Manager home page, click the **Deployments** tab.
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.

18.3.2.6 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. This is a two step process:

1. Configure the SLB
2. Make needed changes on the Management Services

18.3.2.6.1 SLB Setup Use the following table as reference for setting up the SLB with Grid Control Management Services.

Table 18–1 Management Service Ports

Grid Control Service	TCP Port	Monitor Name	Persistence	Pool Name	Load Balancing	Virtual Server Name	Virtual Server Port
Secure Upload	1159	mon_gcsu1159	None	pool_gcsu1159	Round Robin	vs_gcsu1159	1159
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 18–1](#).

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 18–2 Pools

Pool Name	Usage	Members	Persistence	Load Balancing
pool_gcsu1159	Secure upload	HostA:1159 HostB:1159	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 18–3 Virtual Servers

Virtual Server Name	Usage	Virtual Server Port	Pool
vs_gcsu1159	Secure upload	1159	pool_gcsu1159
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 18–4 Monitors

Monitor Name	Configuration	Associate With
mon_gcsu1159	Type: https Interval: 60 Timeout: 181 Send String: GET /em/upload Receive String: Http Receiver Servlet active!	HostA:1159 HostB:1159
mon_gcar4889	Type: http Interval: 60 Timeout: 181 Send String: GET /em/genwallet Receive String: GenWallet Servlet activated	HostA:4889 HostB:4889

Table 18–4 (Cont.) Monitors

Monitor Name	Configuration	Associate With
mon_gcsc7799	Type: https Interval: 5 Timeout: 16 Send String: GET /em/console/home HTTP/1.0\n Receive String: /em/console/logon/logon;jsessionid=	HostA:7799 HostB:7788
mon_gcuc7788 (optional)	Type: https Interval: 5 Timeout: 16 Send String: GET /em/console/home HTTP/1.0\n Receive String: /em/console/logon/logon;jsessionid=	HostA:7788 HostB:7788

Note: If you have SSO configured, use the following alternate definitions for the mon_gcsc7799 and mon_gcuc7788 monitors.

Table 18–5 Monitors for SSO Configuration

Monitor Name	Configuration	Associate With
mon_gcsc7799	Type: https Interval: 5 Timeout: 16 Send String: GET /em/genwallet Receive String: GenWallet Servlet activated	HostA:7799 HostB:7788
mon_gcuc7788 (optional)	Type: https Interval: 5 Timeout: 16 Send String: GET /em/genwallet Receive String: GenWallet Servlet activated	HostA:7788 HostB:7788

Note: F5 SLB monitors expect the "Receive String" within the first 5120 characters of a response. For SSO environments, the "Receive String" may be returned at some point beyond the 5120 limit. The monitor will not function in this situation.

18.3.2.6.2 Enterprise Manager Side Setup

Perform the following steps:

1. Resecure the 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 Management Service through a load balancer. You must run the following command to regenerate the certificate on the first Management Service:

```
emctl secure
  -oms -sysman_pwd <sysman_pwd>
  -reg_pwd <agent_reg_password>
  -host slb.example.com
  -secure_port 1159
  -slb_port 1159
  -slb_console_port 443
  [-lock] [-lock_console]
```

2. Resecure all Management Agents

Management Agents that were installed prior to SLB setup, including the Management Agent that comes with the Management Service install, would be uploading directly to the Management Service. These Management Agents will not be able to upload after SLB is setup. Resecure these Management Agents to upload to the SLB by running the following command on each Management Agent:

```
emctl secure agent -emdWalletSrcUrl https://slb.example.com:<upload port>/em
```

18.3.3 Configuring Additional Management Services

Once your first Management Service is setup for high availability, there are two paths to setting up your additional Management Services for high availability:

- Installing a fresh additional Management Service as per MAA best practices
- Retrofitting MAA best practices on existing Additional Management Service

In either of the two cases, the following considerations should be noted:

- The additional Management Service should be hosted in close network proximity to the Management Repository database for network latency reasons.
- Configure the same path on all Management Services for the directory used for the shared filesystem loader.
- Additional Management Services should be installed using the same OS user and group as the first Management Service. Proper user equivalence should be setup so that files created by each Management Service on the shared loader directory can be accessed and modified by the other Management Service processes.
- Adjust the parameters of your Management Repository database. For example, you will likely need to increase the number of processes allowed to connect to the database at one time. Although the number of required processes will vary depending on the overall environment and the specific load on the database, as a general guideline, you should increase the number of processes by 40 for each additional Management Service.

18.3.3.1 Installing a Fresh Additional Management Service According MAA Best Practices

Install the additional Management Service using the OUI installer. The additional Management Service will inherit most of the HA configuration from the first Management Service. Post installation, do the following step to complete the HA configuration:

- Update the SLB configuration by adding the additional Management Service to the different pools on the SLB. Setup monitors for the new Management Service.

18.3.3.2 Retrofitting MAA Best Practices on Existing Additional Management Service

Once you have the additional Management Service installed, use the following steps to copy over the configuration from the first Management Service.

1. Export the configuration from the first Management Service using `emctl exportconfig oms -dir <location for the export file>`
2. Copy over the exported file to the additional Management Service
3. Shutdown the additional Management Service

4. Import the exported configuration on the additional Management Service using `emctl importconfig oms -file <full path of the export file>`
5. Restart the additional Management Service
6. Setup EMCLI using `emcli setup -url=https://slb.example.com/em -username sysman -password <sysman password> -nocertvalidate`
7. Resecure the Management Agent that is installed with the additional Management Service to upload to SLB using `emctl secure agent -emdWalletSrcUrl https://slb.example.com:<upload port>/em`
8. Update the SLB configuration by adding the additional Management Service to the different pools on the SLB. Setup monitors for the new Management Service. Modify the `ssl.conf` file to set the Port directive to the SLB virtual port used for UI access.

18.3.4 Configuring the Management Agent

The final piece of Enterprise Manager High Availability is the Management Agent configuration. It is worthwhile to note that the Management Agent has high availability built into it out of the box. A 'watchdog' process, created automatically on Management Agent startup, monitors each Management Agent process. In the event of a failure of the Management Agent process, the 'watchdog' will try to automatically re-start the Management Agent process.

Communication between the Management Agent and the Management Service tiers in a default Enterprise Manager Grid Control install is a point-to-point set up. Therefore, the default configuration does not protect from the scenario where the Management Service becomes unavailable. In that scenario, a Management Agent will not be able to upload monitoring information to the Management Service (and to the Management Repository), resulting in the targets becoming unmonitored until that Management Agent is manually configured to point to a second Management Service.

To avoid this situation, use hardware Server Load Balancer (SLB) between the Management Agents and the Management Services. The Load Balancer monitors the health and status of each Management Service and makes sure that the connections made through it are directed to surviving Management Service nodes in the event of any type of failure. As an additional benefit of using SLB, the load balancer can also be configured to manage user communications to Enterprise Manager. The Load Balancer handles this through the creation of 'pools' of available resources.

- Configure the Management Agent to Communicate Through SLB

The load balancer provides a virtual IP address that all Management Agents can use. Once the load balancer is setup, the Management Agents need to be configured to route their traffic to the Management Service through the SLB. This can be achieved through a couple of property file changes on the Management Agents.

Resecure all Management Agents - Management Agents that were installed prior to SLB setup 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
```

- Configure the Management Agent to Allow Retrofitting a SLB

Some installations may not have access to a SLB during their initial install, but may foresee the need to add one later. If that is the case, consider configuring the Virtual IP address that will be used for the SLB as a part of the initial installation and having that IP address point to an existing Management Service. Secure communications between Management Agents and Management Services are based on the host name. If a new host name is introduced later, each Management Agent will not have to be re-secured as it is configured to point to the new Virtual IP maintained by the SLB.

18.3.4.1 Load Balancing Connections Between the Management Agent and the Management Service

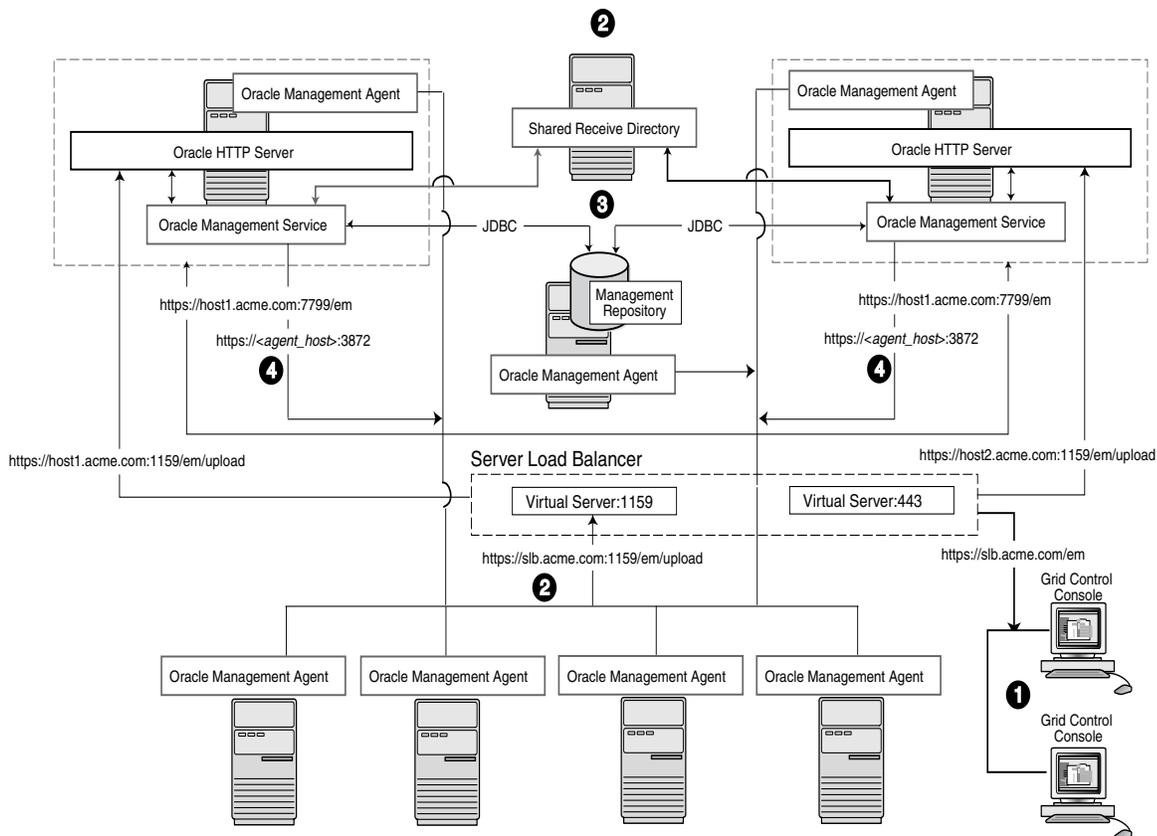
Before you implement a plan to protect the flow of management data from the Management Agents to the Management Service, you should be aware of some key concepts.

Specifically, Management Agents do not maintain a persistent connection to the Management Service. When a Management Agent needs to upload collected monitoring data or an urgent target state change, the Management Agent establishes a connection to the Management Service. If the connection is not possible, such as in the case of a network failure or a host failure, the Management Agent retains the data and reattempts to send the information later.

To protect against the situation where a Management Service is unavailable, you can use a load balancer between the Management Agents and the Management Services.

However, if you decide to implement such a configuration, be sure to understand the flow of data when load balancing the upload of management data.

[Figure 18–5](#) shows a typical scenario where a set of Management Agents upload their data to a load balancer, which redirects the data to one of two Management Service installations.

Figure 18–5 Load Balancing Between the Management Agent and the Management Service

In this example, only the upload of Management Agent data is routed through the load balancer. The Grid Control console still connects directly to a single Management Service by way of the unique Management Service upload URL. This abstraction allows Grid Control to present a consistent URL to both Management Agents and Grid Control consoles, regardless of the loss of any one Management Service component.

When you load balance the upload of Management Agent data to multiple Management Service installations, the data is directed along the following paths:

1. Each Management Agent uploads its data to a common load balancer URL. This URL is defined in the `emd.properties` file for each Management Agent. In other words, the Management Agents connect to a virtual service exposed by the load balancer. The load balancer routes the request to any one of a number of available servers that provide the requested service.
2. Each Management Service, upon receipt of data, stores it temporarily in a local file and acknowledges receipt to the Management Agent. The Management Services then coordinate among themselves and one of them loads the data in a background thread in the correct chronological order.

Also, each Management Service communicates by way of JDBC with a common Management Repository, just as they do in the multiple Management Service configuration defined in [Section 18.2, "Using Multiple Management Service Installations"](#).

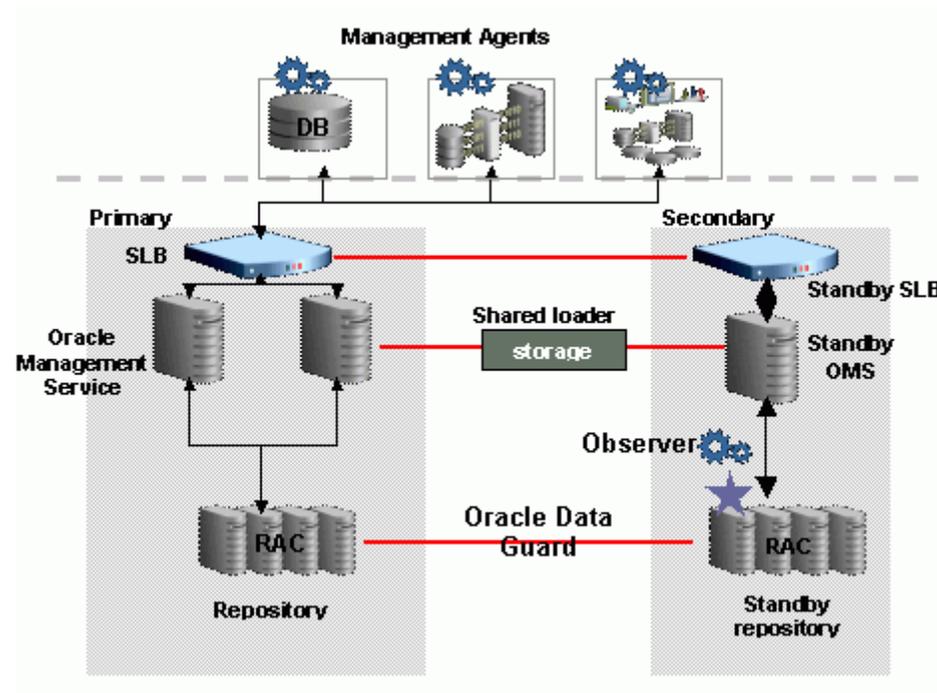
3. Each Management Service communicates directly with each Management Agent by way of HTTP, just as they do in the multiple Management Service configuration defined in [Section 18.2, "Using Multiple Management Service Installations"](#).

18.3.5 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 mandates deploying a remote failover architecture that allows a secondary datacenter to take over the management infrastructure in the event that disaster strikes the primary management infrastructure.

Figure 18–6 Disaster Recovery Architecture



As can be seen in [Figure 18–6](#), setting up disaster recovery for Enterprise Manager essentially consists of installing a standby RAC, a standby Management Service and a standby Server Load Balancer and configuring them to automatically startup when the primary components fail.

The following sections lists the best practices to configure the key Enterprise Manager components for disaster recovery:

- [Prerequisites](#)
- [Setup Standby Database](#)
- [Setup Standby Management Service](#)
- [Switchover](#)
- [Failover](#)
- [Automatic Failover](#)

18.3.5.1 Prerequisites

The following prerequisites must be satisfied.

- The primary site must be configured as per Grid 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.
- There must be sufficient network bandwidth between the primary and standby sites to handle peak redo data generation.
- Configure shared storage used for shared filesystem loader and software library to be replicated at the primary and standby site. In the event of a site outage, the contents of this shared storage must be made available on the standby site using hardware vendor disk level replication technologies.
- 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.

18.3.5.2 Setup Standby Database

The starting point of this step is to have the primary site configured as per Grid Control MAA guidelines. The following steps lay down the procedure for setting up the 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 CRS and Database software on the standby Management Repository hosts. The version used must be the same as that on the primary site.

2. Prepare 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 Physical Standby Database

In Enterprise Manager, the standby Management Repository database must be physical standbys. Use the Enterprise Manager Console to setup a physical standby database in the standby environment. Note that 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 Enterprise Manager Console to convert the single instance standby database to RAC. Also, note that during single instance standby creation, the database files should be created on shared storage to facilitate conversion to RAC later.

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.

4. Add Static Service to 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

6. Verify 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.

18.3.5.3 Setup Standby Management Service

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 Management Services 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.
- The shared storage used for the shared filesystem loader and software library must be made available on the standby site using the same paths as the primary site.

18.3.5.3.1 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. Perform a software-only install by running the installer with the following arguments:

```
runInstaller -noconfig -validationaswarnings
```

3. Apply one-off patches

```
<OMS Oracle Home>/install/oneoffs/apply_NewOneoffs.pl <OMS Oracle Home> OC9321514,9207217
```

4. 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 DOMAIN_NAME GCDomainStby -nostart
```

When prompted for Management Repository details, provide the Primary database details.

5. Configure the virtualization add on by running the following command:

```
addonca -oui -omsonly -name vt -install gc
```
6. Configure the Management Agent that comes with the Management Service install by running:

```
<Agent Oracle Home>/bin/agentca -f
```
7. Export the configuration from the primary Management Service using:

```
emctl exportconfig oms -dir <location for the export file>
```
8. Copy over the exported file to the standby Management Service
9. Import the exported configuration on the standby Management Service using:

```
emctl importconfig oms -file <full path of the export file>
```

18.3.5.3.2 Installing Additional Standby Management Services Install additional standby Management Services as follows:

Specify the primary database details and the standby administration server details on the installer screens. Post installation, do the following steps to complete the HA configuration.

1. Do a software only install by running the installer with following arguments:

```
runInstaller -noconfig -validationaswarnings
```
2. Apply one-off patches

```
<OMS Oracle Home>/install/oneoffs/apply_NewOneoffs.pl <OMS Oracle Home> OC9321514,9207217
```
3. Configure the Management Service by running `omsca`. When prompted for Management Repository details, provide the Primary database details. When prompted for Administration Server details, provide the standby administration server details.

```
omsca add -nostart
```
4. Configure the virtualization add on by running the following command

```
addonca -oui -omsonly -install gc
```
5. Configure the Management Agent that comes with the Management Service install by running:

```
<Agent Oracle Home>/bin/agentca -f
```
6. Export the configuration from the primary Management Service using:

```
emctl exportconfig oms -dir <location for the export file>
```
7. Copy over the exported file to the standby Management Service
8. Import the exported configuration on the standby Management Service using:

```
emctl importconfig oms -file <full path of the export file>
```

18.3.5.3.3 Validating Your Installation and Complete the Setup Validate your installation and complete the setup as follows:

1. Update the standby SLB configuration by adding the standby Management Services to the different pools on the SLB. Setup monitors for the new Management Service.
2. Make the standby Management Services point to the standby Management Repository database by running the following command on the first standby Management Service:


```
emctl config oms -store_repos_details -repos_conndesc
<connect descriptor of standby database> -repos_user sysman
-no_check_db
```
3. Shut down all Management Services by running the following command on each Management Service:


```
emctl stop oms -all
```

18.3.5.4 Switchover

Switchover is a planned activity where operations are transferred from the Primary site to a Standby site. This is usually done for testing and validation of Disaster Recovery (DR) scenarios and for planned maintenance activities on the primary infrastructure. This section describes the steps to switchover to the standby site. The same procedure is applied to switchover in either direction.

Enterprise Manager Console cannot be used to perform switchover of the Management Repository database. Use the Data Guard Broker command line tool DGMGRL instead.

1. Prepare the Standby Database

Verify that recovery is up-to-date. Using the Enterprise Manager Console, you can view the value of the ApplyLag column for the standby database in the Standby Databases section of the Data Guard Overview Page.

2. Shutdown the Primary Enterprise Manager Application Tier

Shutdown all the Management Services in the primary site by running the following command on each Management Service:

```
emctl stop oms -all
```

Shutdown the Enterprise Manager jobs running in Management Repository database:

```
- alter system set job_queue_processes=0;
```

3. Verify Shared Loader Directory / Software Library Availability

Ensure all files from the primary site are available on the standby site.

4. Switchover to the Standby Database

Use DGMGRL to perform a switchover to the standby database. The command can be run on the primary site or the standby site. The switchover command verifies the states of the primary database and the standby database, affects switchover of roles, restarts the old primary database, and sets it up as the new standby database.

```
SWITCHOVER TO <standby database name>;
```

Verify the post switchover states. To monitor a standby database completely, the user monitoring the database must have SYSDBA privileges. This privilege is required because the standby database is in a mounted-only state. A best practice

is to ensure that the users monitoring the primary and standby databases have SYSDBA privileges for both databases.

```
SHOW CONFIGURATION;  
SHOW DATABASE <primary database name>;  
SHOW DATABASE <standby database name>;
```

5. Startup the Enterprise Manager Application Tier

Startup all the Management Services on the standby site:

```
emctl start oms
```

Startup the Enterprise Manager jobs running in Management Repository database on the standby site (the new primary database):

```
- alter system set job_queue_processes=10;
```

6. Relocate Management Services and Management Repository target

The Management Services and Management Repository target is monitored by a Management Agent on one of the Management Services on the primary site. To ensure that the target is monitored after switchover/failover, relocate the target to a Management Agent on the standby site by running the following command on one of the Management Service standby sites.

```
emctl config emrep -agent <agent name> -conn_desc
```

7. Switchover to Standby SLB

Make appropriate network changes to failover your primary SLB to standby SLB that is, all requests should now be served by the standby SLB without requiring any changes on the clients (browser and Management Agents).

This completes the switchover operation. Access and test the application to ensure that the site is fully operational and functionally equivalent to the primary site.

Repeat the same procedure to switchover in the other direction.

18.3.5.5 Failover

A standby database can be converted to a primary database when the original primary database fails and there is no possibility of recovering the primary database in a timely manner. This is known as a manual failover. There may or may not be data loss depending upon whether your primary and target standby databases were synchronized at the time of the primary database failure.

This section describes the steps to failover to a standby database, recover the Enterprise Manager application state by resynchronizing the Management Repository database with all Management Agents, and enabling the original primary database as a standby using flashback database.

The word *manual* is used here to contrast this type of failover with a fast-start failover described later in [Section 18.3.5.6, "Automatic Failover"](#).

1. Verify Shared Loader Directory and Software Library Availability

Ensure all files from the primary site are available on the standby site.

2. Failover to Standby Database

Shutdown the database on the primary site. Use DGMGRL to connect to the standby database and execute the FAILOVER command:

```
FAILOVER TO <standby database name>;
```

Verify the post failover states:

```
SHOW CONFIGURATION;
SHOW DATABASE <primary database name>;
SHOW DATABASE <standby database name>;
```

Note that after the failover completes, the original primary database cannot be used as a standby database of the new primary database unless it is re-enabled.

3. Resync the New Primary Database with Management Agents

Skip this step if you are running in Data Guard Maximum Protection or Maximum Availability level as there is no data loss on failover. However, if there is data loss, synchronize the new primary database with all Management Agents.

On any one Management Service on the standby site, run the following command:

```
emctl resync repos -full -name "<name for recovery action>"
```

This command submits a resync job that would be executed on each Management Agent when the Management Services on the standby site are brought up.

Repository resynchronization is a resource intensive operation. A well tuned Management Repository will help significantly to complete the operation as quickly as possible. Specifically if you are not routinely coalescing the IOTs/indexes associated with Advanced Queueing tables as described in My Oracle Support note 271855.1, running the procedure before resync will significantly help the resync operation to complete faster.

4. Startup the Enterprise Manager Application Tier

Startup all the Management Services on the standby site by running the following command on each Management Service.

```
emctl start oms
```

5. Relocate Management Services and Management Repository target

The Management Services and Management Repository target is monitored by a Management Agent on one of the Management Services on the primary site. To ensure that target is monitored after switchover/failover, relocate the target to a Management Agent on the standby site by running the following command on one of the standby site Management Services.

```
emctl config emrep -agent <agent name> -conn_desc
```

6. Switchover to Standby SLB

Make appropriate network changes to failover your primary SLB to the standby SLB, that is, all requests should now be served by the standby SLB without requiring any changes on the clients (browser and Management Agents).

7. Establish Original Primary Database as Standby Database Using Flashback

Once access to the failed site is restored and if you had flashback database enabled, you can reinstate the original primary database as a physical standby of the new primary database.

- Shutdown all the Management Services in original primary site:

```
emctl stop oms -all
```

- Restart the original primary database in mount state:

```
shutdown immediate;
startup mount;
```

- Reinststate the Original Primary Database

Use DGMGRL to connect to the old primary database and execute the REINSTATE command

```
REINSTATE DATABASE <old primary database name>;
```

- The newly reinstated standby database will begin serving as standby database to the new primary database.
- Verify the post reinststate states:

```
SHOW CONFIGURATION;  
SHOW DATABASE <primary database name>;  
SHOW DATABASE <standby database name>;
```

8. Monitor and complete Repository Resynchronization

Navigate to the Management Services and Repository Overview page of Grid Control Console. Under Related Links, click **Repository Synchronization**. This page shows the progress of the resynchronization operation on a per Management Agent basis. Monitor the progress.

Operations that fail should be resubmitted manually from this page after fixing the error mentioned. Typically, communication related errors are caused by Management Agents being down and can be fixed by resubmitting the operation from this page after restarting the Management Agent.

For Management Agents that cannot be started due to some reason, for example, old decommissioned Management Agents, the operation should be stopped manually from this page. Resynchronization is deemed complete when all the jobs have a completed or stopped status.

This completes the failover operation. Access and test the application to ensure that the site is fully operational and functionally equivalent to the primary site.

Do a switchover procedure if the site operations have to be moved back to the original primary site.

