

Oracle® Enterprise Manager

Advanced Configuration

10g Release 2 (10.2)

B16242-01

October 2005

Oracle Enterprise Manager Advanced Configuration, 10g Release 2 (10.2)

B16242-01

Copyright © 2003, 2005, Oracle. All rights reserved.

Contributor: Raj Aggarwal, Muralidharan Bhoopathy, Diarmuid Cawley, Phil Choi, Leo Cloutier, Sudip Datta, Erik DeMember, Kondayya Duvvuri, James Emmond, Irina Goldshteyn, Jacqueline Gosselin, Scott Grover, Rahul Gupta, Luming Han, Ana Hernandez, Narain Jagathesan, Eunhei Jang, Aparna Kamath, Ramanujam Krishnan, Dennis A. Lee, Conrad Lo, Jaydeep Marfatia, Karen McKeen, Rahul Pandey, Raghu Patti, Ravi Pinnamaneni, Pushpa Raghavachar, Sridhar T. Reddy, Prashanth Shishir, Anu Vale, Steven Viavant, James Viscusi, Jin G. Wang, Julie Wong, Michael Zampiceni

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States 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, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software—Restricted Rights (June 1987). Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle, JD Edwards, PeopleSoft, and Retek are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Contents

Preface	xv
Intended Audience	xv
Documentation Accessibility	xv
Related Documents	xvi
Conventions	xvi
1 Introduction to Enterprise Manager Advanced Configuration	
1.1 Types of Advanced Configuration Tasks	1-1
1.2 Understanding the Enterprise Manager Directory Structure	1-1
1.2.1 Understanding the Enterprise Manager Directories Installed with Oracle Enterprise Manager 10g Grid Control	1-2
1.2.1.1 About the Oracle Management Service Home Directory	1-2
1.2.1.2 About the Oracle Management Agent Home (AGENT_HOME) Directory	1-3
1.2.1.3 Summary of the Important Directories in the Management Service Home	1-3
1.2.2 Understanding the Enterprise Manager Directories Installed with the Management Agent	1-4
1.2.2.1 Summary of the Important Directories in the Management Agent Home	1-5
1.2.2.2 Understanding the Management Agent Directory Structure on Windows	1-5
1.2.3 Understanding the Enterprise Manager Directories Installed with Oracle Application Server	1-6
1.2.4 Understanding the Enterprise Manager Directories Installed with Oracle Database 10g	1-6
1.2.5 Tip for Identifying the Oracle Home When Using the emctl Command	1-7
1.2.6 Configuring Database Control During and After the Oracle Database 10g Installation	1-8
1.2.6.1 Configuring Database Control During Installation	1-8
1.2.6.2 Configuring Database Control with DBCA	1-9
1.2.6.3 Configuring Database Control with EMCA	1-10
1.2.6.4 Using an Input File for EMCA Parameters	1-14
1.2.6.5 Using EMCA with Real Application Clusters	1-15
1.2.6.6 Specifying the Ports Used By the Database Control	1-16
1.2.6.7 EMCA Troubleshooting Tips	1-17
1.2.6.7.1 Using EMCA After Changing the Database Listener Port	1-17
1.2.6.7.2 Upgrading Database or ASM Instances with 10g Release 2 Grid Control Agents	1-17
1.3 Enabling Enterprise Manager Accessibility Features	1-18

1.3.1	Enabling Enterprise Manager Accessibility Mode.....	1-18
1.3.2	Providing Textual Descriptions of Enterprise Manager Charts	1-19

2 Starting and Stopping Enterprise Manager Components

2.1	Controlling the Oracle Management Agent.....	2-1
2.1.1	Starting, Stopping, and Checking the Status of the Management Agent on UNIX...	2-1
2.1.2	Starting and Stopping the Management Agent on Windows.....	2-2
2.1.3	Checking the Status of the Management Agent on Windows	2-3
2.2	Controlling the Oracle Management Service.....	2-4
2.2.1	Controlling the Management Service on UNIX	2-4
2.2.1.1	Using OPMN to Start and Stop the Management Service.....	2-4
2.2.1.2	Using emctl to Start, Stop, and Check the Status of the Oracle Management Service	2-4
2.2.1.3	Starting and Stopping Oracle Application Server Web Cache	2-5
2.2.2	Controlling the Management Service on Windows.....	2-6
2.3	Controlling the Application Server Control.....	2-7
2.3.1	Starting and Stopping the Application Server Control on UNIX.....	2-7
2.3.2	Starting and Stopping the Application Server Control on Windows	2-7
2.4	Controlling the Database Control on UNIX.....	2-8
2.4.1	Starting the Database Control on UNIX.....	2-8
2.4.2	Stopping the Database Control on UNIX.....	2-8
2.4.3	Starting and Stopping the Database Control on Windows	2-9
2.5	Guidelines for Starting Multiple Enterprise Manager Components on a Single Host	2-9
2.6	Starting and Stopping Oracle Enterprise Manager 10g Grid Control	2-10
2.6.1	Starting Grid Control and All Its Components	2-10
2.6.2	Stopping Grid Control and All Its Components	2-11
2.7	Additional Management Agent Commands	2-12
2.7.1	Uploading and Reloading Data to the Management Repository	2-13
2.7.2	Specifying New Target Monitoring Credentials.....	2-13
2.7.2.1	Using the Grid Control Console to Modify the Monitoring Credentials	2-14
2.7.2.2	Using the Enterprise Manager Command Line to Modify the Monitoring Credentials	2-14
2.7.3	Listing the Targets on a Managed Host.....	2-15
2.7.4	Controlling Blackouts.....	2-15
2.7.5	Changing the Management Agent Time Zone.....	2-17
2.7.6	Reevaluating Metric Collections.....	2-18

3 Grid Control Common Configurations

3.1	About Common Configurations.....	3-1
3.2	Summary of the Grid Control Architecture and Components.....	3-2
3.3	Deploying Grid Control Components on a Single Host	3-2
3.4	Managing Multiple Hosts and Deploying a Remote Management Repository	3-4
3.5	Using Multiple Management Service Installations.....	3-6
3.5.1	Determining When to Use Multiple Management Service Installations	3-6
3.5.1.1	Monitoring the Load on Your Management Service Installations	3-6
3.5.1.2	Monitoring the Response Time of the Enterprise Manager Web Application Target.....	3-7

3.5.2	Understanding the Flow of Management Data When Using Multiple Management Services.....	3-8
3.6	High Availability Configurations.....	3-9
3.6.1	Load Balancing Connections Between the Management Agent and the Management Service	3-10
3.6.1.1	Configuring the Management Services for High Availability	3-10
3.6.1.2	Understanding the Flow of Data When Load Balancing the Upload of Management Data.....	3-11
3.6.1.3	Configuring a Server Load Balancer for Management Agent Data Upload.....	3-12
3.6.2	Load Balancing Connections Between the Grid Control Console and the Management Service	3-13
3.6.2.1	Understanding the Flow of Data When Load Balancing the Grid Control Console.....	3-13
3.6.2.2	Configuring a Server Load Balancer for the Grid Control Console.....	3-14
3.6.2.3	Configuring Oracle HTTP Server When Using a Server Load Balancer for the Grid Control Console	3-15
3.6.3	Configuring the Management Repository for High Availability	3-15
3.6.3.1	Understanding the Flow of Data When Configuring the Management Repository for High Availability.....	3-16
3.6.3.2	Installing the Management Repository into a Real Applications Clusters (RAC) Database.....	3-17
3.6.3.3	Specifying the Size of the Management Repository Tablespaces in a RAC Database.....	3-18
3.6.3.4	Configuring the Management Service to Use Oracle Net Load Balancing and Failover.....	3-18

4 Enterprise Manager Security

4.1	About Oracle Enterprise Manager Security	4-1
4.1.1	Oracle Enterprise Manager Security Model.....	4-1
4.1.2	Classes of Users and Their Privileges	4-2
4.1.3	Resources Protected.....	4-2
4.1.4	Authorization and Access Enforcement.....	4-2
4.1.5	Leveraging Oracle Application Server Security Services	4-3
4.1.6	Leveraging Oracle Identity Management Infrastructure.....	4-3
4.2	Configuring Security for Grid Control	4-4
4.2.1	About Enterprise Manager Framework Security	4-4
4.2.2	Overview of the Steps Required to Enable Enterprise Manager Framework Security.....	4-5
4.2.3	Enabling Security for the Oracle Management Service.....	4-6
4.2.3.1	Checking the Security Status	4-8
4.2.4	Enabling Security for the Oracle Management Agent	4-9
4.2.5	Enabling Security with Multiple Management Service Installations.....	4-11
4.2.6	Restricting HTTP Access to the Management Service	4-11
4.2.7	Managing Agent Registration Passwords.....	4-13
4.2.7.1	Using the Grid Control Console to Manage Agent Registration Passwords....	4-13
4.2.7.2	Using emctl to Change the Agent Registration Password	4-14
4.2.8	Enabling Security with a Server Load Balancer	4-14
4.2.9	Enabling Security for the Management Repository Database	4-15

4.2.9.1	About Oracle Advanced Security and the sqlnet.ora Configuration File	4-15
4.2.9.2	Configuring the Management Service to Connect to a Secure Management Repository Database	4-16
4.2.9.3	Enabling Oracle Advanced Security for the Management Repository.....	4-17
4.2.9.4	Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database	4-18
4.3	Configuring Enterprise Manager for Use with Oracle Application Server Single Sign-On.....	4-18
4.3.1	Configuring Enterprise Manager to Use the Single Sign-On Logon Page	4-19
4.3.2	Registering Single Sign-On Users as Enterprise Manager Administrators.....	4-21
4.3.3	Grid Control as a Single Sign-On Partner Application	4-22
4.3.4	Bypassing the Single Sign-On Logon Page	4-22
4.4	Configuring Enterprise Manager for Use with Enterprise User Security	4-23
4.5	Setting Up the Auditing System for Enterprise Manager.....	4-23
4.5.1	Audit Data	4-24
4.5.2	Operation Codes	4-25
4.5.3	Audit APIs	4-25
4.5.4	Configuring the Enterprise Manager Audit System.....	4-26
4.6	Configuring the emkey	4-27
4.6.1	Generating the emkey	4-27
4.6.2	emctl Commands	4-28
4.6.2.1	emctl status emkey	4-28
4.6.2.2	emctl config emkey -repos	4-29
4.6.2.3	emctl config emkey -emkeyfile	4-29
4.6.2.4	emctl config emkey -emkey	4-30
4.6.2.5	emctl config emkey -remove_from_repos	4-30
4.6.2.6	emctl config emkey -copy_to_repos	4-30
4.6.3	Install and Upgrade Scenarios	4-31
4.6.3.1	Installing the Management Repository	4-31
4.6.3.2	Installing the First Oracle Management Service	4-31
4.6.3.3	Installing Additional Oracle Management Service	4-31
4.6.3.4	Upgrading from 10.1 to 10.2	4-31
4.6.3.5	Recreating the Management Repository	4-32
4.7	Additional Security Considerations.....	4-32
4.7.1	Responding to Browser-Specific Security Certificate Alerts	4-32
4.7.1.1	Responding to the Internet Explorer Security Alert Dialog Box	4-32
4.7.1.2	Responding to the Netscape Navigator New Site Certificate Dialog Box	4-34
4.7.1.3	Preventing the Display of the Internet Explorer Security Information Dialog Box.....	4-35
4.7.2	Configuring Beacons to Monitor Web Applications Over HTTPS.....	4-35
4.8	Other Security Features.....	4-37
4.8.1	Using ORACLE_HOME Credentials	4-37
4.8.2	Patching Oracle Homes When the User is Locked	4-39
4.8.3	Cloning Oracle Homes.....	4-39
4.8.4	Using the sudo Command.....	4-40

5 Configuring Enterprise Manager for Firewalls

5.1	Considerations Before Configuring Your Firewall	5-1
-----	---	-----

5.2	Firewall Configurations for Enterprise Management Components.....	5-2
5.2.1	Firewalls Between Your Browser and the Grid Control Console.....	5-2
5.2.2	Configuring the Management Agent on a Host Protected by a Firewall.....	5-3
5.2.2.1	Configuring the Management Agent to Use a Proxy Server	5-4
5.2.2.2	Configuring the Firewall to Allow Incoming Communication From the Management Service	5-4
5.2.3	Configuring the Management Service on a Host Protected by a Firewall.....	5-5
5.2.3.1	Configuring the Management Service to Use a Proxy Server.....	5-6
5.2.3.2	About the dontProxyfor Property	5-7
5.2.3.3	Configuring the Firewall to Allow Incoming Management Data From the Management Agents.....	5-7
5.2.4	Firewalls Between the Management Service and the Management Repository	5-8
5.2.5	Firewalls Between the Grid Control and a Managed Database Target	5-8
5.2.6	Firewalls Used with Multiple Management Services.....	5-9
5.2.7	Configuring Firewalls to Allow ICMP and UDP Traffic for Beacons.....	5-10
5.2.8	Configuring Firewalls When Managing Oracle Application Server.....	5-10
5.3	Viewing a Summary of the Ports Assigned During the Application Server Installation	5-10
5.4	Checking and Configuring Firewall Settings for HTTP/HTTPS.....	5-11

6 Configuring Services

6.1	Summary of Service Management Tasks	6-1
6.2	Setting up the System.....	6-3
6.3	Creating a Service	6-4
6.4	Configuring a Service	6-5
6.4.1	Availability Definition	6-6
6.4.2	Performance Metrics	6-7
6.4.3	Usage Metrics.....	6-8
6.4.4	Service Tests and Beacons	6-9
6.4.4.1	Configuring the Beacons	6-9
6.4.5	Root Cause Analysis Configuration.....	6-10
6.4.5.1	Getting the Most From Root Cause Analysis.....	6-11
6.5	Recording Transactions.....	6-12
6.6	Monitoring Settings	6-12
6.7	Configuring Aggregate Services.....	6-13
6.8	Configuring End-User Performance Monitoring	6-13
6.8.1	Configuring End-User Performance Monitoring Using Oracle HTTP Server Based on Apache 2.0 or Apache HTTP Server 2.0	6-14
6.8.1.1	Setting up the Third Party Apache Server	6-16
6.8.2	Configuring End-User Performance Monitoring Using Oracle Application Server Web Cache.....	6-16
6.8.2.1	Configuring Oracle Application Server Web Cache 10.1.2	6-17
6.8.2.2	Configuring Oracle Application Server Web Cache 9.0.4	6-18
6.8.2.3	Configuring End-User Performance Monitoring Using Earlier Versions of Oracle Application Server Web Cache.....	6-19
6.8.2.3.1	Using the chronos_setup.pl Configuration Script.....	6-19
6.8.2.3.2	Configuring the Document Root For Each Web Server.....	6-20

6.8.2.3.3	Configuring Oracle Application Server Web Cache for End-User Performance Monitoring	6-20
6.8.2.3.4	Starting End-User Performance Monitoring	6-21
6.8.2.4	Configuring End-User Performance Monitoring Using Standalone Oracle Application Server Web Cache	6-22
6.8.2.4.1	Installing Standalone Oracle Application Server Web Cache	6-22
6.8.2.4.2	Configuring Standalone Oracle Application Server Web Cache	6-23
6.8.2.4.3	Enabling End-User Performance Monitoring for Standalone Oracle Application Server Web Cache	6-23
6.8.3	Starting and Stopping End-User Performance Monitoring.....	6-24
6.8.4	Verifying and Troubleshooting End-User Performance Monitoring.....	6-24
6.8.5	Setting Up the Forms Application for End-User Performance Monitoring.....	6-26
6.8.5.1	Configuring Forms Server for End-User Performance Monitoring	6-26
6.8.5.2	Changing the Logging Format for OracleAS Web Cache.....	6-26
6.9	Configuring OC4J for Request Performance Diagnostics	6-27
6.9.1	Selecting OC4J Targets for Request Performance Diagnostics	6-27
6.9.2	Configuring Interactive Transaction Tracing	6-27
6.9.3	Configuring OC4J Tracing for Request Performance Data	6-28
6.9.4	Additional Configuration for Monitoring UIX Applications.....	6-29
6.10	Setting Up Monitoring Templates	6-30
6.10.1	Configuring Service Tests and Beacons.....	6-30
6.11	Configuring Service Levels.....	6-31
6.11.1	Defining Service Level Rules	6-32
6.11.2	Viewing Service Level Details	6-32
6.12	Configuring a Service Using the Command Line Interface.....	6-33

7 Locating and Configuring Enterprise Manager Log Files

7.1	Locating and Configuring Management Agent Log and Trace Files.....	7-1
7.1.1	About the Management Agent Log and Trace Files.....	7-1
7.1.2	Locating the Management Agent Log and Trace Files.....	7-2
7.1.3	About Management Agent Rollover Files.....	7-2
7.1.4	Controlling the Size and Number of Management Agent Log and Trace Files	7-3
7.1.5	Controlling the Contents of the Management Agent Trace File.....	7-4
7.1.6	Controlling the Size and Number of Fetchlet Log and Trace Files	7-4
7.1.7	Controlling the Contents of the Fetchlet Trace File	7-5
7.2	Locating and Configuring Management Service Log and Trace Files	7-6
7.2.1	About the Management Service Log and Trace Files	7-6
7.2.2	Locating the Management Service Log and Trace Files.....	7-6
7.2.3	Controlling the Size and Number of Management Service Log and Trace Files	7-6
7.2.4	Controlling the Contents of the Management Service Trace File	7-8
7.2.5	Controlling the Oracle Application Server Log Files	7-8

8 Maintaining and Troubleshooting the Management Repository

8.1	Management Repository Deployment Guidelines	8-1
8.2	Management Repository Data Retention Policies.....	8-2
8.2.1	Management Repository Default Aggregation and Purging Policies.....	8-2

8.2.2	Management Repository Default Aggregation and Purging Policies for Other Management Data	8-3
8.2.3	Modifying the Default Aggregation and Purging Policies.....	8-3
8.2.4	Modifying Data Retention Policies When Targets Are Deleted	8-4
8.3	Changing the SYSMAN Password	8-5
8.4	Dropping and Recreating the Management Repository	8-6
8.4.1	Dropping the Management Repository.....	8-6
8.4.2	Recreating the Management Repository	8-7
8.4.2.1	Using the RepManager Script to Create the Management Repository.....	8-7
8.4.2.2	Using a Connect Descriptor to Identify the Management Repository Database	8-8
8.5	Troubleshooting Management Repository Creation Errors	8-9
8.5.1	Package Body Does Not Exist Error While Creating the Management Repository...	8-9
8.5.2	Server Connection Hung Error While Creating the Management Repository.....	8-9
8.5.3	General Troubleshooting Techniques for Creating the Management Repository	8-9
8.6	Improving the Login Performance of the Console Home Page	8-10