18.3.5.6 Automatic Failover

This section details the steps to achieve complete automation of failure detection and failover procedure by utilizing Fast-Start Failover and Observer process. At a high level the process works like this:

- Fast-Start Failover (FSFO) determines that a failover is necessary and initiates a failover to the standby database automatically
- When the database failover has completed the DB_ROLE_CHANGE database event is fired
- The event causes a trigger to be fired which calls a script that configures and starts Enterprise Manager Application Tier

Perform the following steps:

1. Develop Enterprise Manager Application Tier Configuration and Startup Script

Develop a script that will automate the Enterprise Manager Application configuration and startup process. See the sample shipped with Grid Control in the OH/sysman/ha directory. A sample script for the standby site is included here and should be customized as needed. Make sure `ssh` equivalence is setup so that remote shell scripts can be executed without password prompts. Place the script in

a location accessible from the standby database host. Place a similar script on the primary site.

```
#!/bin/sh
# Script: /scratch/EMSBY_start.sh
# Primary Site Hosts
# Repos: earth, OMS: jupiter1, jupiter2
# Standby Site Hosts
# Repos: mars, # OMS: saturn1, saturn2
LOGFILE="/net/mars/em/failover/em_failover.log"
OMS_ORACLE_HOME="/scratch/OracleHomes/em/oms11"
CENTRAL_AGENT="saturn1.example.com:3872"

#log message
echo "#####" >> $LOGFILE
date >> $LOGFILE
echo $OMS_ORACLE_HOME >> $LOGFILE
id >> $LOGFILE 2>&1

#startup all OMS
#Add additional lines, one each per OMS in a multiple OMS setup
ssh orausr@saturn1 "$OMS_ORACLE_HOME/bin/emctl start oms" >> $LOGFILE 2>&1
ssh orausr@saturn2 "$OMS_ORACLE_HOME/bin/emctl start oms" >> $LOGFILE 2>&1

#relocate Management Services and Repository target
#to be done only once in a multiple OMS setup
#allow time for OMS to be fully initialized
ssh orausr@saturn1 "$OMS_ORACLE_HOME/bin/emctl config emrep -agent $CENTRAL_
AGENT -conn_desc -sysman_pwd <password>" >> $LOGFILE 2>&1

#always return 0 so that dbms scheduler job completes successfully
exit 0
```

2. Automate Execution of Script by Trigger

Create a database event "DB_ROLE_CHANGE" trigger, which fires after the database role changes from standby to primary. See the sample shipped with Grid Control in OH/sysman/ha directory.

```
--
--
-- Sample database role change trigger
--
--
CREATE OR REPLACE TRIGGER FAILOVER_EM
AFTER DB_ROLE_CHANGE ON DATABASE
DECLARE
    v_db_unique_name varchar2(30);
    v_db_role varchar2(30);
BEGIN
    select upper(VALUE) into v_db_unique_name
    from v$parameter where NAME='db_unique_name';
    select database_role into v_db_role
    from v$database;

    if v_db_role = 'PRIMARY' then

        -- Submit job to Resync agents with repository
        -- Needed if running in maximum performance mode
        -- and there are chances of data-loss on failover
        -- Uncomment block below if required
```

```

-- begin
-- SYSMAN.setemusercontext('SYSMAN', SYSMAN.MGMT_USER.OP_SET_
IDENTIFIER);
-- SYSMAN.emd_maintenance.full_repository_resync('AUTO-FAILOVER to '||v_
db_unique_name||' - '||systimestamp, true);
-- SYSMAN.setemusercontext('SYSMAN', SYSMAN.MGMT_USER.OP_CLEAR_
IDENTIFIER);
-- end;

-- Start the EM mid-tier
dbms_scheduler.create_job(
  job_name=>'START_EM',
  job_type=>'executable',
  job_action=>'<location>' || v_db_unique_name|| '_start_oms.sh',
  enabled=>TRUE
);
end if;
EXCEPTION
WHEN OTHERS
THEN
  SYSMAN.mgmt_log.log_error('LOGGING', SYSMAN.MGMT_GLOBAL.UNEXPECTED_ERR,
SYSMAN.MGMT_GLOBAL.UNEXPECTED_ERR_M || 'EM_FAILOVER: ' ||SQLERRM);
END;
/

```

Note: Based on your deployment, you might require additional steps to synchronize and automate the failover of SLB and shared storage used for loader receive directory and software library. These steps are vendor specific and beyond the scope of this document. One possibility is to invoke these steps from the Enterprise Manager Application Tier startup and configuration script.

3. Configure Fast-Start Failover and Observer

Use the Fast-Start Failover configuration wizard in Enterprise Manager Console to enable FSFO and configure the Observer.

This completes the setup of automatic failover.

18.4 Installation Best Practices for Enterprise Manager High Availability

The following sections document best practices for installation and configuration of each Grid Control component.

18.4.1 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.

18.4.2 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 3 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.

18.4.3 Installing the Management Agent Software on Redundant Storage

The Management Agent persists its intermediate state and collected information using local files in the `$AGENT_HOME/$HOSTNAME/sysman/emd` sub tree under the Management Agent home 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.

At a minimum, configure these sub-directories on striped redundant or mirrored storage. Availability would be further enhanced by placing the entire `$AGENT_HOME` on redundant storage. The Management Agent home directory is shown by entering the command `'emctl getemhome'` on the command line, or from the Management Services and Repository tab and Agents tab in the Grid Control console.

18.4.4 Install the Management Service Shared File Areas on Redundant Storage

The Management Service contains results of the intermediate collected data before it is loaded into the Management Repository. The loader receive directory contains these files and is typically empty when the Management Service is able to load data as quickly as it is received. Once the files are received by the Management Service, the Management Agent considers them committed and therefore removes its local copy. In the event that these files are lost before being uploaded to the Management Repository, data loss will occur. At a minimum, configure these sub-directories on striped redundant or mirrored storage. When Management Services are configured for the Shared Filesystem Loader, all services share the same loader receive directory. It is recommended that the shared loader receive directory be on a clustered file system like NetApps Filer.

18.5 Configuration With Grid Control

Grid Control comes preconfigured with a series of default rules to monitor many common targets. These rules can be extended to monitor the Grid Control infrastructure as well as the other targets on your network to meet specific monitoring needs.

18.5.1 Console Warnings, Alerts, and Notifications

The following list is a set of recommendations that extend the default monitoring performed by Enterprise Manager. Use the Notification Rules link on the Preferences page to adjust the default rules provided on the Configuration/Rules page:

- Ensure the Agent Unreachable rule is set to alert on all Management Agents unreachable and Management Agents clear errors.
- Ensure the Repository Operations Availability rule is set to notify on any unreachable problems with the Management Service or Management Repository nodes. Also modify this rule to alert on the Targets Not Providing Data condition and any database alerts that are detected against the database serving as the Management Repository.

Modify the Agent Upload Problems Rule to alert when the Management Service status has hit a warning or clear threshold.

18.5.2 Configure Additional Error Reporting Mechanisms

Enterprise Manager provides error reporting mechanisms through e-mail notifications, PL/SQL packages, and SNMP alerts. Configure these mechanisms based on the infrastructure of the production site. If using e-mail for notifications, configure the notification rule through the Grid Control console to notify administrators using multiple SMTP servers if they are available. This can be done by modifying the default e-mail server setting on the Notification Methods option under Setup.

18.5.3 Component Backup

Backup procedures for the database are well established standards. Configure backup for the Management Repository using the RMAN interface provided in the Grid Control console. Refer to the RMAN documentation or the Maximum Availability architecture document for detailed implementation instructions.

In addition to the Management Repository, the Management Service and Management Agent should also have regular backups. Backups should be performed after any configuration change.

18.5.4 Troubleshooting

In the event of a problem with Grid Control, the starting point for any diagnostic effort is the console itself. The Management System tab provides access to an overview of all Management Service operations and current alerts. Other pages summarize the health of Management Service processes and logged errors. These pages are useful for determining the causes of any performance problems as the summary page shows at a historical view of the amount of files waiting to be loaded to the Management Repository and the amount of work waiting to be completed by Management Agents.

18.5.4.1 Upload Delay for Monitoring Data

When assessing the health and availability of targets through the Grid Control console, information is slow to appear in the UI, especially after a Management Service outage. The state of a target in the Grid Control console may be delayed after a state change on the monitored host. Use the Management System page to gauge backlog for pending files to be processed.

18.5.4.2 Notification Delay of Target State Change

The model used by the Management Agent to assess the state of health for any particular monitored target is poll based. Management Agents immediately post a notification to the Management Service as soon as a change in state is detected. This infers that there is some potential delay for the Management Agent to actually detect a change in state.

18.6 Configuring Oracle Enterprise Manager for Active and Passive Environments

Active and Passive environments, also known as Cold Failover Cluster (CFC) environments, refer to one type of high availability solution that allows an application to run on one node at a time. These environments generally use a combination of *cluster* software to provide a logical host name and IP address, along with interconnected host and storage systems to share information to provide a measure of high availability for applications.

Note: The database for hosting the Management Repository and the WebLogic Server software must be installed before running Grid Control. For information on installing WebLogic Server, refer to *Oracle Enterprise Manager Grid Control Basic Installation Guide*.

This chapter contains the following sections:

- [Using Virtual Host Names for Active and Passive High Availability Environments in Enterprise Manager Database Control](#)
- [Configuring Grid Control Repository in Active/Passive High Availability Environments](#)
- [How to Configure Grid Control OMS in Active/Passive Environment for High Availability Failover Using Virtual Host Names](#)
- [Configuring Targets for Failover in Active/Passive Environments](#)

18.7 Using Virtual Host Names for Active and Passive High Availability Environments in Enterprise Manager Database Control

This section provides information to database administrators about configuring an Oracle Database release 11gR1 in Cold Failover Cluster environments using Enterprise Manager Database Control.

The following conditions must be met for Database Control to service a database instance after failing over to a different host in the cluster:

- The installation of the database must be done using a Virtual IP address.
- The installation must be conducted on a shared disk or volume which holds the binaries, configuration, and runtime data (including the database).
- Configuration data and metadata must also failover to the surviving node.
- Inventory location must failover to the surviving node.
- Software owner and time zone parameters must be the same on all cluster member nodes that will host this database.

The following items are configuration and installation points you should consider before getting started.

- To override the physical host name of the cluster member with a virtual host name, software must be installed using the parameter `ORACLE_HOSTNAME`.
- For inventory pointer, software must be installed using the command parameter `invPtrLoc` to point to the shared inventory location file, which includes the path to the shared inventory location.
- The database software, the configuration of the database, and Database Control are done on a shared volume.

18.7.1 Set Up the Alias for the Virtual Host Name and Virtual IP Address

You can set up the alias for the virtual host name and virtual IP address by either allowing the clusterware to set it up automatically or by setting it up manually 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 similar to `nslookup` and `traceroute` commands can be used to verify the set up.

18.7.2 Set Up Shared Storage

Shared storage can be managed by the clusterware that is in use or you can use any shared file system volume as long as it is supported. The most common shared file system is NFS. You can also use the Oracle Cluster File System software.

18.7.3 Set Up the Environment

Some operating system versions require specific operating system patches to be applied prior to installing release 11gR1 of the Oracle database. You must also have sufficient kernel resources available when you conduct the installation.

Before you launch the installer, specific environment variables must be verified. Each of the following variables must be identically set for the account you are using to install the software on all machines participating in the cluster.

- Operating system variable `TZ`, time zone setting. You should unset this prior to the installation.
- PERL variables. Variables like `PERL5LIB` should be unset to prevent the installation and Database Control from picking up the incorrect set of PERL libraries.
- Paths used for dynamic libraries. Based on the operating system, the variables can be `LD_LIBRARY_PATH`, `LIBPATH`, `SHLIB_PATH`, or `DYLD_LIBRARY_PATH`. These variables should *only* point to directories that are visible and usable on each node of the cluster.

18.7.4 Ensure That the Oracle USERNAME, ID, and GROUP NAME Are Synchronized on All Cluster Members

The user and group of the software owner should be defined identically on all nodes of the cluster. You can verify this using the following command:

```
$ id -a
uid=1234(oracle) gid=5678(dba) groups=5678(dba)
```

18.7.5 Ensure That Inventory Files Are on the Shared Storage

To ensure that inventory files are on the shared storage, follow these steps:

- Create you new ORACLE_HOME directory.
 - Create the Oracle Inventory directory under the new Oracle home
- ```
cd <shared oracle home>
mkdir oraInventory
```
- Create the oraInst.loc file. This file contains the Inventory directory path information required by the Universal Installer:
    1. vi oraInst.loc
    2. Enter the path information to the Oracle Inventory directory and specify the group of the software owner as the dba user. For example:

```
inventory_loc=/app/oracle/product/11.1/oraInventory
inst_group=dba
```

Depending on the type of operating system, the default directory for the oraInst.loc file is either /etc (for example, on Linux) or /var/opt/oracle (for example, on Solaris and HP-UX).

## 18.7.6 Start the Installer

To start the installer, point to the inventory location file oraInst.loc, and specify the host name of the virtual group. The debug parameter in the example below is optional:

```
$ export ORACLE_HOSTNAME=lxdb.acme.com
$ runInstaller -invPtrloc /app/oracle/share1/oraInst.loc ORACLE_
HOSTNAME=lxdb.acme.com -debug
```

### 18.7.6.1 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.

## 18.7.7 Start Services

You must start the services in the following order:

1. Establish IP address on the active node
2. Start the TNS listener
3. Start the database
4. Start dbconsole
5. Test functionality

In the event that services do not start, do the following:

1. Establish IP on failover box
2. Start TNS listener  
`lsnrctl start`
3. Start the database  
`dbstart`
4. Start Database Control  
`emctl start dbconsole`
5. Test functionality

To manually stop or shutdown a service, follow these steps:

1. Stop the application.
2. Stop Database Control  
`emctl stop dbconsole`
3. Stop TNS listener  
`lsnrctl stop`
4. Stop the database  
`dbshut`
5. Stop IP

## 18.8 Configuring Grid Control Repository in Active/Passive High Availability Environments

In order for Grid Control repository to fail over to a different host, the following conditions must be met:

- The installation must be conducted using a Virtual Hostname and an associated unique IP address
- Installation must occur on a shared disk/volume which holds the binaries, the configuration, and the runtime data (including the repository database)
- Configuration data and metadata must also failover to the surviving node
- Inventory location must failover to the surviving node
- Software owner and time zone parameters must be the same on all cluster member nodes that will host this OMS

### 18.8.1 Installation and Configuration

The following installation and configuration requirements should be noted:

- To override the physical host name of the cluster member with a virtual host name, software must be installed using the parameter `ORACLE_HOSTNAME`.
- For inventory pointer, 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.

- The database software, the configuration of the database, and data files are on a shared volume.

If you are using an NFS mounted volume for the installation, ensure that you specify `rsize` and `wsize` in your mount command to prevent I/O issues. See My Oracle Support note 279393.1 Linux.NetApp: RHEL/SUSE Setup Recommendations for NetApp Filer Storage.

Example:

```
grid-repo.acme.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* could also be true for non-shared failover volumes, which can be mounted on active hosts after failover.

---

## 18.8.2 Set 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 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 commands listed below:

```
nslookup <virtual hostname>
```

This command returns the virtual IP address and fully 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 to verify that the correct information is returned.

## 18.8.3 Set Up the Environment

Some operating system versions require specific patches to be applied prior to installing 11gR1. The user installing and using the 11gR1 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 must be verified. Each of these variables must be set up identically 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 prevent inadvertently picking up the wrong set of PERL libraries.
- Same operating system, operating system patches, and version of the kernel.
  - Therefore, RHEL 3 and RHEL 4 are *not* allowed for a CFC system.
- System libraries

For example, `LIBPATH`, `LD_LIBRARY_PATH`, `SHLIB_PATH`, and so on. The same system libraries must be present.

### 18.8.4 Synchronize Operating System User 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)
```

### 18.8.5 Set Up Inventory

You can set up the inventory by using the following steps:

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 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:

Example:

```
inventory_loc=/app/oracle/product/11.1/oraInventory
inst_group=oinstall
```

### 18.8.6 Install the Software

Follow these steps to install 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:

```
$ export ORACLE_HOSTNAME=grid-repo.acme.com
$ runInstaller -invPtrLoc /app/oracle/share1/oraInst.loc ORACLE_
HOSTNAME=grid-repo.acme.com
```

3. Install the repository DB software only on the shared location. For example:

```
/oradbnas/app/oracle/product/oradb111 using Host1
```
4. Start DBCA and create all the data files be on the shared location. For example:

```
/oradbnas/oradata
```
5. Continue the rest of the installation normally.
6. Once completed, copy the files `oraInst.loc` and `oratab` to `/etc`. Also copy `/opt/oracle` to all cluster member hosts (`Host2`, `Host3`, and so on).

### 18.8.6.1 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.

## 18.8.7 Startup of Services

Be sure you start your services in the proper order:

1. Establish 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

In case of failover, follow these steps:

1. Establish IP address on the failover box
2. Start TNS listener (`lsnrctl start`) if it is part of the same failover group
3. Start the database (`dbstart`) if it is part of the same failover group

## 18.8.8 Summary

The Grid Control Management Repository can now be deployed in a CFC environment that utilizes a floating host name.

To deploy the OMS midtier in a CFC environment, please see [Section 18.9, "How to Configure Grid Control OMS in Active/Passive Environment for High Availability Failover Using Virtual Host Names"](#).

## 18.9 How to Configure Grid Control OMS in Active/Passive Environment for High Availability Failover Using Virtual Host Names

This section provides a general reference for Grid Control administrators who want to configure Enterprise Manager 11gR1 Grid Control in Cold Failover Cluster (CFC) environments.

### 18.9.1 Overview and Requirements

The following conditions must be met for Grid 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, the configuration and the runtime data (including the recv directory).
- Configuration data and metadata must also failover to the surviving node.
- Inventory location must failover to the surviving node.

- Software owner and time zone parameters must be the same on all cluster member nodes that will host this Oracle Management Service (OMS).

## 18.9.2 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`. For inventory pointer, 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.acme.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.

---

---

## 18.9.3 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.

## 18.9.4 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 as long as it is not an unsupported type, such as OCFS V1. The most common shared file system is NFS.

---

---

**Note:** 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.

---

---

## 18.9.5 Setting Up the Environment

Some operating system versions require specific operating system patches be applied prior to installing 11gR1. The user installing and using the 11gR1 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

## 18.9.6 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)
```

## 18.9.7 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 Inventory directory path information needed by the Universal Installer.
  - a. `vi oraInst.loc`
  - b. 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/product/11.1/oraInventory
inst_group=oinstall
```

## 18.9.8 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. Install WebLogic Server. For information on installing WebLogic Server, refer to *Oracle Enterprise Manager Grid Control Basic Installation Guide*.
3. 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:

```
$ export ORACLE_HOSTNAME=lxdb.acme.com
```

```
$ runInstaller -invPtrloc /app/oracle/share1/oraInst.loc
ORACLE_HOSTNAME=lxdb.acme.com -debug
```

4. Install Oracle Management Services on cluster member *Host1* using the option, "EM install using the existing DB"
5. Continue the remainder of the installation normally.
6. Once completed, copy the files *oraInst.loc* and *oratab* to */etc*. Also copy */opt/oracle* to all cluster member hosts (*Host2*, *Host3*, and so on).

#### 18.9.8.1 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.

### 18.9.9 Starting Up Services

Ensure that you start your services in the proper order. Use the order listed below:

1. Establish 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 Grid Control using `emctl start oms`
5. Test functionality

In case of failover, refer to the following steps:

1. Establish IP on failover box
2. Start TNS listener using the command `lsnrctl start` if it is part of the same failover group
3. Start the database using the command `dbstart` if it is part of the same failover group
4. Start Grid Control using the command `emctl start oms`
5. Test the functionality

### 18.9.10 Summary

The OMS mid-tier component of Grid Control can now be deployed in a CFC environments that utilize a floating host name.

To deploy the repository database in a CFC environment, see [Section 18.8, "Configuring Grid Control Repository in Active/Passive High Availability Environments"](#).

## 18.10 Configuring Targets for Failover in Active/Passive Environments

This section provides a general reference for Grid Control administrators who want to relocate Cold Failover Cluster (CFC) targets from one existing Management Agent to another. Although the targets are capable of running on multiple nodes, these targets run only on the active node in a CFC environment.

CFC environments generally use a combination of cluster software to provide a virtual host name and IP address along with interconnected host and storage systems to share information and provide high availability for applications. Automating failover of the virtual host name and IP, in combination with relocating the Enterprise Manager targets and restarting the applications on the passive node, requires the use of Oracle Enterprise Manager command-line interface (EM CLI) and Oracle Clusterware (running Oracle Database release 10g or 11g) or third-party cluster software. Several Oracle partner vendors provide clusterware solutions in this area.

The Enterprise Manager Command Line Interface (EM CLI) allows you to access Enterprise Manager Grid Control functionality from text-based consoles (terminal sessions) for a variety of operating systems. Using EM CLI, you can perform Enterprise Manager Grid Control console-based operations, like monitoring and managing targets, jobs, groups, blackouts, notifications, and alerts. See the *Oracle Enterprise Manager Command Line Interface* manual for more information.

### 18.10.1 Target Relocation in Active/Passive Environments

Beginning with Oracle Enterprise Manager 10g release 10.2.0.5, a single Oracle Management Agent running on each node in the cluster can monitor targets configured for active / passive high availability. Only one Management Agent is required on each of the physical nodes of the CFC cluster because, in case of a failover to the passive node, Enterprise Manager can move the HA monitored targets from the Management Agent on the failed node to another Management Agent on the newly activated node using a series of EMCLI commands. See the *Oracle® Enterprise Manager Command Line Interface* manual for more information.

If your application is running in an active/passive environment, the clusterware brings up the applications on the passive node in the event that the active node fails. For Enterprise Manager to continue monitoring the targets in this type of configuration, the existing Management Agent needs additional configuration.

The following sections describe how to prepare the environment to automate and restart targets on the new active node. Failover and fallback procedures are also provided.

### 18.10.2 Installation and Configuration

The following sections describe how to configure Enterprise Manager to support a CFC configuration using the existing Management Agents communicating with the Oracle Management Service processes:

- [Prerequisites](#)
- [Configuration Steps](#)

#### 18.10.2.1 Prerequisites

Prepare the Active/Passive environments as follows:

- Ensure the operating system clock is synchronized across all nodes of the cluster. (Consider using Network Time Protocol (NTP) or another network synchronization method.)
- Use the EM CLI RELOCATE\_TARGETS command only with Enterprise Manager Release 10.2.0.5 (and higher) Management Agents.

### 18.10.2.2 Configuration Steps

The following steps show how to configure Enterprise Manager to support a CFC configuration using the existing Management Agents that are communicating with the OMS processes. The example that follows is based on a configuration with a two-node cluster that has one failover group. For additional information about targets running in CFC active/passive environments, see My Oracle Support note 406014.1.

#### 1. Configure EM CLI

To set up and configure target relocation, use the Oracle Enterprise Manager command-line interface (EM CLI). See the *Oracle Enterprise Manager Command Line Interface* manual and the *Oracle Enterprise Manager Extensibility* manual for information about EM CLI and Management Plug-Ins.

#### 2. Install Management Agents

Install the Management Agent on a local disk volume on each node in the cluster. Once installed, the Management Agents are visible in the Grid Control console.

#### 3. Discover Targets

After the Active / Passive targets have been configured, use the Management Agent discovery screen in the Grid Control console to add the targets (such as database, listener, application server, and so on). Perform the discovery on the active node, which is the node that is currently hosting the new target.

## 18.10.3 Failover Procedure

To speed relocation of targets after a node failover, configure the following steps using a script that contains the commands necessary to automatically initiate a failover of a target. Typically, the clusterware software has a mechanism with which you can automatically execute the script to relocate the targets in Enterprise Manager. Also, see [Section 18.10.6, "Script Examples"](#) for sample scripts.

#### 1. Shut down the target services on the failed active node.

On the active node where the targets are running, shut down the target services running on the virtual IP.

#### 2. If required, disconnect the storage for this target on the active node.

Shut down all the applications running on the virtual IP and shared storage.

#### 3. Enable the target's IP address on the new active node.

#### 4. If required, connect storage for the target on the currently active node.

#### 5. Relocate the targets in Grid Control using EM CLI.

To relocate the targets to the Management Agent on the new active node, issue the EM CLI RELOCATE TARGET command for each target type (listener, application servers, and so on) that you must relocate after the failover operation. For example:

```
emcli relocate_targets
-src_agent=<node 1>:3872
```

```
-dest_agent=<node 2>:3872
-target_name=<database_name>
-target_type=oracle_database
-copy_from_src
-force=yes
```

In the example, port 3872 is the default port for the Management Agent. To find the appropriate port number for your configuration, use the value for the EMD\_URL parameter in the emd.properties file for this Management Agent.

**Note:** In case of a failover event, the source agent will not be running. However, there is no need to have the source Management Agent running to accomplish the RELOCATE operation. EM CLI is an OMS client that performs its RELOCATE operations directly against the Management Repository.

### 18.10.4 Fallback Procedure

To return the HA targets to the original active node or to any other cluster member node:

1. Repeat the steps in [Section 18.10.3, "Failover Procedure"](#) to return the HA targets to the active node.
2. Verify the target status in the Grid Control console.

### 18.10.5 EM CLI Parameter Reference

Issue the same command for each target type that will be failed over to (or be switched over) during relocation operations. For example, issue the same EM CLI command to relocate the listener, the application servers, and so on. [Table 18–6](#) shows the EM CLI parameters you use to relocate targets:

**Table 18–6** EM CLI Parameters

| EM CLI Parameter | Description                                                                                                                                                                                                                 |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -src_agent       | Management Agent on which the target was running before the failover occurred.                                                                                                                                              |
| -dest_agent      | Management Agent that will be monitoring the target after the failover.                                                                                                                                                     |
| -target_name     | Name of the target to be failed over.                                                                                                                                                                                       |
| -target_type     | Type of target to be failed over (internal Enterprise Manager target type). For example, the Oracle database (for a standalone database or an Oracle RAC instance), the Oracle listener for a database listener, and so on. |
| -copy_from_src   | Use the same type of properties from the source Management Agent to identify the target. This is a <b>MANDATORY</b> parameter! If you do not supply this parameter, you can corrupt your target definition!                 |
| -force           | Force dependencies (if needed) to failover as well.                                                                                                                                                                         |

### 18.10.6 Script Examples

The following sections provide script examples:

- [Relocation Script](#)
- [Start Listener Script](#)

## ■ Stop Listener Script

### 18.10.6.1 Relocation Script

```
#!/bin/ksh

#get the status of the targets

emcli get_targets -
targets="db1:oracle_database;listener_db1:oracle_listener" -noheader

 if [[$? != 0]]; then exit 1; fi

blackout the targets to stop false errors. This blackout is set to expire in 30
minutes

emcli create_blackout -name="relocating active passive test targets" -
add_targets="db1:oracle_database;listener_db1:oracle_listener" -
reason="testing failover" -
schedule="frequency:once;duration:0:30"
 if [[$? != 0]]; then exit 1; fi

stop the listener target. Have to go out to a OS script to use the 'lsnrctl set
current_listener' function

emcli execute_hostcmd -cmd="/bin/ksh" -osscript="FILE" -
input_file="FILE:/scratch/oraha/cfc_test/listener_stop.ksh" -
credential_set_name="HostCredsNormal" -
targets="host1.us.oracle.com:host"
 if [[$? != 0]]; then exit 1; fi

now, stop the database

emcli execute_sql -sql="shutdown abort" -
targets="db1:oracle_database" -
credential_set_name="DBCredsSYSDBA"
 if [[$? != 0]]; then exit 1; fi

relocate the targets to the new host

emcli relocate_targets -
src_agent=host1.us.oracle.com:3872 -
dest_agent=host2.us.oracle.com:3872 -
target_name=db1 -target_type=oracle_database -
copy_from_src -force=yes -
changed_param=MachineName:host1vip.us.oracle.com
 if [[$? != 0]]; then exit 1; fi

emcli relocate_targets -
src_agent=host1.us.oracle.com:3872 -
dest_agent=host2.us.oracle.com:3872 -
target_name=listener_db1 -target_type=oracle_listener -
copy_from_src -force=yes -
changed_param=MachineName:host1vip.us.oracle.com
 if [[$? != 0]]; then exit 1; fi

Now, restart database and listener on the new host

emcli execute_hostcmd -cmd="/bin/ksh" -osscript="FILE" -
input_file="FILE:/scratch/oraha/cfc_test/listener_start.ksh" -
```

```

credential_set_name="HostCredsNormal" -
targets="host2.us.oracle.com:host"
 if [[$? != 0]]; then exit 1; fi

emcli execute_sql -sql="startup" -
targets="db1:oracle_database" -
credential_set_name="DBCredsSYSDBA"
 if [[$? != 0]]; then exit 1; fi

Time to end the blackout and let the targets become visible

emcli stop_blackout -name="relocating active passive test targets"
 if [[$? != 0]]; then exit 1; fi

and finally, recheck the status of the targets

emcli get_targets -
targets="db1:oracle_database;listener_db1:oracle_listener" -noheader
 if [[$? != 0]]; then exit 1; fi