9 Sizing and Maximizing the Performance of Oracle Enterprise Manager

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	9-7
9.2.3.1	Online Weekly Tasks.....	9-7
9.2.3.2	Offline Monthly Tasks	9-8
9.2.4	Step 4: Eliminate Bottlenecks Through Tuning.....	9-10
9.2.4.1	High CPU Utilization.....	9-10
9.2.4.2	Loader Vital Signs	9-11
9.2.4.3	Rollup Vital Signs	9-12
9.2.4.4	Job, Notification, and Alert Vital Signs	9-13
9.2.4.5	I/O Vital Signs	9-13
9.2.4.6	The Oracle Enterprise Manager Performance Page.....	9-14
9.2.5	Step 5: Extrapolating Linearly Into the Future for Sizing Requirements	9-15
9.3	Oracle Enterprise Manager Backup, Recovery, and Disaster Recovery Considerations.....	9-15
9.3.1	Best Practices for Backup and Recovery.....	9-16
9.3.1.1	Oracle Management Service	9-16
9.3.1.2	Management Agent.....	9-16
9.3.2	Best Practice for Disaster Recovery (DR)	9-17
9.3.2.1	Management Repository	9-17
9.3.2.2	Oracle Management Service	9-17
9.3.2.3	Management Agent.....	9-17
9.4	Configuring Enterprise Manager for High Availability	9-18
9.4.1	Architectural Overview	9-18
9.4.2	Installation and Configuration for High Availability	9-18
9.4.2.1	Management Agent.....	9-18

9.4.2.1.1	Configuring the Management Agent to Automatically Start on Boot and Restart on Failure.....	9-19
9.4.2.1.2	Configuring Restart for the Management Agent.....	9-19
9.4.2.1.3	Configuring the Connection Between Management Agents and the Management Service	9-19
9.4.2.1.4	Installing the Management Agent Software on Redundant Storage	9-20
9.4.2.1.5	Configuring All Out-of-band Notifications.....	9-20
9.4.2.2	Management Service	9-21
9.4.2.2.1	Configuring the Shared Filesystem Loader.....	9-21
9.4.2.2.2	Configuring SLB to Abstract the Underlying Management Service Host Names for Easier Reconnect After Failure	9-21
9.4.2.2.3	Management Service Installation Should Be Done to Non-Clustered Servers.....	9-21
9.4.2.2.4	Configuring Management Service to Use Client Side Oracle Net Load Balancing for Failover and Load Balancing.....	9-21
9.4.2.2.5	Install the Management Service Software on Redundant Storage.....	9-22
9.4.2.3	Management Repository	9-22
9.4.2.3.1	Install Into an Existing RAC Management Repository	9-22
9.4.2.3.2	Consider (Physical) Data Guard for Redundancy	9-23
9.4.3	Configuration Within Grid Control	9-23
9.4.3.1	Console Warnings, Alerts, and Notifications	9-23
9.4.3.2	Configure Additional Error Reporting Mechanisms	9-23
9.4.3.3	Component Backup.....	9-24
9.4.3.4	Troubleshooting.....	9-24
9.4.3.4.1	Upload Delay for Monitoring Data	9-24
9.4.3.4.2	Notification Delay of Target State Change.....	9-24

10 Reconfiguring the Management Agent and Management Service

10.1	Reconfiguring the Oracle Management Agent.....	10-1
10.1.1	Configuring the Management Agent to Use a New Management Service.....	10-1
10.1.2	Changing the Management Agent Port.....	10-2
10.1.3	Controlling the Amount of Disk Space Used by the Management Agent	10-3
10.1.4	About the Management Agent Watchdog Process.....	10-4
10.1.5	Setting the Management Agent Time Zone	10-4
10.1.5.1	Understanding How the Management Agent Obtains Time Zone Information.....	10-4
10.1.5.2	Resetting the Time Zone of the Management Agent Due to Inconsistency of Time Zones.....	10-5
10.1.5.3	Troubleshooting Management Agent Time Zone Problems.....	10-5
10.1.5.4	Troubleshooting Management Service Time Zone Problems	10-6
10.1.6	Adding Trust Points to the Management Agent Configuration.....	10-7
10.2	Reconfiguring the Oracle Management Service	10-8
10.2.1	Configuring the Management Service to Use a New Management Repository	10-8
10.2.1.1	Changing the Repository Properties in the emoms.properties File	10-8
10.2.1.2	About Changing the Repository Password	10-9
10.2.2	Configuring the Management Service to Use a New Port.....	10-9
10.2.3	Configuring the Management Service to Prompt You When Using Execute Commands.....	10-10

11 Migrating from Previous Versions of Enterprise Manager

11.1	Overview of the Enterprise Manager Migration Process.....	11-1
11.2	Requirements for Migrating from Previous Versions of Enterprise Manager.....	11-1
11.3	The Oracle Enterprise Manager 10g Migration Process.....	11-2
11.3.1	Deploying and Configuring Oracle Enterprise Manager 10g Management Agents.....	11-2
11.3.1.1	Deploying the Oracle Enterprise Manager 10g Management Agents Using the Release 2.2, Release 9.0.1, or Release 9.2 Job System	11-3
11.3.1.1.1	More About the Directory Type Parameter.....	11-4
11.3.1.2	Configuring the Oracle Enterprise Manager 10g Management Agents for Use with the Oracle Enterprise Manager 10g Job System (UNIX Systems Only)....	11-6
11.3.2	Migrating Management Repository Data	11-7
11.4	Configuring Metric Thresholds	11-8
11.4.1	Copying Metric Thresholds to Multiple Targets.....	11-8

12 Configuring Notifications

12.1	Setting Up Notifications.....	12-1
12.1.1	Setting Up a Mail Server for Notifications.....	12-1
12.1.2	Setting Up E-mail for Yourself.....	12-4
12.1.2.1	Defining E-mail Addresses	12-4
12.1.2.2	Setting Up a Notification Schedule.....	12-5
12.1.2.3	Subscribe to Receive E-mail for Notification Rules	12-6
12.1.3	Setting Up E-mail for Other Administrators	12-9
12.2	Extending Notification Beyond E-mail.....	12-10
12.2.1	Custom Notification Methods Using Scripts and SNMP Traps	12-11
12.2.1.1	Adding a Notification Method based on an OS Command or Script.....	12-11
12.2.1.2	Adding a Notification Method Based on a PL/SQL Procedure.....	12-15
12.2.1.3	Adding a Notification Method Based on an SNMP Trap.....	12-19
12.3	Passing Corrective Action Status Change Information.....	12-21
12.3.1	Passing Corrective Action Execution Status to an OS Command or Script.....	12-21
12.3.2	Passing Corrective Action Execution Status to a PLSQL Procedure.....	12-22
12.4	Passing Job Execution Status Information.....	12-24
12.4.1	Passing Job Execution Status to a PLSQL Procedure	12-25
12.4.2	Passing Job Execution Status to an OS Command or Script.....	12-26
12.5	Assigning Methods to Rules.....	12-27
12.6	Assigning Rules to Methods.....	12-28
12.7	Management Information Base (MIB).....	12-29
12.7.1	About MIBs.....	12-29
12.7.2	Reading the MIB Variable Descriptions	12-30
12.7.2.1	Variable Name	12-30
12.7.2.2	MIB Definition.....	12-30
12.8	Troubleshooting Notifications	12-36
12.8.1	General Setup	12-37
12.8.2	Notification System Errors	12-37
12.8.3	Notification System Trace Messages.....	12-37
12.8.4	E-mail Errors.....	12-38
12.8.5	OS Command Errors.....	12-39

12.8.6	SNMP Trap Errors	12-39
12.8.7	PL/SQL Errors	12-39

13 User-Defined Metrics

13.1	Extending Monitoring Capability.....	13-1
13.2	Creating OS-Based User-Defined Metrics	13-2
13.2.1	Create Your OS Monitoring Script	13-2
13.2.1.1	Code to check the status of monitored objects	13-2
13.2.1.2	Code to return script results to Enterprise Manager.....	13-2
13.2.1.3	Script Runtime Environment	13-4
13.2.2	Register the Script as a User-Defined Metric	13-5
13.2.3	OS-Based User-Defined Metric Example	13-8
13.3	Creating a SQL-Based User-Defined Metric	13-10
13.3.1	SQL-Based User-Defined Metric Examples	13-14
13.3.1.1	Example 1: Query Returning Tablespace Name and Percent Used	13-14
13.3.1.2	Example 2: Query Returning Segment Name/Type and Extent Count.....	13-15
13.4	Notifications, Corrective Actions, and Monitoring Templates	13-15
13.4.1	Getting Notifications for User-Defined Metrics.....	13-15
13.4.2	Setting Corrective Actions for User-Defined Metrics.....	13-16
13.4.3	Deploying User-Defined Metrics across many targets using Monitoring Templates.....	13-16

14 Additional Configuration Tasks

14.1	Understanding Default and Custom Data Collections.....	14-1
14.1.1	How Enterprise Manager Stores Default Collection Information	14-1
14.1.2	Restoring Default Collection Settings	14-2
14.2	Enabling Multi-Inventory Support for Configuration Management	14-2
14.2.1	AGENT_HOME Versus AGENT_STATE Directories.....	14-3
14.3	Manually Configuring a Database Target for Complete Monitoring	14-4
14.4	Modifying the Default Login Timeout Value	14-6
14.5	Configuring Clusters and Cluster Databases in Grid Control	14-7
14.5.1	Configuring Clusters.....	14-7
14.5.2	Configuring Cluster Databases.....	14-7
14.5.3	Discovering Instances Added to the Cluster Database	14-8
14.5.3.1	Troubleshooting.....	14-9
14.6	Collecting Client Configurations	14-9
14.6.1	Configuring the Client System Analyzer	14-10
14.6.1.1	Client System Analyzer in Oracle Grid Control	14-10
14.6.1.2	Deploying Client System Analyzer Independently	14-10
14.6.2	Configuration Parameters	14-11
14.6.2.1	Associating the Parameters with an Application	14-15
14.6.3	Rules	14-15
14.6.4	Customization	14-17
14.6.5	CSA Deployment Examples.....	14-17
14.6.5.1	Using Multiple Collection Tags.....	14-17
14.6.5.2	Privilege Model for Viewing Client Configurations	14-18
14.6.5.3	Using the Customization API Example	14-19

14.6.5.4	Using the CSA Servlet Filter Example.....	14-20
14.6.5.5	Sample Deployments	14-21
14.6.5.5.1	Example 1: Helpdesk	14-21
14.6.5.5.2	Example 2: Inventory	14-22
14.6.5.5.3	Example 3: Problem Detection	14-23

Index

Preface

This *Advanced Configuration* guide describes the advanced configuration tasks you can perform after you have installed Oracle Enterprise Manager and have started using the software. These tasks are optional and provide additional functionality for specific types of Oracle Enterprise Manager customers.

Note that later versions of this and other Enterprise Manager books may be available on the Oracle Technology Network:

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

Intended Audience

This guide is written for system administrators who want to configure the advanced features of Oracle Enterprise Manager 10g. You should already be familiar with Oracle and the administrative tasks you want to perform.

You should also be familiar with the operation of your specific UNIX or Windows system. Refer to the documentation for your platform-specific documentation, if necessary.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For additional information, visit the Oracle Accessibility Program Web site at

<http://www.oracle.com/accessibility/>

Accessibility of Code Examples in Documentation JAWS, a Windows screen reader, may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, JAWS may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation This documentation may contain links to Web sites of other companies or organizations that Oracle does

not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Related Documents

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

- *Oracle Enterprise Manager Concepts*
- *Oracle Enterprise Manager Grid Control Quick Installation Guide*
- *Oracle Enterprise Manager Grid Control Installation and Basic Configuration*
- *Oracle Enterprise Manager Configuration for Oracle Collaboration Suite*
- *Oracle Enterprise Manager Policy Reference Manual*
- *Oracle Enterprise Manager Metric Reference Manual*
- *Oracle Enterprise Manager Extensibility*
- *Oracle Enterprise Manager Command Line Interface*

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.

Introduction to Enterprise Manager Advanced Configuration

This chapter introduces you to Enterprise Manager advanced configuration and provides basic information about your Enterprise Manager installation. It describes the directory structure and how to make Enterprise Manager accessible to all your users.

After you review this chapter, you can move on to the other advanced configuration tasks described in this manual.

Specifically, this chapter includes the following topics:

- [Types of Advanced Configuration Tasks](#)
- [Understanding the Enterprise Manager Directory Structure](#)
- [Enabling Enterprise Manager Accessibility Features](#)

1.1 Types of Advanced Configuration Tasks

Enterprise Manager is designed to install easily with a set of standard configuration settings so you can get up and running with the software quickly.

However, Oracle realizes that hardware and software management requirements vary dramatically among business enterprises. As a result, Enterprise Manager can be reconfigured after installation so you can:

- Implement Enterprise Manager security and firewall features.
- Enable End-User Performance Monitoring for your Web applications.
- Reconfigure Enterprise Manager components when you need to modify the topology of your network environment.
- Maintain and troubleshoot the Enterprise Manager components as your business grows.

1.2 Understanding the Enterprise Manager Directory Structure

Before you perform maintenance and advanced configuration tasks, you must be familiar with the directories and files that are copied to disk when you install Enterprise Manager. Understanding where specific files are located can help you if you need to troubleshoot installation or configuration problems.

The directories and files installed by Enterprise Manager vary, depending upon the installation options you select during the Enterprise Manager installation. The location of Enterprise Manager files and directories also varies slightly when Enterprise

Manager is installed as part of an Oracle Application Server or Oracle Database 10g installation.

Use the following sections to become familiar with the directories that are created on your disk when you install Enterprise Manager:

- [Understanding the Enterprise Manager Directories Installed with Oracle Enterprise Manager 10g Grid Control](#)
- [Understanding the Enterprise Manager Directories Installed with the Management Agent](#)
- [Understanding the Enterprise Manager Directories Installed with Oracle Application Server](#)
- [Understanding the Enterprise Manager Directories Installed with Oracle Database 10g](#)
- [Tip for Identifying the Oracle Home When Using the emctl Command](#)
- [Configuring Database Control During and After the Oracle Database 10g Installation](#)

1.2.1 Understanding the Enterprise Manager Directories Installed with Oracle Enterprise Manager 10g Grid Control

When you install Oracle Enterprise Manager 10g Grid Control, you can select from four installation types. All of these installation types, except the Oracle Management Agent installation type, install the Oracle Management Service.

When you install the Oracle Management Service, you actually install three Oracle home directories:

- The Management Service home directory
- The Management Agent home directory
- The Database home directory

Note: When you install Oracle Enterprise Manager 10g Grid Control, Oracle Database is also installed, but will not contain Enterprise Manager Configuration Assistant (EMCA) in the Oracle Database Home.

1.2.1.1 About the Oracle Management Service Home Directory

The Oracle Management Service is a J2EE application that is installed and deployed using Oracle Application Server. As a result, when you install the Oracle Management Service, the installation procedure first installs Oracle Application Server. Specifically, the installation procedure installs the Oracle Application Server J2EE and Web Cache installation type, which is used to deploy the Oracle Management Service.

The installation procedure installs the Enterprise Manager components within the Oracle Application Server Home, including:

- The Oracle Management Service
- Optionally, the Oracle Management Repository

Information about the directories that are specific to the Oracle Application Server installation can be found in the Oracle Application Server documentation. For

example, the location of the most of the Oracle Application Server configuration and log files are described in the Oracle Application Server documentation.

See Also: "Configuration Files and Log Files" in the *Oracle Application Server 10g Administrator's Guide*

1.2.1.2 About the Oracle Management Agent Home (AGENT_HOME) Directory

In addition to the Management Service home directory, the installation procedure installs the Oracle Management Agent that is used to gather management data and perform administration tasks for the targets on the Management Service host.

By default, if the Oracle Universal Installer (or the account used to run the Universal Installer) has the proper privileges to write to the install directories, the Management Agent is installed in a separate Oracle home directory at the same level as the Oracle Application Server home directory.

However, if the Oracle Universal Installer does not have the necessary privileges, the Management Agent is installed in a subdirectory of the Oracle Application Server home directory.

1.2.1.3 Summary of the Important Directories in the Management Service Home

Figure 1–1 shows some of the important directories you should be familiar with in a typical Grid Control Console installation. You can use this information as you begin to maintain, troubleshoot, and configure the Oracle Management Service installation.

Figure 1–1 Important Oracle Management Service Installation Directories

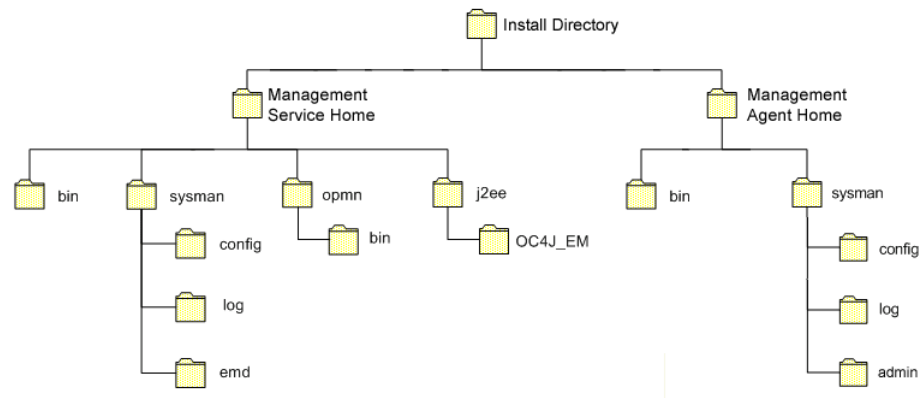


Table 1–1 describes in more detail the Management Service directories shown in Figure 1–1. In the table, ORACLE_HOME refers to the Management Service home directory in which the Oracle Management Service is installed and deployed.

Table 1–1 Important Directories in the Management Service Oracle Home

Directory	Description
ORACLE_HOME/bin	<p>The bin directory in the Oracle Application Server Home contains commands used to control the components of the Oracle Application Server J2EE and Web Cache installation, including the Application Server Control Console, which is used to monitor and configure Oracle Application Server instances.</p> <p>Use the <code>emctl</code> command in this directory to start and stop the Application Server Control Console. For more information about the Application Server Control Console, see the <i>Oracle Application Server 10g Administrator's Guide</i>.</p>
ORACLE_HOME/sysman	<p>The sysman directory in the Oracle Application Server Home contains the system management files associated with this Oracle Application Server Release 2 (9.0.4) installation.</p> <p>Note that the ORACLE_HOME/sysman/log directory contains the Oracle Management Service log files (<code>emoms.log</code>) and trace files (<code>emoms.trc</code>).</p>
ORACLE_HOME/opmn	<p>This directory contains files used to control the Oracle Process Manager and Notification Server (OPMN) utility. OPMN can be used to start and stop the instances of Oracle Application Server Containers for J2EE (OC4J) associated with this instance of Oracle Application Server. The Oracle Management Service runs as an application in one of those OC4J instances.</p>
ORACLE_HOME/j2ee	<p>This directory contains the files associated with the OC4J instances running in this instance of Oracle Application Server. For example, you will notice a directory for the OC4J_EM instance, which is the OC4J instance used to deploy the Management Service J2EE Web application.</p>
ORACLE_HOME/hostname	<p>For real application cluster agent install, this directory contains sysman files.</p>

1.2.2 Understanding the Enterprise Manager Directories Installed with the Management Agent

The Management Agent is installed automatically when you install the Grid Control Console. This local instance of the Management Agent gathers management information about the targets on the Management Service host. You can then manage those targets, such as the host itself, from the Grid Control Console.