```

### 18.10.6.2 Start Listener Script

```

#!/bin/ksh

export
ORACLE_HOME=/oradbshare/app/oracle/product/11.1.0/db
export PATH=$ORACLE_HOME/bin:$PATH

lsnrctl << EOF
set current_listener listener_db1
start
exit
EOF

```

### 18.10.6.3 Stop Listener Script

```

#!/bin/ksh

export
ORACLE_HOME=/oradbshare/app/oracle/product/11.1.0/db
export PATH=$ORACLE_HOME/bin:$PATH

lsnrctl << EOF
set current_listener listener_db1
stop
exit
EOF

```



---

# Management Agent and Management Services

This chapter describes how to reconfigure Enterprise Manager if you later revisit your configuration decisions after you have installed the software.

- [Reconfiguring the Oracle Management Agent](#)
- [Reconfiguring the Oracle Management Service](#)

## 19.1 Reconfiguring the Oracle Management Agent

The following sections describe reconfiguration and tuning changes you can make to the Management Agent after you have installed Enterprise Manager. Refer to the following sections for more information:

- [Configuring the Management Agent to Use a New Management Service](#)
- [Changing the Management Agent Port](#)
- [Controlling the Amount of Disk Space Used by the Management Agent](#)
- [About the Management Agent Watchdog Process](#)
- [Setting the Management Agent Time Zone](#)
- [Adding Trust Points to the Management Agent Configuration](#)

### 19.1.1 Configuring the Management Agent to Use a New Management Service

When you install the Management Agent on a managed host, you associate the Management Agent with a particular Management Service. The Management Agent uses the Management Service URL address and port to identify and communicate with the Management Service.

After you install the Management Agent, you can later reconfigure the Management Agent so it is associated with a different Management Service. Reconfiguring the Management Agent requires no changes to the Management Service. The reconfigured Management Agent will begin communicating with the new Management Service after the Management Agent is restarted.

If you are associating the Management Agent with a new Management Service that is locked in secure mode, then first associate the Management Agent with the new Management Service and then secure the Management Agent.

To associate the Management Agent with a new Management Service after you have installed the Management Agent:

1. Stop the Management Agent.
2. Locate the `emd.properties` file in the Management Agent home directory:  
`AGENT_HOME/sysman/config/emd.properties`
3. Use a text editor to open the file and locate the `REPOSITORY_URL` property.
4. Modify the value for the `REPOSITORY_URL` property so it references the new Management Service. For example:  
`REPOSITORY_URL=http://mgmthost2.acme.com:4889/em/upload`
5. Modify the value for the `emdWalletSrcUrl` property so it references the new Management Service. For example, if the new Management Service is on a host called `mgmthost2.acme.com`, modify the property as follows:  
`emdWalletSrcUrl=http://mgmthost2.acme.com:4889/em/wallets/emd`
6. Save your changes and close the `emd.properties` file.
7. To ensure that the Management Agent is no longer holding any specific data or settings from the previous Management Service, delete all the files in the following directories:

```
AGENT_HOME/sysman/emd/upload/
AGENT_HOME/sysman/emd/state/
AGENT_HOME/sysman/emd/collection/*
AGENT_HOME/sysman/emd/lastupld.xml
AGENT_HOME/sysman/emd/agtstmp.txt
AGENT_HOME/sysman/emd/blackouts.xml
AGENT_HOME/sysman/emd/protocol.ini
```

Note that this action removes all user-defined metrics (UDM)s and custom changes to metric and policy collections.

---

---

**Note:** You can use the `emctl clearstate agent` command to delete the files in the state directory.

---

---

8. Restart the Management Agent.

## 19.1.2 Securing the Management Agent

To secure the Management Agent of the new Management Service, use the following command:

```
emctl secure agent [registration password] [-emdWalletSrcUrl <url>]
```

## 19.1.3 Changing the Management Agent Port

The Management Agent uses a predefined port number to receive requests from the Management Service. This port number is defined by default when you install the Management Agent on a managed host. If you later need to modify this port, you can use the following procedure. You might need to modify this port number if you have existing software that uses the default Management Agent port.

To change the Management Agent port:

1. Stop the Management Agent.

2. Locate the `emd.properties` file in the Management Agent home directory:

```
AGENT_HOME/sysman/config/emd.properties
```

3. Use a text editor to open the file and locate the `EMD_URL` property.

For example:

```
EMD_URL=http://managed_host1.acme.com:1813/emd/main
```

4. Modify the port number in the `EMD_URL` property so the Management Agent uses a new unused port on the managed host.

For example:

```
EMD_URL=http://managed_host1.acme.com:1913/emd/main
```

You can also use the `netstat` command to check for the unused port:

On Windows:

```
netstat -an | findstr <new port number>
```

On UNIX:

```
netstat -an | grep <new port number>
```

5. Start the Management Agent.

---



---

**Note:** After the changed URL is processed, the old Management Agent should not have any targets. Oracle recommends that you remove the old Management Agent from the Management Service to ensure that there are no unwanted targets in the Grid Control console.

---



---

## 19.1.4 Controlling the Amount of Disk Space Used by the Management Agent

Oracle designed the Management Agent to work within a set of disk space limits. These limits prevent the Management Agent from using too much disk space and causing performance or resource issues on your enterprise systems. However, if disk space becomes an issue, you can adjust the default settings that are used to control the amount of disk space used by the Management Agent.

As the Management Agent on a particular host gathers management data about the targets on the host, it saves the collected data on the local disk until the data is uploaded to the Management Repository. The Management Agent saves this collected data and metadata in the following directory:

```
AGENT_HOME/sysman/emd/upload
```

By default, the Management Agent will save up to 50MB of collected data in the upload directory. If the amount of collected data exceeds 50MB, data collection is stopped temporarily until the data is uploaded to the repository and more disk space becomes available.

To verify how much space is available:

- Use the `emctl status agent` command. For example:

```
Available disk space on upload filesystem : 1.18%
Collection Status : Disabled by Upload Manager
Last successful heartbeat to OMS : 2007-07-31 11:22:07
```

- Investigate the `<AGENT_HOME>/sysman/log/emagent.trc` file. The file will have errors such as :

```
24.995519 MB Data. 34.06% of disk used. Disabling collections.
2006-10-19 10:41:23 Thread-19 WARN collector: Disable collector
```

In addition, the Management Agent checks to be sure that the percentage of disk space currently in use on the local disk does not exceed 98 percent. If this value is exceeded, the Management Agent stops collecting data and stops saving information to the Management Agent log and trace files.

You can modify these default settings as follows:

1. Stop the Management Agent.
2. Locate the `emd.properties` file in the Management Agent home directory:  
`AGENT_HOME/sysman/config/emd.properties`
3. Use a text editor to open the file and modify the entries shown in [Table 19-1](#).
4. Save your changes and exit the file.
5. Restart the Management Agent.

**Table 19-1 Properties for Controlling the Disk Space Used by the Management Agent**

| Property                               | Explanation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>UploadMaxBytesXML</code>         | Use this property in the <code>emd.properties</code> file to specify the maximum number of megabytes (MB) used by the collected data in the Management Agent upload directory. When this limit is exceeded, the Management Agent will stop collecting additional management data until the next upload to the Management Repository reduces the amount of collected data in the upload directory.                                                                                                                                                   |
| <code>UploadMaxDiskUsedPct</code>      | Use this property in the <code>emd.properties</code> file to specify the maximum percentage of disk space that can be in use on the local disk before the Management Agent temporarily stops collecting additional data and stops saving information to the Management Agent log and trace files.<br><br>The Management Agent will begin collecting data again when the percentage of disk space in use falls to less than the percentage specified in the <code>UploadMaxDiskUsedPctFloor</code> property in the <code>emd.properties</code> file. |
| <code>UploadMaxDiskUsedPctFloor</code> | Use this property in the <code>emd.properties</code> file to specify the amount (%) of disk space that can be used on the EMD filesystem before the following items are re-enabled after being disabled: <ul style="list-style-type: none"> <li>■ Collection of data (upload manager)</li> <li>■ Logging and tracing</li> <li>■ Diagnosability tracing does minimal dumps above this limit</li> </ul>                                                                                                                                               |
| <code>UploadMaxNumberXML</code>        | Use this property in the <code>emd.properties</code> files to specify the maximum number of files the upload manager will support in the upload directory. When this limit is exceeded, the Management Agent will temporarily disable collections, logging, and tracing.                                                                                                                                                                                                                                                                            |

### 19.1.5 About the Management Agent Watchdog Process

The Management Agent is the Enterprise Manager component that gathers the data you need to manage your enterprise efficiently. As a result, Enterprise Manager includes software that keeps track of the Management Agent processes and makes sure the Management Agent stays running.

For example, if the Management Agent quits unexpectedly, this self-monitoring process—referred to as the watchdog process—will restart the Management Agent automatically.

In most situations, the watchdog process works in the background and requires no configuration or maintenance. The watchdog process is controlled by the `emwd.pl` script located in the following directory of the Management Agent home directory:

```
AGENT_HOME/bin
```

You can identify the watchdog process by using the following commands:

```
$PROMPT> ps -ef | grep emwd
```

## 19.1.6 Setting the Management Agent Time Zone

In today's global economy, it is not uncommon for the systems you manage to reside in multiple locations throughout the world. For example, if your company headquarters are in New Hampshire, USA, you may need to manage systems that reside in California, Canada, and in Europe.

As Enterprise Manager collects monitoring data from Management Agents running on these remote systems, it is important that the data is correlated accurately. A software failure on a machine in Ontario, Canada might be the cause of a performance problem on a machine in Hoboken, New Jersey.

To correlate this data, it is important that Enterprise Manager obtains the correct time zone for each Management Agent that you install. The following sections describe how the Management Agent obtains the time zone and how to correct the problem if the time zone for a Management Agent is incorrect:

- [Understanding How the Management Agent Obtains Time Zone Information](#)
- [Resetting the Time Zone of the Management Agent Due to Inconsistency of Time Zones](#)
- [Troubleshooting Management Agent Time Zone Problems](#)

### 19.1.6.1 Understanding How the Management Agent Obtains Time Zone Information

When you install the Management Agent, the software attempts to obtain the current time zone of the host computer. If successful, the installation procedure updates the `agentTZRegion` property setting in the following configuration file:

```
AGENT_HOME/sysman/config/emd.properties
```

The `agentTZRegion` property can be set to any of the values listed in the following file, which is installed in the Management Agent home directory:

```
AGENT_HOME/sysman/admin/supportedtzs.lst
```

To reconfigure a different time zone, perform the following steps. These steps assume that the original time zone used was EST and the target time zone is CST.

1. Set the environment correctly.
  - On Windows XP
    - From the Start Menu, access the Control Panel. Click **Date and Time**, then click the **Time Zone** tab.
    - Select GMT-06:00 Central Time (US & Canada) from the list.

- Click **OK**.
- Open a command line screen (cmd.exe).
- Set the following environment variables:

```
SET TZ=CST
SET ORACLE_HOME=< your oracle home directory >
SET PATH=%ORACLE_HOME%\bin;%PATH%
```

- On UNIX

- Login to your UNIX server as the Oracle user.
- Set the following environment variables:

```
$ export TZ=CST
$ export ORACLE_HOME=< your oracle home directory >
$ export PATH=%ORACLE_HOME%\bin;%PATH%
```

2. Execute the following commands:

- On Windows

```
%ORACLE_HOME%\bin\emctl config agent getTZ
%ORACLE_HOME%\bin\emctl stop iasconsole
%ORACLE_HOME%\bin\emctl resetTZ agent
```

- On UNIX

```
$/ORACLE_HOME/bin/emctl config agent getTZ
$/ORACLE_HOME/bin/emctl stop iasconsole
$/ORACLE_HOME/bin/emctl resetTZ agent
```

3. Delete all the files under the following directory:

- On Windows

```
%ORACLE_HOME%\sysman\logs
```

- On UNIX

```
$/ORACLE_HOME$/sysman/logs
```

4. Start the console again.

- On Windows

```
%ORACLE_HOME%\bin\emctl start iasconsole
%ORACLE_HOME%\bin\emctl status iasconsole
%ORACLE_HOME%\bin\emctl status agent
%ORACLE_HOME%\bin\emctl config agent getTZ
```

- On UNIX

```
$/ORACLE_HOME/bin/emctl start iasconsole
$/ORACLE_HOME/bin/emctl status iasconsole
$/ORACLE_HOME/bin/emctl status agent
$/ORACLE_HOME/bin/emctl config agent getTZ
```

5. Check the timestamp in the log file.

6. Check the em.properties file. The agentTZRegion parameter should now look like this:

- On Windows

```
%ORACLE_HOME%\sysman\config\em.properties
agentTZRegion=America/Chicago
```

- On UNIX

```
$ORACLE_HOME/sysman/config/em.properties
agentTZRegion=America/Chicago
```

### 19.1.6.2 Resetting the Time Zone of the Management Agent Due to Inconsistency of Time Zones

You need to reset the time zone of the Management Agent when *both* of the following situations are true:

- The Management Agent has been running with a particular time zone
- Subsequently a change occurs to the time zone of the host where the Management Agent is running

To propagate the time zone change to the `emd.properties` file, perform the following:

1. Execute the following script:

```
ORACLE_HOME/bin/emctl resetTZ agent
```

This script updates `ORACLE_HOME/sysman/config/emd.properties` so that the value of `agentTZRegion` matches that of the current time zone setting of the machine.