The Management Agent is also available as its own install type. This enables you to install the Management Agent on the hosts throughout your enterprise. The Management Agent can then gather management data about the targets on each host so those targets can be managed from the Grid Control Console.

When you select the Additional Management Agent installation type, you install only the files required to run the Management Agent.

Specifically, the Management Agent files are installed into the same directory structure shown in the `agent` directory when you install the Oracle Management Service (Figure 1–1).

The directory that contains the files required to run the Management Agent is referred to as the `AGENT_HOME` directory. For example, to start or stop an Oracle Management Agent, you use the `emctl` command located in the `bin` directory of the

AGENT_HOME. Similarly, to configure the log files for the Management Agent, you modify the configuration files in the `sysman/config` directory of the AGENT_HOME.

1.2.2.1 Summary of the Important Directories in the Management Agent Home

[Table 1–2](#) describes some of the important subdirectories inside the AGENT_HOME directory.

Table 1–2 Important Directories in the AGENT_HOME Directory

Directory	Description
AGENT_HOME	<p>The agent directory contains all the files required to configure and run the Oracle Management Agent on this host.</p> <p>This directory serves as the Oracle Home for the Management Agent. Later in this document, this directory is referred to as the AGENT_HOME.</p> <p>If you install only the Management Agent on a managed host, only the files in this directory are installed. For more information, see "Understanding the Enterprise Manager Directories Installed with the Management Agent" on page 1-4.</p>
AGENT_HOME/bin	<p>The agent/bin directory in the Oracle Application Server Home contains the <code>emctl</code> command that controls the Management Agent for this host.</p> <p>You use the <code>emctl</code> command in this directory to start and stop the Oracle Management Agent on this host.</p>
AGENT_HOME/sysman/admin	<p>This directory contains the files used by the Management Agent to define target types (such as databases, hosts, and so on), to run configuration scripts, and other administrative tasks.</p>
AGENT_HOME/sysman/config	<p>This directory contains the configuration files for the Management Agent. For example, this is where Enterprise Manager stores the <code>emd.properties</code> file. The <code>emd.properties</code> file defines settings such as the Management Service upload URL for this particular agent.</p>
AGENT_HOME/sysman/log	<p>This directory contains the log files for the Management Agent.</p>
AGENT_HOME/hostname	<p>For real application clusters, this directory contains all configuration, log files, and system files.</p>

1.2.2.2 Understanding the Management Agent Directory Structure on Windows

When you install the Management Agent on a Windows system, the directory structure of the AGENT_HOME directory is the same as the directory structure for installations on a UNIX system.

For example, if you installed the Management Agent in the `E:\oracle\em10gAgent` directory of your Windows system, you can locate the `emctl` command for the Management Agent on a Windows system, by navigating to the following directory:

```
$PROMPT> E:\oracle\em10gAgent\bin
```

1.2.3 Understanding the Enterprise Manager Directories Installed with Oracle Application Server

When you install Oracle Application Server (Oracle Application Server), you also install the Oracle Enterprise Manager 10g Application Server Control Console. The Application Server Control Console provides you with the Enterprise Manager features required to manage your Oracle Application Server installation. As a result, the Oracle Application Server installation procedure installs a set of Enterprise Manager directories and files into each Oracle Application Server home directory.

In particular, the `emctl` commands required to control the Application Server Control Console are installed into the `ORACLE_HOME/bin` directory. The configuration and log files for the Application Server Control Console are installed into the `ORACLE_HOME/sysman` directory structure.

See Also: ["Starting and Stopping Oracle Enterprise Manager 10g Grid Control"](#) on page 2-10

["Locating and Configuring Enterprise Manager Log Files"](#) on page 7-1

1.2.4 Understanding the Enterprise Manager Directories Installed with Oracle Database 10g

When you install Oracle Database 10g, you also install Oracle Enterprise Manager 10g Database Control. Database Control provides the tools you need to manage your Oracle Database 10g immediately after you install the database. As a result, the Oracle Database 10g installation procedure installs a set of Enterprise Manager directories and files into each Oracle Database 10g home directory.

In particular, the `emctl` commands required to control Database Control are installed into the `ORACLE_HOME/bin` directory.

The Management Agent and Management Service support files are installed in two locations in an Oracle Database 10g installation:

- Files that are common and shared among all instances of the database are stored in the following directory of the Oracle Database 10g home:

`ORACLE_HOME/sysman`

For example, the administration files, which define the supported target types and the scripts used to perform Management Agent configuration tasks are stored in the `ORACLE_HOME/sysman/admin` directory.

- Files that are unique to each instance of the database are stored in following directory of the Oracle Database 10g home:

`ORACLE_HOME/hostname_sid/` (for a single instance database)

`ORACLE_HOME/nodename_sid/` (for a cluster database)

Throughout the rest of this guide, `ORACLE_HOME/hostname_sid/` and `ORACLE_HOME/nodename_sid/` may be used interchangeably. Both paths refer to the same concept – the Enterprise Manager directory for the specific database instance. The difference is that `ORACLE_HOME/hostname_sid/` is used for single instance databases, while `ORACLE_HOME/nodename_sid/` is used for cluster (RAC) databases. In cluster databases, `nodename` refers to the public name of the node, as specified during Cluster Ready Services (CRS) configuration for cluster environments.

For example, if the database host name is `mgmt1.acme.com` and the system identifier for the database instance is `db42`, the log files for the Management Agent and Management Service for that instance are installed in the following directory:

```
ORACLE_HOME/mgmt1.acme.com_db42/sysman/log
```

See Also: ["Locating and Configuring Enterprise Manager Log Files"](#) on page 7-1


If a `hostname_sid` directory does not exist in the Oracle Database 10g home directory, then Oracle Enterprise Manager 10g Database Control was never configured for the database instance.

See Also: ["Configuring Database Control During and After the Oracle Database 10g Installation"](#) on page 1-8

In addition, the files required to deploy the Database Control as a J2EE application are installed into the `ORACLE_HOME/oc4j/j2ee` directory structure. Database Control is a J2EE application that is deployed using the standalone version of Oracle Application Server Containers for J2EE (OC4J). The `OC4J_DBConsole` directory contains the template files that are used to create database-specific deployment directories for each Database Control instance deployed in the Oracle home.

[Figure 1-2](#) summarizes the location of the important Enterprise Manager directories in a typical Oracle Database 10g home directory. Note that references to `hostname_sid` are for single instance databases; cluster databases have paths of the form `nodename_sid` instead.

Figure 1-2 Important Enterprise Manager Directories in an Oracle Database 10g Installation



1.2.5 Tip for Identifying the Oracle Home When Using the `emctl` Command

When you install Grid Control, Oracle Application Server, or Oracle Database 10g, the resulting directory structure can often include multiple subdirectories with the same name. For example, you can have a `bin` directory within the `AGENT_HOME` directory. Use the `emctl` command within the `AGENT_HOME/bin` directory to control the Management Agent.

In addition, you can have a `bin` directory within the Management Service Oracle home. Use the `emctl` command in this directory to control the Management Service.

To quickly identify the Oracle home that is controlled by the files in a particular `bin` directory, use the following command:

```
$PROMPT> emctl getemhome
```

This command displays the path to the current Oracle home that will be affected by commands executed by this instance of the `emctl` command. For example, the following example shows how the current `emctl` command can be used to control the Management Service installed in the `/dev1/private/em_ms_home1/` Oracle home:

```
$PROMPT> emctl getemhome
Copyright (c) 1996, 2004 Oracle Corporation. All rights reserved.
EMHOME=/dev1/private/em_ms_home1
```

1.2.6 Configuring Database Control During and After the Oracle Database 10g Installation

The following sections describe how Oracle Enterprise Manager 10g Database Control is configured during the Oracle Database 10g installation. These sections also describe how you can configure Database Control after the installation:


- [Configuring Database Control During Installation](#)
- [Configuring Database Control with DBCA](#)
- [Configuring Database Control with EMCA](#)
- [Using EMCA with Real Application Clusters](#)
- [EMCA Troubleshooting Tips](#)

1.2.6.1 Configuring Database Control During Installation

If you create a database while installing Oracle Database 10g, you have the option of configuring your database so it can be managed by Oracle Enterprise Manager 10g Grid Control Console or by Oracle Enterprise Manager 10g Database Control Console.

[Figure 1–3](#) shows the Management Options page, which allows you to select your database management options while installing Oracle Database 10g.

Figure 1–3 *Selecting Your Management Options While Installing Oracle Database 10g*



To select Grid Control Console as your management option, the Oracle Management Service must be installed on a network host. In addition, the Oracle Management Agent must be installed on the host where you are installing the database. Otherwise, the Grid Control Console option is unavailable and you must instead choose to manage your database with Database Control.

For most of the Oracle Database 10g installation types, you must choose either Database Control or Grid Control as your management option when you create a database during the installation.

However, if you create a database using one of the following methods, you can choose not to configure Database Control:

- Choosing to create a database during a custom installation
- Choosing the Advanced database configuration option during an Enterprise or Standard Edition installation
- Running Database Configuration Assistant (DBCA) after the installation

If you do not configure Database Control during the Oracle Database 10g installation, no *hostname_sid* directory is created in the resulting Oracle home directory (Figure 1-2).

1.2.6.2 Configuring Database Control with DBCA

The primary method for configuring an existing Oracle Database 10g database so it can be managed with Database Control is to use DBCA. You can use DBCA to create a new database or to reconfigure an existing database.

See Also: "Installing Oracle Software and Building the Database" in *Oracle Database 2 Day DBA* for more information about using DBCA to create a new database instance

To use DBCA to reconfigure your database so it can be managed with Database Control:

1. Log into the database host as a member of the administrative group that is authorized to install Oracle software and create and run the database.
2. Start DBCA, as follows:
 - On Windows, select **Start, point to Programs, Oracle - home_name, Configuration and Migration Tools, and then select Database Configuration Assistant.**
 - On UNIX, change directory to the `ORACLE_HOME/bin` directory and enter the following command:

```
$PROMPT> ./dbca
```

The DBCA Welcome page appears.

3. Advance to the Operations page and select **Configure Database Options.**
4. Advance to the Database page and select the database you want to configure.
5. Advance to the Management Options page (Figure 1-4) and select the following options:
 - **Configure the Database with Enterprise Manager**
 - **Use Database Control for Database Management**

6. Optionally, select the options for enabling e-mail notifications and enabling daily backups.

For more information about Enterprise Manager notifications and daily backups, click **Help** on the Management Options page.

7. Advance until the **Finish** button is available.
8. Click **Finish** to reconfigure the database so it uses Database Control.

After DBCA reconfigures the database, a new subdirectory appears in the Oracle home. This directory is named using the following format and contains Database Control configuration and state files specific to the database you just configured:

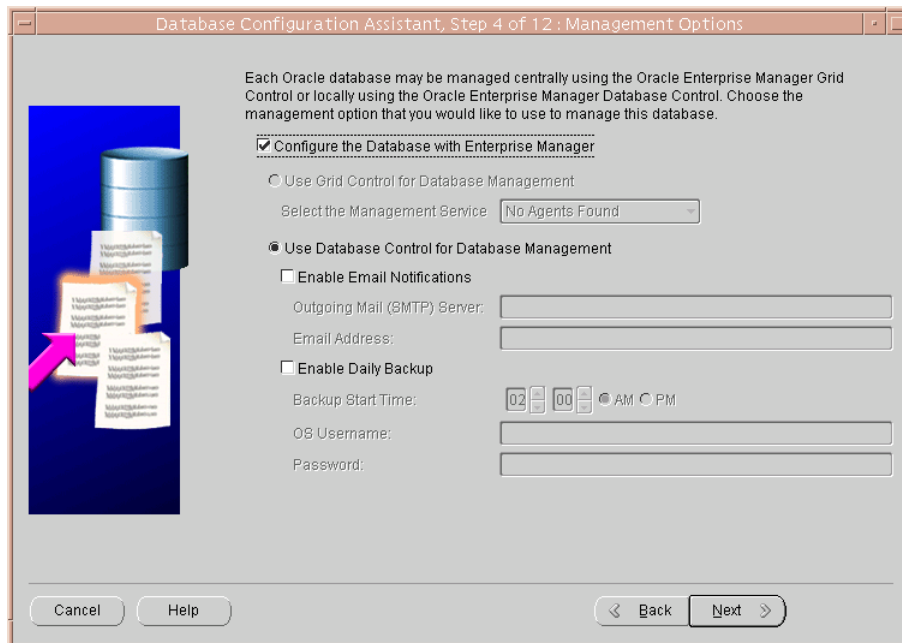
hostname_sid

For example:

mgmthost1.acme.com_myNewDB

Note that for cluster databases, the directories are named nodename_sid.

Figure 1–4 Management Options Page in DBCA



1.2.6.3 Configuring Database Control with EMCA

When you use DBCA to configure Oracle Database 10g, DBCA provides a graphical user interface to help you select Database Control options and to configure other aspects of your database.

However, if you want to use the operating system command line to configure Database Control, you can use the Enterprise Manager Configuration Assistant (EMCA).

To configure Database Control with EMCA:

1. Set the following environment variables to identify the Oracle home and the system identifier (SID) for the database you want to manage:
 - ORACLE_HOME

- ORACLE_SID
2. Change directory to the ORACLE_HOME/bin directory.
 3. Start EMCA by entering the following command with any of the optional command-line arguments shown in [Table 1-3](#):

```
$PROMPT> ./emca
```

Depending upon the arguments you include on the EMCA command line, EMCA prompts you for the information required to configure Database Control.

For example, enter the following command to configure Database Control so it will perform automatic daily backups of your database:

```
$PROMPT> ./emca -config dbcontrol db -backup
```

EMCA commands are of the form:

```
emca [operation] [mode] [type] [flags] [parameters]
```

[Table 1-3](#) describes the valid execution operations and modes, and lists the optional parameters in brackets. [Table 1-4](#) discusses the flags and their behavior, while [Table 1-5](#) defines the optional parameters in detail. EMCA parameters are of the form [-parameterName parameterValue]. Multiple parameters can be used in combination at the command line.

Table 1-3 EMCA Command-Line Operations

Command	Description
emca -h --h -help --help	Use this option to display the Help message for the EMCA utility. The options described in Table 1-3 , Table 1-4 , and Table 1-5 , and the valid parameters you may include are listed.
emca --version	Prints the version information associated with EMCA.
emca -config dbcontrol db [-repos (create recreate)] [-cluster] [-silent] [-backup] [parameters]	Configures Database Control for a database. Options include creating (or recreating) Database Control repository, configuring automatic backups, and performing these operations on a cluster database.
emca -config centralAgent (db asm) [-cluster] [-silent] [parameters]	Configures central agent management for a database or an Automatic Storage Management (ASM) instance. Options include performing this operation on a cluster environment. This operation will configure the database so that it can be centrally managed by the Oracle Enterprise Manager 10g Grid Control Console. To use this option, you must have previously installed the Oracle Management Service component of Enterprise Manager on a network host. In addition, the Oracle Management Agent must be installed on the host where you are running the database.
emca -config all db [-repos (create recreate)] [-cluster] [-silent] [-backup] [parameters]	Configures both Database Control and central agent management for a database. The possible configuration options are similar to those described above.
emca -deconfig dbcontrol db [-repos drop] [-cluster] [-silent] [parameters]	Deconfigures Database Control for a database. Options include dropping the Database Control repository and performing these operations on a cluster database. For example, you might use this command to remove the Database Control configuration from a database you are planning to delete. In such a scenario, remove the Database Control configuration before physically deleting the database. This operation does not remove the actual database or its data files.

Table 1–3 (Cont.) EMCA Command-Line Operations

Command	Description
<code>emca -deconfig centralAgent (db asm) [-cluster] [-silent] [parameters]</code>	Deconfigures central agent management for a database or an ASM instance. Options include performing this operation on a cluster environment. For example, you might use this command to remove the central agent management configuration from a database you are planning to delete. In such a scenario, remove the central agent management configuration before physically deleting the database. This operation does not remove the actual database or its data files.
<code>emca -deconfig all db [-reposit drop] [-cluster] [-silent] [parameters]</code>	Deconfigures both Database Control and central agent management for a database. The possible deconfiguration options are similar to those described above.
<code>emca -addInst (db asm) [-silent] [parameters]</code>	Configures Enterprise Manager for a new cluster instance of a database or ASM storage. For more information, refer to Section 1.2.6.5 .
<code>emca -deleteInst (db asm) [-silent] [parameters]</code>	Deconfigures Enterprise Manager for a specific instance of a cluster database or ASM storage. This is discussed further below, in Section 1.2.6.5 .
<code>emca -reconfig ports [-cluster] [parameters]</code>	Explicitly reassigns Database Control ports. Options include performing this operation on a cluster environment. For more information, refer to Section 1.2.6.6 .
<code>emca -reconfig dbcontrol -cluster [-silent] [parameters]</code>	Reconfigures Database Control deployment for a cluster database. Note that this command must be used with the "-cluster" option. For more information, refer to Section 1.2.6.5 .
<code>emca -displayConfig dbcontrol -cluster [-silent] [parameters]</code>	Displays information about the current deployment configuration of Database Control in a cluster environment. Note that this command must be used with the "-cluster" option. For more information, refer to Section 1.2.6.5 .
<code>emca -upgrade (db asm db_ asm) [-cluster] [-silent] [parameters]</code>	Upgrades the configuration of an earlier version of Enterprise Manager to the current version. This operation can be performed for database, ASM, or database and ASM instances together simultaneously. This does not upgrade the actual database or ASM instances, nor does it upgrade the Enterprise Manager software. Instead, it upgrades the configuration files for the specified instance so that they are compatible with the current version of the Enterprise Manager software. EMCA will attempt to upgrade all instances of the specified database and/or ASM target on the host, across all Oracle Homes (since it is likely that certain target properties, such as listener port or Oracle Home, have changed).
<code>emca -restore (db asm db_ asm) [-cluster] [-silent] [parameters]</code>	Restores the current version of Enterprise Manager configuration to an earlier version. This is the inverse of the "-upgrade" option (and will reverse any changes that result from an "-upgrade" operation), and as such, the options are similar.

Table 1–4 EMCA Command-Line Flags

Flag	Description
<code>db</code>	Performs the operation for a database (including cluster databases). Use this option for databases that use Automatic Storage Management (ASM) to store the data files. If a database is using ASM, all the configuration operations and modes described above (except for "-upgrade" and "-restore") will detect this automatically and apply the changes to both the database and ASM instance(s).

Table 1–4 (Cont.) EMCA Command-Line Flags

Flag	Description
asm	Performs the operation for an ASM-only instance (including cluster ASM instances).
db_asm	This flag can only be used in "-upgrade" and "-restore" mode. Performs the upgrade/restore operation for a database and an ASM instance together. Database and ASM instances may be upgraded or restored separately (that is, upgrading an ASM instance does not require upgrading the database instances it services). Hence, the Enterprise Manager configuration can be upgraded or restored separately for a database and its respective ASM instance.
-repos create	Creates a new Database Control management repository.
-repos drop	Drops the current Database Control management repository.
-repos recreate	Drops the current Database Control management repository and then recreates a new one.
-cluster	Performs the operation for a cluster database or ASM instance.
-silent	Performs the operation without prompting for additional information. If this mode is specified, all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. You can view a list of the available parameters by entering <code>emca -help</code> at the command line.
-backup	Configures automatic backup for a database. EMCA will prompt for daily automatic backup options. The default Enterprise Manager settings will be used to backup the database files. Note: If you use this option, EMCA will use the value of the <code>db_recovery_file_dest</code> initialization parameter to identify the flashback recovery area for the automated backups. If that parameter is not set, EMCA will generate an error. You can modify these settings later using the Maintenance page in Database Control. For more information, see the Database Control online Help.