---

**Note:** The location of the `emd.properties` file depends on the Control Console being used:

- For the Database Control Console, the location is usually:  
`ORACLE_HOME/<host>_<sid>/sysman/config`
  - For the Application Server Control Console, the location is:  
`ORACLE_HOME/sysman/config`
  - For the Grid Control Management Agent, the location is  
`ORACLE_HOME/sysman/config`
  - For the Real Application Cluster central Management Agent, the location is usually: `ORACLE_HOME/<host>/sysman/config`
- 

2. In addition, this command prompts you to run a script against the Enterprise Manager Repository. You must log in to the database as the Enterprise Manager repository user and run the script `mgmt_target.set_agent_tzrgn`. An example follows:

```
SQL> exec mgmt_target.set_agent_tzrgn('em.oracle.com:1830','PST8PDT');
SQL> commit;
SQL> exit
```

`em.oracle.com:1830` represents the name of the emd target.

### 19.1.6.3 Troubleshooting Management Agent Time Zone Problems

Sometimes, during the Management Agent installation, the time zone detected by the Management Agent configuration tool is not recognized by the Management Agent. In

other words, the time zone obtained by the configuration tool is not listed in the Management Agent list of supported time zones.

This problem prevents the Management Agent from starting and results in an error similar to the following:

```
Could not determine agent time zone. Please refer to the file:
ORACLE_HOME/sysman/admin/supportedtzs.lst and pick a timezone region with a
standard offset of +5:0 from GMT and update the property 'agentTZRegion' in the
file: ORACLE_HOME/sysman/config/emd.properties
```

This error appears in one of the log files shown in [Table 19-2](#), depending upon which Enterprise Manager product you are using.

**Table 19-2 Location of Time Zone Error in the Enterprise Manager Log Files**

| <b>If you are using...</b>         | <b>Look for the Time Zone Error in This File...</b> |
|------------------------------------|-----------------------------------------------------|
| Grid Control Console               | emagent.nohup                                       |
| Application Server Control Console | em.nohup                                            |
| Database Control Console           | emdb.nohup                                          |

To configure the Management Agent to use a valid time zone:

1. Enter the following command in the Management Agent home directory to identify the time zone currently being used by the host computer:
 

```
AGENT_HOME/bin/emctl config agent getTZ
```
2. Note the time zone that is returned by the `emctl config agent getTZ` command.
 

This is the time zone of the host computer.
3. Use a text editor to open the following file in the Management Agent home directory:
 

```
AGENT_HOME/sysman/admin/supportedtzs.lst
```

This file contains a list of all the time zones supported by the Management Agent.
4. Browse the contents of the `supportedtzs.lst` file and note the supported time zone closest to the time zone of the host computer.
5. Use a text editor to open the following Management Agent configuration file:
 

```
AGENT_HOME/sysman/config/emd.properties
```
6. Locate the following property near the end of the `emd.properties` file:
 

```
agentTZRegion=
```
7. Set the value of this property to the time zone you identified as closest to the host time zone in the `supportedtzs.lst` file.
 

For example:

```
agentTZRegion=Europe/Warsaw
```
8. Save your changes and close the `emd.properties` file.
 

You should now be able to start the Management Agent without generating the error in the log file.

## 19.1.7 Adding Trust Points to the Management Agent Configuration

Perform these steps to add the relevant security certificate:

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 as follows:

```
-----BEGIN CERTIFICATE-----
MIIDBzCCAnCgAw...
..... base 64 certificate content
-----END CERTIFICATE-----
```

2. In the Oracle Home of the Management Agent monitoring the wallet, run the following command 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
```

## 19.2 Reconfiguring the Oracle Management Service

The following sections describe configuration changes you can make to the Management Service after you install Enterprise Manager:

- [Configuring the Management Service to Use a New Management Repository](#)
- [Configuring the Management Service to Prompt You When Using Execute Commands](#)
- [Troubleshooting Management Service Time Zone Problems](#)

### 19.2.1 Configuring the Management Service to Use a New Management Repository

When you install and deploy the Management Service, you associate the Management Service with a Management Repository. The Management Service uses the database host, database system identifier (SID), database port, management user, and management password to identify and communicate with the Repository.

The following sections describe how to modify the repository information and provide details about how Enterprise Manager keeps the Management Repository password secure.

#### 19.2.1.1 Changing the Repository Properties

To associate the Management Service with a new repository, you must modify the repository properties:

1. Stop the Management Service.
2. Use the following commands:

```
emctl config oms -list_repos_details
emctl config oms -store_repos_details (-repos_host <host> -repos_port <port>
-repos_sid <sid> | -repos_comndesc <connect descriptor>) -repos_user <username>
[-repos_pwd <pwd>] [-no_check_db]
```

3. [Example 19–1](#) shows sample entries.
4. Restart the Management Service.

**Table 19–3 Repository Properties**

| Property                              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -repos_user <username>                | The Management Repository user name. The default value is SYSMAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| -repos_pwd <pwd>                      | The Management Repository password. See <a href="#">"About Changing the Repository Password"</a> on page 19-10 for information of how to change the password value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| -repos_conndesc <connect description> | The Management Repository Oracle Net Connect String for the repository database. The values specified for properties -repos_sid, -repos_host, and -repos_port must be the same as that of HOST, PORT, and SERVICE_NAME in the connect string. If this property is not specified, then -repos_sid, -repos_host, and -repos_port properties are used to construct the connect descriptor. If the database hosting the repository is a RAC database, then the value must be configured as explained in <a href="#">"Configuring the Management Services"</a> on page 18-7.<br><br><b>Note:</b> Connect descriptor should be enclosed in quotes, for example, "<conn_descr>" or "<conn_descr>" |
| -repos_sid <sid>                      | The System Identifier (SID) for the database where the Management Repository schema resides.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| -repos_host <host>                    | The name of the server or host computer where the repository database resides.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| -repos_port <port>                    | The port number for the repository database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Example 19–1 Sample Repository Properties**

```
oracle.sysman.eml.mntr.emdRepUser=SYSMAN
oracle.sysman.eml.mntr.emdRepPwd=sysman
oracle.sysman.eml.mntr.emdRepConnectDescriptor=(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=system12.mycompany.com)(PORT=1521))
(CONNECT_DATA=(SERVICE_NAME=oemrep1)))
oracle.sysman.eml.mntr.emdRepSID=oemrep1
oracle.sysman.eml.mntr.emdRepServer=system12.mycompany.com
oracle.sysman.eml.mntr.emdRepPort=1521
```

**19.2.1.2 About Changing the Repository Password**

For security reasons, the password is encrypted as soon as you start the Management Service. To change the repository password, use the `emctl config oms change_repos_pwd` command line utility. This utility prompts you for the new password for the repository. When you press ENTER after supplying the password, the utility automatically updates the password.

To modify the repository password, do the following:

1. Stop the Management Service using the following command:

```
ORACLE_HOME/bin/emctl stop oms
```

2. Change the repository password by using the following command:

```
emctl config oms -change_repos_pwd [-change_in_db] [-old_pwd <old_pwd>] [-new_pwd <new_pwd>] [-use_sys_pwd [-sys_pwd <sys_pwd>]]
```

-change\_in\_db This option will change in repository too. If not specified, only Credential Store will be updated.

```
emctl config oms -change_view_user_pwd [-sysman_pwd <sysman_pwd>] [-user_pwd <user_pwd>] [-auto_generate]
```

-auto\_generate This option will generate a random password.

- Restart the Management Service using the following command:

```
ORACLE_HOME/bin/emctl start oms
```

## 19.2.2 Configuring the Management Service to Prompt You When Using Execute Commands

The Execute Host Command and Execute SQL applications enable you to execute commands against multiple hosts and multiple databases respectively.

The default, when you click the Execute button of these applications, is for the command execution to begin immediately on the specified targets. If desired, you can set up the Management Service so that a confirmation page displays when you click the Execute button.

To enable the confirmation page for each application, perform the following:

- Stop the Management Service using the following command:

```
ORACLE_HOME/bin/emctl stop oms
```

- Change the properties as follows:

- For the Execute Host Command:

```
emctl set property
oracle.sysman.cmd.tgt.multiTarget.confirmExecuteHostCommand=true
```

- For Execute SQL:

```
emctl set property oracle.sysman.cmd.tgt.multiTarget.confirmExecuteSQL=true
```

---

**Note:** The text in the commands is case-sensitive.

---

- Restart the Management Service using the following command:

```
ORACLE_HOME/bin/emctl start oms
```

## 19.2.3 Troubleshooting Management Service Time Zone Problems

[Section 19.1.6.3](#) describes how to correct potential problems that result when the Management Agent cannot determine the proper time zone. Similar problems can occur when the Management Agent finds the correct time zone, but the time zone is not recognized by the Management Service or the database where the Management Repository resides.

When the Management Service does not recognize the time zone established by the Management Agent, Enterprise Manager generates the following error:

```
OMS does not understand the timezone region of the agent.
Either start the OMS using the extended list of time zones supported by
the database or pick a value of time zone from
ORACLE_HOME/emdw/sysman/admin/nsupportedtzs.lst, update the property
'agentTZRegion' in the file
ORACLE_HOME/sysman/config/emd.properties and restart the agent.
A value which is around an offset of -05:00 from GMT should be picked.
```

This error appears in one of the log files shown in [Table 19–2, "Location of Time Zone Error in the Enterprise Manager Log Files"](#), depending upon which Enterprise Manager product you are using.

There are two ways to correct this problem:

- Restart the Management Repository database using the more extensive list of time zones in the `timez1rg.dat` database configuration file, and then start the Management Agent.

**See Also:** "Specifying the Database Time Zone File" in the *Oracle Database Administrator's Guide*

- Specify a new time zone for the Management Agent that the Management Repository database will recognize.

**See Also:** "[Troubleshooting Management Agent Time Zone Problems](#)" on page 19-7 for instructions on changing the time zone assigned to the Management Agent

---

---

# Index

## A

---

accessing  
  compliance management pages, 13-2  
adaptive alert thresholds, purpose of, 1-4  
ADDM, 1-1  
  purpose of, 1-1  
Agent Registration Password, 2-18  
  changing, 2-27  
Agent Upload Problems  
  default notification rule, 3-8  
AGENT\_HOME/network/admin, 2-32  
Agents Unreachable  
  default notification rule, 3-8  
aggregation and purging policies  
  *See* data retention policies  
Alert Details page, picture of, 1-12  
alerts  
  automated responses, 1-6  
  corrective actions, 1-6  
  notification methods, 1-5  
  notifications for, 1-5  
  purpose of, 1-4  
  server-generated, 1-1  
  Warning Alerts page, 1-10  
analyzing  
  job activity, 6-15  
Application Performance Management, 2-53  
Application Server Availability and Critical/Warning States  
  default notification rule, 3-8  
Application Server Control  
  starting and stopping on Windows systems, 7-5  
archive logging  
  for Management Repository database, 10-2  
asynchronous I/O, 9-16  
automated  
  responses to alerts, 1-6  
Availability History Report, picture of, 5-6, 8-2

## B

---

Backup, 9-18  
baseline normalized views, 1-4  
baselines, 9-5  
  adaptive thresholds, 1-4

  for metrics, 1-4  
  normalized views, 1-4  
BEA WebLogic  
  Oracle ecosystem and, 1-2  
Beacons, 9-17  
  introduction, 1-2  
  monitoring Web Applications over HTTPS, 2-53  
benefits of  
  extending Enterprise Manager, 15-1  
  Information Publisher, 8-1  
blackouts  
  command-line interface and, 1-6  
  controlling with emctl, 7-14  
  examples, 7-16  
  functionality of, 1-6  
  retroactive, 1-7  
buffer cache, 9-14

## C

---

capacity  
  predicting, 9-2  
Certificate dialog box  
  Internet Explorer, 2-51, 2-54  
Checkpoint Firewall, Oracle ecosystem and, 1-2  
Command Line Interface (EMCLI)  
  blackouts and, 1-6  
Common Configurations  
  overview, 17-1  
common configurations  
  deploying a remote management repository, 18-1  
  deploying Grid Control on a single host, 17-2  
  firewalls and other security considerations, 17-1  
  high availability configurations, 18-5  
  managing multiple hosts, 18-1  
  using multiple Management Services, 18-3  
Compare Monitoring Template feature, 1-7  
Compare Targets, picture of, 1-12  
compliance  
  management pages, accessing, 13-2  
  violations, investigating, 13-2  
compliance evaluation  
  scheduling, 13-4  
  setting up, 13-4  
configuring  
  blackouts, functionality of, 1-6

- monitoring templates, 1-7
- Configuring Services, 14-1
  - Availability, 14-4, 14-6
    - Beacons, 14-5
      - Key Beacons, 14-5
      - Local Beacon, 14-5
      - Service Test-Based, 14-5, 14-6
      - System-Based, 14-5, 14-6
  - Command Line Interface, 14-46
  - Create, 14-4, 14-31
  - End-User Performance Monitoring, 14-2, 14-16
    - Access Log Format, 14-20
    - chronos\_setup.sh, 14-22, 14-24
    - EUM, 14-18
    - Manage Web Server Data Collection, 14-19, 14-21
    - Set URLs, 14-19
    - Standalone Web Cache, 14-25
    - Unprocessed Samples, 14-29
    - URL Pattern, 14-20
    - Web Cache Log Format, 14-21
    - Web Cache Manager, 14-21
  - Interactive Transaction Tracing, 14-2
  - J2EE Server Activity, 14-2
  - Metrics
    - Performance, 14-5
    - Usage, 14-8
    - Usage Metrics, 14-5
  - Monitoring Settings, 14-15
    - Beacon Overrides, 14-15
    - Collection Settings, 14-15
    - Data Granularity, 14-15
    - Frequency, 14-15
  - Monitoring Templates, 14-43
    - Beacons, 14-44
    - Service Tests, 14-44
    - Service Tests and Beacons, 14-44
    - Variables, 14-44
  - Oracle Application Server 10g (9.0.4), 14-3
  - Performance, 14-5
  - Performance Metrics, 14-7
    - Aggregation Function, 14-7
  - Recording Transactions, 14-2, 14-14, 14-32
  - Request Performance, 14-3, 14-40
    - Correlate Requests, 14-3
    - Database Connections, 14-41
    - Enterprise Java Beans (EJBs), 14-41
    - JDeveloper, 14-43
    - Manage OC4J Data Collection, 14-41
    - OC4J Cluster, 14-40
    - OC4J Instances, 14-40
    - OC4J Tracing, 14-42
    - Oracle Application Server 10g (9.0.4), 14-3
    - Oracle User Interface XML (UIX), 14-43
    - Service Tests and Beacons, 14-41
    - Tracing Properties, 14-41
  - Root Cause Analysis, 14-3, 14-7, 14-13
    - Component Tests, 14-13
    - Topology, 14-13
    - Topology Viewer, 14-3
  - Service Level Rules, 14-44
    - Actual Service Level, 14-45
    - Availability, 14-45
    - Business Hours, 14-44
    - Expected Service Level, 14-45
    - Information Publisher, 14-46
    - Performance Criteria, 14-45
    - Services Dashboard, 14-46
  - Service Tests and Beacons, 14-9
    - Configuring Dedicated Beacons, 14-10
    - SSL Certificate, 14-10
    - Tests, 14-9
      - Web Proxy, 14-11
  - Service-Test Based Availability
    - Key Service Tests, 14-6
  - System, 14-3
    - Key Components, 14-4
  - System-Based Availability
    - Key Components, 14-6
  - Test Performance, 14-2
  - Thresholds
    - Critical, 14-5
    - Warning, 14-5
  - Time Zone, 14-4
  - Types
    - Aggregate Service, 14-15
  - Types of Service
    - Generic Service, 14-4
  - Types of Services
    - Aggregate Service, 14-4
    - OCS Service, 14-4
    - Web Application, 14-4
  - connect descriptor
    - using to identify the Management Repository database, 10-9, 10-10
  - corrective actions
    - alerts and, 1-6
    - policies, 13-5
    - privileges required for, 1-6
  - creating
    - custom reports, 8-4
    - report definitions, 8-2
    - user-defined metric, 1-8
  - creating a monitoring script, 4-2
  - Critical URL Monitoring, as substitute for
    - Management Agent, 1-2
  - custom reports, 8-4
  - customizing
    - notifications, 1-5
    - policies, 13-4

## D

- dashboard
  - groups, 5-5
- Data Guard
  - configuring Enterprise Manager availability, 10-1
  - data retention policies
    - for Application Performance Management data, 10-3

- for other Management data, 10-3
- modifying default, 10-3
- of the Management Repository, 10-2
- when targets are deleted, 10-4
- Database Availability and Critical/Warning States
  - default notification rule, 3-9
- Database Control
  - starting on UNIX, 7-6
  - stopping on UNIX, 7-6
- DBSNMP database user, 7-10
  - setting the password for, 7-10
- deleting targets
  - data retention policies when, 10-4
- diagnosing
  - policy violations, 13-2
- Disaster Recovery, 9-21
- disk mirroring and stripping
  - Management Repository guideline, 10-1
- disk space management
  - controlling Management Agent disk space, 19-3
- documentation, getting from OTN, 1-8
- dropping the Management Repository, 10-8

## E

- E2E monitoring, 9-17
- em\_message, 4-4
- em\_result, 4-3
- E-mail Customization, 3-14
- e-mail notifications
  - upper limits, 3-5
- e-mails, formats of, 1-5
- emctl
  - controlling blackouts, 7-14
  - listing targets on a managed host, 7-11
  - security commands, 2-18
  - setting monitoring credentials, 7-11
  - starting, stopping, and checking the Management Service, 7-4
- emctl config agent credentials, 7-11
- emctl config agent listtargets, 7-12
- emctl config oms
  - sample output, 2-40
- emctl istop, 7-3
- emctl reload, 7-10
- emctl secure agent, 2-23
  - sample output, 2-24
- emctl secure oms, 2-18
  - sample output, 2-19
- emctl secure setpwd, 2-28
- emctl secure unlock, 2-27
- emctl start agent, 7-2
- emctl start blackout, 7-15
- emctl start dbconsole, 7-6
- emctl start oms, 7-4
- emctl status agent, 7-2
- emctl status blackout, 7-16
- emctl status oms, 7-5
- emctl stop agent, 7-2
- emctl stop blackout, 7-15

- emctl stop dbconsole, 7-6
- emctl stop oms, 7-5
- emctl upload, 7-10
- EMD\_URL
  - property in the emd.properties file, 19-3
- emd.properties, 19-2, 19-3, 19-4
  - EMD\_URL, 19-3
  - emdWalletSrcUrl, 19-2
  - REPOSITORY\_URL, 17-3, 18-2, 19-2
  - UploadMaxBytesXML, 19-4
  - UploadMaxDiskUsedPct, 19-4
- emdRepPort
  - property in the emoms.properties file, 19-10
- emdRepPwd
  - property in the emoms.properties file, 19-10
- emdRepServer
  - property in the emoms.properties file, 19-10
- emdRepSID
  - property in the emoms.properties file, 19-10
- emdRepUser
  - property in the emoms.properties file, 19-10
- emdWalletSrcUrl
  - property in emd.properties, 19-2
- em.notification.emails\_per\_minute
  - property in emoms.properties, 3-4
- em.notification.os\_cmd\_timeout
  - property in emoms.properties, 3-20
- emoms.properties, 10-7
  - emdRepPort, 19-10
  - emdRepPwd, 19-10
  - emdRepServer, 19-10
  - emdRepSID, 19-10
  - emdRepUser, 19-10
  - em.notification.emails\_per\_connection, 3-3
    - property in emoms.properties, 3-3
  - em.notification.emails\_per\_minute, 3-4
  - em.notification.os\_cmd\_timeout, 3-20
  - oracle.net.crypto\_checksum\_client, 2-31
  - oracle.net.crypto\_checksum\_types\_client, 2-31
  - oracle.net.encrypted\_client, 2-31
  - oracle.net.encrypted\_types\_client, 2-31
  - oracle.sysman.eml.mntr.emdRepPwd, 10-7
  - oracle.sysman.eml.mntr.emdRepPwdEncrypted, 10-7
  - oracle.sysman.emRep.dbConn.enableEncryption, 2-31
  - oracle.sysman.emSDK.sec.DirectoryAuthentication Type, 2-10
- emwd watchdog script
  - in the AGENT\_HOME/bin directory, 19-5
- End-User Performance Monitoring
  - Web Server
    - Apache HTTP Server 2.0, 14-2, 14-3, 14-17
    - Oracle Application Server Web Cache, 14-2, 14-3, 14-19
    - Oracle HTTP Server, 14-16, 14-17
    - Oracle HTTP Server Based on Apache 2.0, 14-2, 14-3, 14-16
    - OracleAS Web Cache, 14-16
- Enterprise Manager

- blackouts, functionality of, 1-6
  - monitoring templates, 1-7
- Enterprise Manager Console, picture of, 1-9
- Enterprise Manager Framework Security
  - about, 2-16
  - configuring, 2-16
  - enabling for Management Repository, 2-29
  - enabling for multiple Management Services, 2-25
  - enabling for the Management Agent, 2-23
  - overview of steps required, 2-17
  - restricting HTTP access, 2-26
  - types of secure connections, 2-17
- Extended Network, as substitute for Management Agent, 1-2
- extending Enterprise Manager
  - benefits of, 15-1

## F

---

- figures
  - Availability History Report, 5-7, 8-2
  - Enterprise Manager Console, 1-9
  - Enterprise Manager Grid Control Console, 1-8
  - Group Administration Page, 5-5
  - Group Charts Page, 5-4
  - Group Dashboard, 5-6
  - Group Home Page, 5-3
  - Group Members Page, 5-5
  - Policy Trend Overview Page, 5-3
  - Summary of Target Jobs, 6-16
  - Warning Alert Metric Details, 1-11
  - Warning Alerts Page, 1-10

## G

---

- generating HTML reports, 8-2
- Grid Control
  - architecture overview, 9-1
  - components, 9-1
  - console home page, 1-8
  - deploying on a single host, 17-2
  - sizing, 9-2
  - starting, 7-7
  - starting all components of, 7-7
  - stopping, 7-8
  - stopping all components of, 7-8
- Group Administration page, picture of, 5-4
- Group Charts page, picture of, 5-4
- Group Home page, picture of, 5-2
- Group Members page, picture of, 5-5
- groups
  - administrative tasks, 5-4
  - central monitoring location, 5-2
  - dashboard, 5-5
  - description and purpose, 5-1
  - management features, 5-2
  - member targets, 5-5
  - monitoring collective performance, 5-4
  - policy trends, 5-3
  - redundancy, 5-7

- guidelines
  - for deploying the Management Repository, 10-1

## H

---

- Host Availability and Critical/Warning States
  - default notification rule, 3-9
- HTTP access
  - restricting, 2-26
- HTTP Server Availability and Critical/Warning States
  - default notification rule, 3-9
- HTTPS, 2-17
- Hyper-Threading, 9-12

## I

---

- IBM WebSphere
  - Oracle ecosystem and, 1-2
- Information Publisher
  - Create Like function, 8-4
  - generating HTML reports, 8-2
  - overview of, 8-1
  - predefined report definitions, 8-3
  - predefined reports, 5-6
  - report definitions, 8-2
  - report elements, 8-5
  - reporting framework, 8-1
  - sharing reports, 8-7
  - viewing reports, 8-7
- initialization parameter
  - adjusting when using multiple Management Services, 18-3
- Installing a New Grid Control
  - installation type, 17-2
- Internet Explorer
  - Certificate dialog box, 2-51, 2-54
  - security alert dialog box, 2-50
  - Security Information dialog box, 2-52
- I/O Channels
  - monitoring, 9-16
- istop
  - emctl command, 7-3

## J

---

- javax.net.ssl.SSLException
  - SSL handshake failed, 2-53
- Job Activity page, 6-1
- jobs
  - analyzing job activity, 6-15
  - definition of, 6-1
  - Job Activity page, 6-1
  - job executions, 6-2
  - job runs, 6-2
  - multitask, 6-14