Table 1–5 EMCA Command-Line Parameters

Parameter	Description
-respFile	Specifies the path of an input file listing parameters for EMCA to use while performing its configuration operation. For more information, refer to Section 1.2.6.4 .
-SID	Database system identifier
-PORT	Port number for the listener servicing the database
-ORACLE_HOME	Database Oracle Home, as an absolute path
-LISTENER_OH	Oracle Home from where the listener is running. If the listener is running from an Oracle Home other than the one on which the database is running, the parameter LISTENER_OH must be specified.
-HOST_USER	Host machine user name (for automatic backup)
-HOST_USER_PWD	Host machine user password (for automatic backup)
-BACKUP_SCHEDULE	Schedule in the form of "HH:MM" (for daily automatic backups)
-EMAIL_ADDRESS	E-mail address for notifications

Table 1–5 (Cont.) EMCA Command-Line Parameters

Parameter	Description
-MAIL_SERVER_NAME	Outgoing Mail (SMTP) server for notifications
-ASM_OH	Automatic Storage Management Oracle Home
-ASM_SID	System identifier for ASM instance
-ASM_PORT	Port number for the listener servicing the ASM instance
-ASM_USER_ROLE	User role for connecting to the ASM instance
-ASM_USER_NAME	User name for connecting to the ASM instance
-ASM_USER_PWD	Password for connecting to the ASM instance
-DBSNMP_PWD	Password for DBSNMP user
-SYSMAN_PWD	Password for SYSMAN user
-SYS_PWD	Password for SYS user
-SRC_OH	Oracle Home of the database with Enterprise Manager configuration to be upgraded/restored
-DBCONTROL_HTTP_PORT	Use this parameter to specify the port you use to display the Database Control Console in your Web browser. For more information, refer to Section 1.2.6.6 .
-AGENT_PORT	Use this parameter to specify the Management Agent port for Database Control. For more information, refer to Section 1.2.6.6 .
-RMI_PORT	Use this parameter to specify the RMI port for Database Control. For more information, refer to Section 1.2.6.6 .
-JMS_PORT	Use this parameter to specify the JMS port for Database Control. For more information, refer to Section 1.2.6.6 .
-CLUSTER_NAME	Cluster name (for cluster databases)
-DB_UNIQUE_NAME	Database unique name (for cluster databases)
-SERVICE_NAME	Database service name (for cluster databases)
-EM_NODE	Node from which Database Control console is to be run (for cluster databases). For more information, refer to Section 1.2.6.5 .
-EM_SID_LIST	Comma-separated list of SIDs for agent-only configurations, uploading data to -EM_NODE. For more information, refer to Section 1.2.6.5 .

1.2.6.4 Using an Input File for EMCA Parameters

Instead of answering a series of prompts when you run EMCA, you can use the `-respFile` argument to specify an input file. The input file you create must be in a format similar to the following example:

```
PORT=1521
SID=DB
DBSNMP_PWD=xpE234D
SYSMAN_PWD=KD0dk432
```

After you create an EMCA input file, you can use it on the command line as follows:

```
$PROMPT> ./emca -config dbcontrol db -respFile input_file_path
```

For example, to configure the Database Control to perform daily backups and create the Database Control Management Repository, create an input file similar to the one

shown in [Example 1–1](#) and enter the following command at the operating system prompt:

```
$PROMPT> ./emca -config dbcontrol db -repos create -backup -respFile input_file_path
```

Example 1–1 EMCA Input File that Configures Database Control for Automatic Backup and Creates the Database Control Management Repository

```
PORT=1521
SID=DB
DBSNMP_PWD=dow31224
SYSMAN_PWD=squN3243
HOST_USER=johnson
HOST_USER_PWD=diTf32of
SYS_PWD=q1Kj4352
BACKUP_SCHEDULE=06:30
```

1.2.6.5 Using EMCA with Real Application Clusters

Oracle Real Application Clusters (RAC) provides a high availability database environment spanning multiple hosts. Each cluster may be made up of multiple cluster databases, each of which consists of multiple cluster database instances. A cluster database is available as long as one of its instances is available.

Each EMCA command can be used in Real Application Clusters environments; certain commands are only applicable in cluster setups. To indicate that you have a cluster database, use the `-cluster` flag which is available in almost every EMCA operational mode.

When you use EMCA to configure Database Control for Real Application Clusters, you configure the Database Control for each instance in the cluster. However, by default, the Database Control Console will only start on the local node. On every other node of the cluster, only the Enterprise Manager agent will start. This is because the Database Control Console opens a number of connections to the database. If an instance of the console is running on every host in the cluster, then you may easily exceed the maximum number of permitted open connections on a 32-node or 64-node environment.

To remedy this, the Database Control Console is only started on the local node. On every other node, the commands `emctl start dbconsole` and `emctl stop dbconsole` only start and stop the agent. Each of the remote agents will upload their respective data to the console running on the local node, from where you can monitor and manage all the targets in the cluster.

However, note that if you upgrade a 10g Release 1 cluster database (configured with Database Control) to 10g Release 2, the 10g Release 1 Database Control configuration will be retained. The console will still be started on each individual node. If you wish to modify the configuration, use the following command:

```
emca -reconfig dbcontrol -cluster -EM_NODE <nodename> -EM_SID_LIST <SID list>
```

where `<nodename>` is the public name of the node and `<SID list>` is a comma-separated list of database system identifiers. This command reconfigures the current Database Control setup and:

1. Starts a Database Control Console on `<nodename>`, if one has not been started yet.
2. Redirects the agents monitoring the database instances in `<SID list>` so that they upload their data to the console running on `<nodename>`. Also, agents monitoring

database instances on <nodename> will also upload their data to the local console. Note that if you do not pass `-EM_NODE` or `-EM_SID_LIST` at the command line, you will be prompted for them.

`-EM_NODE` defaults to the local node if not specified when prompted. `-EM_SID_LIST` defaults to all database instances if not specified.

You may use this command to start the console on more than one node. For instance, on an 8-node cluster with <node1, node2, node3, node4, node5, node6, node7, node8> and database instances <oradb1, oradb2, oradb3, oradb4, oradb5, oradb6, oradb7, oradb8>, you can run the following commands in succession:

```
$PROMPT> emca -reconfig dbcontrol -cluster -EM_NODE node1 -EM_SID_LIST
oradb2,oradb3,oradb4
$PROMPT> emca -reconfig dbcontrol -cluster -EM_NODE node5 -EM_SID_LIST
oradb6,oradb7,oradb8
```

In this scenario, there will be two Database Control consoles running, one on node1 and the other on node5. From either of these consoles, you can manage and monitor all targets in the cluster.

For information on the current cluster configuration, you can run:

```
emca -displayConfig dbcontrol -cluster
```

The above command prompts for the database unique name for the cluster database. This will print the current configuration onto the screen, indicating the nodes that have consoles running on them and the consoles where each agent is uploading.

On cluster databases, another common operation is the creation and deletion of database instances. After you create a new instance, you can run EMCA to configure Database Control or central agent management for that instance using the command `emca -addInst db`. Running EMCA does not create the actual database instance; it only configures Enterprise Manager so that you can manage the instance in a way consistent with the rest of the cluster database instances. When configuring Enterprise Manager for a new instance, run the EMCA command only after you have created the instance. Also, run the command from a node in the cluster that already has Enterprise Manager configured for its associated database instance, as these configuration settings will be propagated to the new instance. Do not run this command from the node on which the new instance was created. Note that this option can be used only in a Real Application Clusters environment, so you do not need to use the `-cluster` option on the command line.

To deconfigure Enterprise Manager for a specific database instance (typically before the database instance is deleted), use the inverse command, `emca -deleteInst db`. Running EMCA does not delete the database instance; it only removes the Enterprise Manager configuration so that you will no longer be able to manage the instance with Enterprise Manager. Ensure that you run the EMCA command before you delete the actual cluster database instance. Also, ensure that you run the command from a different node and not from the node on which the database instance will be deleted. Note that this option can be used only in a Real Application Clusters environment, so you do not need to use the `-cluster` option on the command line.

For more information, see [Table 1–3](#) which describes EMCA command-line operations.

1.2.6.6 Specifying the Ports Used By the Database Control

When you initially install Oracle Database 10g or configure the Database Control with EMCA, the Database Control uses a set of default system ports. For example, by default, you access Database Control using port 1158 in 10g Release 2, as in:


```
http://host.domain:1158/em
```

This is the default port assigned to Database Control by the Internet Assigned Numbers Authority (IANA). Likewise, the default Database Control Agent port, as assigned by the IANA, is 3938.

To use ports other than the default ports, use the following EMCA command-line arguments when you initially configure the Database Control with EMCA. Alternatively, you can explicitly assign ports after configuring Database Control using the following command:

```
emca -reconfig ports [-cluster]
```

Note: You can also use the following EMCA command-line arguments to configure Database Control after you have installed and configured Oracle Database 10g.