  - notification rules for e-mail, 6-9
  - operations on runs and executions, 6-2
  - predefined tasks, 6-3
  - privileges for sharing job responsibilities, 6-4
  - purpose of, 6-1

## L

---

- Listener Availability
  - default notification rule, 3-10
- load balancer switches
  - BIG-IP, Oracle ecosystem and, 1-2
- Loader, 9-12
- loader threads, 9-13
- LVM (Logical Volume Manager), 10-1

## M

---

- Management Agent, 9-1, 9-4, 9-5
  - additional Management Agent commands, 7-9
  - checking the status on UNIX, 7-2
  - checking the status on Windows, 7-4
  - configuring trust points, 19-9
  - Critical URL Monitoring as substitute, 1-2
  - Extended Network as substitute, 1-2
    - purpose of, 1-2
    - reinstalling, 9-22
  - starting and stopping on UNIX, 7-1
  - starting and stopping on Windows, 7-2
- Management Information Base (MIB), 3-37
  - definition, 3-37
  - MIB variable descriptions, 3-38
- Management Plug-ins
  - extending Enterprise Manager with, 15-1
- Management Repository
  - introduction of, 1-2
  - See* Oracle Management Repository
- Management Server, 9-11
- Management Servers
  - adding, 9-16
- Management Service, 9-1, 9-4
  - starting and stopping on Windows systems, 7-5
  - using a server load balancer, 18-15
- managing
  - compliance, 13-1
  - groups, 5-2
  - policies, 13-1
- master agent
  - Oracle Peer SNMP Master Agent service, 7-3
- metric
  - baselines, 1-4
  - snapshots, 1-3
  - threshold value and alerts, 1-4
  - thresholds, 1-3
- metric baselines
  - adaptive thresholds, 1-4
  - normalized views, 1-4
- metrics
  - creating user-defined, 1-8
  - snapshots, 1-3
  - threshold values, 1-3
  - thresholds, 1-3
  - user-defined, creating, 1-8
  - user-defined, types of, 1-7
  - warning alert details page, 1-11
- MGMT\_ADMIN.DISABLE\_METRIC\_DELETION, 10-5

- MGMT\_ADMIN.ENABLE\_METRIC\_DELETION, 10-5
- MGMT\_METRICS\_1DAY table, 10-4
- MGMT\_METRICS\_1HOUR table, 10-4
- MGMT\_METRICS\_RAW table, 10-4
- MGMT\_PARAMETERS table, 10-3
- MGMT\_RT\_datatype\_1DAY table, 10-4
- MGMT\_RT\_datatype\_1HOUR table, 10-4
- MGMT\_RT\_datatype\_DIST\_1DAY table, 10-4
- MGMT\_RT\_datatype\_DIST\_1HOUR table, 10-4
- MGMT\_RT\_METRICS\_RAW table, 10-4
- MIB

- <italic>See* Management Information Base (MIB)
- monitoring
  - alerts as they occur, 5-5
  - basics of, 1-2
  - information, accessing, 1-8
  - systems, 1-1
  - templates
    - function of, 1-7
- monitoring credentials
  - defined, 7-10
  - example of setting, 7-11
  - setting, 7-10
  - setting in Grid Control, 7-11
  - setting with emctl, 7-11
- monitoring script creation, 4-2
- monitoring template, comparing, 1-7
- monitoring templates, 4-20
- multitask jobs, 6-14

## N

---

- NetApp Filers
  - Oracle ecosystem and, 1-2
- Netscape Navigator
  - New Site Certificate dialog box, 2-51
- network/admin, 2-30, 2-32
- New Site Certificate dialog box
  - Netscape Navigator, 2-51
- Notification Methods, 3-18
- notification methods
  - based on a PL/SQL Procedure, 3-23
  - based on an SNMP trap, 3-26
  - based on operating system commands, 3-19
- notification rules
  - definition, 3-7
  - out-of-box, 3-7
  - out-of-the-box notification rules, 3-6
  - subscribing to, 3-7
- notification schedules, 3-6
- notification system
  - e-mail errors, 3-47
- notification system errors, 3-45
- notification system, trace messages, 3-45
- notifications
  - alerts, 1-5
  - assigning methods to rules, 3-35
  - assigning rules to methods, 3-36
  - customizing, 1-5

- defining multiple mail servers, 3-3
- for jobs, 6-9
- long e-mail notifications, 3-6
- mail server settings, 3-2
- mail server settings in emoms.properties, 3-3
- management information base (MIB), 3-37
- methods, 1-5
- notification method, 1-5
- notification schedules, 3-6
- sample Operating System command script, 3-22
- setting up, 3-1
- short email notifications, 3-6
- using custom notification methods, 3-19

Notification Rules

- Custom, 3-11

## O

---

OC4J Availability and Critical/Warning States

- default notification rule, 3-10

operating system

- user-defined metrics, 1-8

Operating System command

- sample notification method for, 3-20
- sample script, 3-22

Operating System scripts, 3-18

- while creating notification methods, 3-19

ORA-12645

- Parameter does not exist, 2-30

Oracle

- ecosystem, 1-2
- Technology Network (OTN), 1-8

Oracle Advanced Security, 2-17, 2-30

- enabling for Management Repository, 2-32
- enabling for the Management Agent, 2-32

Oracle Application Server Web Cache

- Web Cache Manager, 14-21

Oracle Enterprise Manager

- components, 9-5
- rollup process, 9-14

Oracle Enterprise Manager 10g Grid Control

- See* Grid Control

Oracle Enterprise Manger

- tuning, 9-11

Oracle Internet Directory, 2-6

Oracle Management Agent

- changing the port, 19-2
- controlling disk space used by, 19-3
- enabling security for, 2-23, 2-32
- reconfiguring to use a new Management Service, 19-1
- starting and stopping, 7-1
- Watchdog process, 19-4

Oracle Management Repository, 9-2

- changing the Management Repository password, 19-10
- configuring for high availability, 18-6
- data retention policies, 10-2
- deploying on a remote host, 18-1
- deployment guidelines, 10-1
- dropping, 10-8
- enabling Oracle Advanced Security, 2-32
- enabling security for, 2-29
- identifying with a connect descriptor, 10-9, 10-10
- recreating, 10-8, 10-9
- reloading data, 7-9
- restoring, 9-21
- starting the Management Repository database, 7-8
- troubleshooting, 10-11
- uploading data, 7-9

Oracle Management Service, 9-22

- adjusting the PROCESSES initialization parameter, 18-3
- configuring to use a new Repository, 19-9
- enabling security for, 2-18
- enabling security for multiple Management Services, 2-25
- modifying monitoring credentials, 7-10
- reconfiguring, 19-9
- restoring, 9-22
- starting, stopping, and checking, 7-4
- using multiple management services, 18-3

Oracle Process Management and Notification (OPMN)

- using to start and stop the Management Service, 7-5

Oracle Technology Network (OTN), 14-25

ORACLE\_HOME/network/admin, 2-30, 2-32

oracle.net.crypto\_checksum\_client

- property in emoms.properties, 2-31

oracle.net.crypto\_checksum\_types\_client

- property in emoms.properties, 2-31

oracle.net.encryption\_client

- property in emoms.properties, 2-31

oracle.net.encryption\_types\_client

- property in emoms.properties, 2-31

oracle.sysman.eml.mntr.emdRepPwd

- property in emoms.properties, 10-7

oracle.sysman.eml.mntr.emdRepPwdEncrypted

- property in emoms.properties, 10-7

oracle.sysman.emRep.dbConn.enableEncryption

- entry in emoms.properties, 2-31

oracle.sysman.emSDK.sec.DirectoryAuthenticationType

- property in emoms.properties, 2-10

OS scripts

- <italic>See* Operating System scripts

OTN (Oracle Technology Network), 14-25

OTN and documentation, 1-8

out-of-box

- monitoring, 1-2
- policies, 13-4
- reports, 8-2

## P

---

password

- changing the Management Repository password, 19-10

- changing the SYSMAN password, 10-6
- peer encapsulator service
  - SNMP, 7-3
- Performance Metrics
  - Beacon Aggregation Function
    - Average, 14-7, 14-8
    - Maximum, 14-7
    - Minimum, 14-7, 14-8
    - Sum, 14-7, 14-8
  - System Aggregation Function
    - Maximum, 14-8
- planning outage periods, blackouts, 1-6
- PL/SQL procedures, 3-18
  - sample, 3-26
  - while creating a notification method, 3-23
  - while creating notification methods, 3-19
- policies
  - assessing security policies, 13-3
  - corrective actions, defining, 13-5
  - customizing, 13-4
  - investigating violations, 13-2
  - managing, 13-1
  - out-of-box, 13-4
  - reports on violations, 13-4
  - using templates for monitoring, 13-5
  - violations reports, 13-4
- policy group
  - Oracle database, 13-5
  - Oracle Listener, 13-6
  - RAC, 13-6
- Policy Groups reports, 13-4
- policy management, 13-1
- Policy Trend Overview page, picture of, 5-3
- Policy Violations reports, 13-4
- ports
  - changing the Management Agent port, 19-2
  - default port for the Management Agent upload URL, 17-3
- predefined
  - job tasks, 6-3
- privileges
  - for corrective actions, 1-6
  - for sharing job responsibilities, 6-4
- PROCESSES, 18-3
- ProcessManager
  - service used to control the Management Service on Windows systems, 7-5
- Public Key Infrastructure (PKI), 2-17, 2-53
- purging policies
  - See* data retention policies

## R

---

- RAID-capable disk
  - Management Repository guideline, 10-1
- Recovery, 9-18
- redundancy groups, 5-7
- Repeat Notifications, 3-4
- Repeat Notifications for Rules, 3-4
- RepManager script, 10-8, 10-9

- reports
  - creating custom reports, 8-4
  - custom, 8-4
  - definitions, Information Publisher, 8-2
  - e-mailing, 8-7
  - generating HTML report, 8-2
  - Information Publisher, 8-1
  - out-of-box, Information Publisher, 8-2
  - policy violations, 13-4
  - predefined, 5-6
  - predefined definitions, 8-3
  - predefined report definitions, 8-2
  - report elements, 8-5
  - scheduling, 8-6
  - sharing, 8-7
  - storing and purging, 8-6
  - viewing, 8-7
- Repository Operations Availability
  - default notification rule, 3-10
- REPOSITORY\_URL
  - property in emd.properties, 17-3, 18-2
  - property in the emd.properties file, 19-2
- retroactive blackouts, 1-7
- rollup process, 9-14
- Root Cause Analysis
  - Mode
    - Automatic, 14-13
    - Manual, 14-13
- root password
  - See also* SYSMAN
  - when enabling security for the Management Service, 2-18

## S

---

- scheduled maintenance with blackouts, 1-6
- scheduling
  - reports, 8-6
  - reports, flexibility, 8-6
- script registration, UDM, 4-5
- script results, returning, 4-2
- security
  - about Enterprise Manager security, 2-1
  - overview of steps required to enable Enterprise Manager Framework Security, 2-17
  - policies and, 13-3
  - See also* Enterprise Manager Framework Security
- security alert dialog box
  - Internet Explorer, 2-50
- Security At a Glance feature, policies and, 13-3
- security certificate alerts
  - responding to, 2-50
- security features
  - See* Enterprise Manager Framework Security
- Security Information dialog box
  - Internet Explorer, 2-52
- self-monitoring
  - feature of the Management Agent, 19-5
- Server Connection Hung
  - error while creating the repository, 10-11

- Server Load Balancer, 2-29
- server load balancer, 18-16
  - using with Management Services, 18-15
- server-generated alerts, 1-1
- Service Tests and Beacons
  - Tests
    - DNS, 14-9
    - FTP, 14-9
    - SOAP, 14-9
    - Web Transaction, 14-9
- Services control panel
  - using to start and stop the Management Agent, 7-6
  - using to start the Management Service, 7-5
- setting
  - metric threshold values, 1-3
- sharing reports, 8-7
- snapshots of metrics, 1-3
- SNMP
  - Oracle Peer SNMP Master Agent service, 7-3
  - Oracle SNMP Peer Encapsulator service, 7-3
- SNMP traps, 3-18, 3-19, 3-26
  - sample, 3-27
- SQL
  - user-defined metrics, 1-8
- SQL UDM, character limit, 4-15
- SQL UDM, long statements, 4-15
- SQLNET.CRYPTO\_SEED
  - entry in sqlnet.ora, 2-32
- SQLNET.ENCRYPTION\_SERVER
  - entry in sqlnet.ora, 2-32
- sqlnet.ora, 2-30
  - SQLNET.CRYPTO\_SEED, 2-32
  - SQLNET.ENCRYPTION\_SERVER, 2-32
- state directory
  - in the Management Agent home, 19-2
- Status Codes, Corrective Actions, 3-32, 3-34
- support of third-party components, 1-1
- SYSMAN
  - changing the SYSMAN password, 10-6
  - checking for existence of, 10-11
  - entering SYSMAN password when enabling security, 2-18
- System Dashboard, picture of, 5-6
- system errors, notification, 3-45
- systems
  - monitoring, 1-1

## T

---

- target
  - definition of, 1-2
- target monitoring credentials
  - defined, 7-10
  - example of setting, 7-11
  - setting, 7-10
  - setting in Grid Control, 7-11
- targets
  - listing targets on a managed host, 7-11
- templates, policies and, 13-5

- third-party
  - components, support of, 1-1
- thresholds, 9-5, 9-9
  - adaptive, 1-4
  - adaptive alert, 1-4
  - alerts and, 1-4
  - definition of, 1-3
  - for metrics, 1-3
- troubleshooting
  - general techniques while creating the Management Repository, 10-11
  - while creating the Management Repository, 10-10
- Troubleshooting Service Tests, 14-46
  - Forms Transactions, 14-47
- troubleshooting, notifications, 3-45
- trust points
  - Management Agent Configuration, 19-9

## U

---

- upload directory
  - in the Management Agent home, 19-2, 19-3
- UploadMaxBytesXML
  - property in the emd.properties file, 19-4
- UploadMaxDiskUsedPct
  - property in the emd.properties file, 19-4
- Usage Metrics
  - Aggregation Function
    - Average, 14-8, 14-9
    - Maximum, 14-8, 14-9
    - Minimum, 14-8, 14-9
    - Sum, 14-8, 14-9
- user-defined
  - metrics, 1-7
  - metrics, creating, 1-8
- User-defined metric page
  - Command Line, 4-6
  - Comparison Operator, 4-7, 4-12
  - Consecutive Occurrences Preceding Notification, 4-8, 4-13
  - Critical, 4-8, 4-13
  - Environment, 4-7
  - Metric Name, 4-6, 4-11
  - Metric Type, 4-6, 4-11
  - Operating System User Name and Password, 4-7
  - Response Action, 4-8, 4-14
  - Warning, 4-7, 4-13
- User-defined metric page, Central Console, 4-5, 4-10
- user-defined metric, example, 4-8
- User-defined metrics, 4-1
- user-defined metrics, 4-20

## V

---

- viewing
  - reports, 8-7

## W

---

- Warning Alert Metric Details page, picture of, 1-11
- Warning Alerts page, picture of, 1-10

- watchdog process
  - for the Management Agent, 19-4
- Web Application
  - Source
    - Step, 14-7
    - Step Group, 14-7
    - Transaction, 14-7
- Web Applications
  - monitoring over HTTPS, 2-53
- Web Cache Availability and Critical/Warning States
  - default notification rule, 3-11