The following list summarizes the EMCA command-line arguments that control the standard Database Control port assignments:

- `-DBCONTROL_HTTP_PORT <port_number>`

This port number is used in the Database Control Console URL. For example, if you set this port to 5570, you can then display the Database Control Console using the following URL:

```
http://host.domain:5570/em
```

- `-RMI_PORT <port_number>`

This port number is used by the Remote Method Invocation (RMI) system, which is part of the J2EE software required by Database Control.

- `-JMS_PORT <port_number>`

This port is used by the OC4J Java Message Service (JMS), which is part of the J2EE software required by Database Control.

- `-AGENT_PORT <port_number>`

This port is used by the Database Control Management Agent, which is monitoring and administering the database for the Database Control.

1.2.6.7 EMCA Troubleshooting Tips

The following section describes some troubleshooting tips to consider when using EMCA to configure the Database Control:

- [Using EMCA After Changing the Database Listener Port](#)
- [Upgrading Database or ASM Instances with 10g Release 2 Grid Control Agents](#)

1.2.6.7.1 Using EMCA After Changing the Database Listener Port If you change the listener port of the database after you have configured Database Control, the database status will appear as down. To reconfigure Database Control so it uses the new listener port, run the EMCA command using the `-config dbcontrol db [-cluster]` command-line arguments.

1.2.6.7.2 Upgrading Database or ASM Instances with 10g Release 2 Grid Control Agents When upgrading a 10g Release 1 database and/or ASM instance that was configured for Oracle Enterprise Manager (either Database Control or a Grid Control central agent) to

10g Release 2, all Enterprise Manager targets on the relevant host(s) referring to the upgraded instance(s) will be updated automatically. This is because the upgrade involves altering the instance's Oracle Home, port, or other target-associated properties. However, some of these targets on the host(s) will not be updated successfully during the upgrade if they are managed by a 10g Release 2 Grid Control Agent. To update these targets, in the Home page for the upgraded database (or ASM) target, click the "Monitoring Configuration" link. On this page, you can update the required properties such as Oracle Home, listener port and so on to the correct values.

1.3 Enabling Enterprise Manager Accessibility Features

As part of the effort to make Oracle products, services, and supporting documentation accessible and usable to the disabled community, Enterprise Manager offers several features that make management data available to users of assistive technology.

To enable these features and provide for full accessibility, you must modify two configuration settings, which are described in the following sections:

- [Enabling Enterprise Manager Accessibility Mode](#)
- [Providing Textual Descriptions of Enterprise Manager Charts](#)

1.3.1 Enabling Enterprise Manager Accessibility Mode

Enterprise Manager takes advantage of user interface development technologies that improve the responsiveness of some user operations. For example, when you navigate to a new record set in a table, Enterprise Manager does not redisplay the entire HTML page.

However, this performance-improving technology is generally not supported by screen readers. To disable this feature, and as a result, make the Enterprise Manager HTML pages more accessible for disabled users, use the following procedure.

Note: The following procedure is valid for both Grid Control Console and Database Control installations. Differences in the location of configuration files is noted where applicable.

For information on enabling accessibility for the Application Server Control Console, see "Managing and Configuring the Application Server Control" in the *Oracle Application Server 10g Administrator's Guide*.

1. Locate the `uix-config.xml` configuration file.

To locate the `uix-config.xml` file in a Grid Control Console installation, change directory to the following location in the Management Service home:

```
ORACLE_HOME/j2ee/OC4J_EM/applications/em/em/WEB-INF (Grid Control)
```

To locate the `uix-config.xml` file in a Oracle Database 10g installation, change directory to the following location in the database home:

```
ORACLE_HOME/oc4j/j2ee/oc4j_applications/applications/em/em/WEB-INF (Database Control)
```

2. Open the `uix-config.xml` file using a text editor and locate the following entry:

```
<!-- An alternate configuration that disables accessibility features -->  
<default-configuration>
```

```
<accessibility-mode>inaccessible</accessibility-mode>
</default-configuration>
```

3. Change the value of the `accessibility-mode` property from `inaccessible` to `accessible`.
4. Save and close the file.
5. Restart the Oracle Management Service (if you are modifying a Grid Control Console installation) or restart the Database Control (if you are modifying an Oracle Database 10g installation).


1.3.2 Providing Textual Descriptions of Enterprise Manager Charts

Throughout Enterprise Manager, charts are used to display performance data. For most users, these charts provide a valuable graphical view of the data that can reveal trends and help identify minimum and maximum values for performance metrics.

However, charts do not convey information in a manner that can be read by a screen reader. To remedy this problem, you can configure Enterprise Manager to provide a complete textual representation of each performance chart. By default, support for the textual representation of charts is disabled. When textual description for charts is enabled, Enterprise Manager displays a small icon for each chart that can be used as a drill-down link to the textual representation.

Figure 1–5 shows an example of the icon that displays beneath Enterprise Manager charts when you have enabled the textual representation of charts.

Figure 1–5 Icon Representing the Textual Representation of a Chart



To enable the drill-down icon for the textual representation of charts:

1. Locate the `web.xml` configuration file.

To locate the `web.xml` file in a Grid Control Console installation, change directory to the following location in the Management Service home:

```
ORACLE_HOME/j2ee/OC4J_EM/applications/em/em/WEB-INF
```

To locate the `web.xml` file in a Oracle Database 10g installation, change directory to the following location in the database home:

```
ORACLE_HOME/oc4j/j2ee/oc4j_applications/applications/em/em/WEB-INF
```

2. Open the `web.xml` file with your favorite text editor and locate the following six lines of the file:

```
<!-- Uncomment this to enable textual chart descriptions
```

```
<context-param>
<param-name>enableChartDescription</param-name>
<param-value>>true</param-value>
</context-param>
-->
```

3. Remove comments from this section by deleting the first line and the last line of this section so that the section consists of only these 4 lines:

```
<context-param>
<param-name>enableChartDescription</param-name>
<param-value>>true</param-value>
</context-param>
```

4. Save and exit the file.
5. Restart the Management Service (if you are modifying a Grid Control Console installation) or restart the Database Control (if you are modifying an Oracle Database 10g installation).

Starting and Stopping Enterprise Manager Components

To start and stop the Management Service, the Management Agent, the Grid Control Console, the Application Server Control Console, and Database Control, use the Enterprise Manager command line utility (`emctl`).

The capabilities of the command-line utility can be broken down into the following categories:

- [Controlling the Oracle Management Agent](#)
- [Controlling the Oracle Management Service](#)
- [Controlling the Application Server Control](#)
- [Controlling the Database Control on UNIX](#)
- [Guidelines for Starting Multiple Enterprise Manager Components on a Single Host](#)
- [Starting and Stopping Oracle Enterprise Manager 10g Grid Control](#)
- [Additional Management Agent Commands](#)

2.1 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](#)

2.1.1 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 2-1](#).

For example, to stop the Management Agent, enter the following commands:

```
$PROMPT> cd AGENT_HOME/bin
$PROMPT> ./emctl stop agent
```

Table 2–1 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 2–1).

Example 2–1 Checking the Status of the Management Agent

```

$PROMPT> ./emctl status agent
Oracle Enterprise Manager 10g Release 10.2.0.0.0
Copyright (c) 1996, 2005 Oracle Corporation. All rights reserved.
-----
Agent Version      : 10.2.0.0.0
OMS Version       : 10.2.0.0.0
Protocol Version  : 10.2.0.0.0
Agent Home        : /scratch/OracleHomesX/agent10g
Agent binaries   : /scratch/OracleHomesX/agent10g
Agent Process ID  : 17604
Parent Process ID : 17587
Agent URL        : https://stadj32.us.oracle.com:3872/emd/main/
Repository URL   : https://stadj32.us.oracle.com:1159/em/upload
Started at      : 2005-09-13 01:31:11
Started by user  : test
Last Reload     : 2005-09-13 01:31:11
Last successful upload          : 2005-09-13 01:39:01
Total Megabytes of XML files uploaded so far : 0.28
Number of XML files pending upload          : 0
Size of XML files pending upload(MB)       : 0.00
Available disk space on upload filesystem   : 8.36%
Last successful heartbeat to OMS           : 2005-09-13 01:38:51
-----
Agent is Running and Ready
$PROMPT>

```

2.1.2 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 2.2.1](#) 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 2–2](#) describes the Windows services that you use to control the Management Agent.

Table 2–2 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.

2.1.3 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 2–1](#)).

2.2 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](#)

2.2.1 Controlling the Management Service on UNIX

There are two methods for starting and stopping the Oracle Management Service on UNIX systems. You can use the Oracle Process Management and Notification (OPMN) utility, or you can use a set of `emctl` commands.

The following sections describe these two approaches to controlling the Management Service, as well as information about starting and stopping OracleAS Web Cache, which is also required by the Grid Control Console:

- [Using OPMN to Start and Stop the Management Service](#)
- [Using emctl to Start, Stop, and Check the Status of the Oracle Management Service](#)
- [Starting and Stopping Oracle Application Server Web Cache](#)

2.2.1.1 Using OPMN to Start and Stop the Management Service

One method of starting and stopping the Management Service by using the Oracle Process Management and Notification (OPMN) utility. The OPMN utility (`opmnctl`) is a standard command used to start and stop components of the Oracle Application Server instance.

The Management Service is a J2EE application running in an Oracle Application Server Containers for J2EE (OC4J) instance within the Application Server. As a result, the following command will start all the components of the Oracle Application Server instance, including the OC4J_EM instance and the Management Service application:

```
$PROMPT> cd opmn/bin
$PROMPT> ./opmnctl startall
```

Similarly, the following command will stop all the components of the Oracle Application Server instance:

```
$PROMPT> ./opmnctl stopall
```

If you want to start only the components necessary to run the Management Service, you can use the Enterprise Manager command-line utility.

2.2.1.2 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 2-3](#).

For example, to stop the Management Service, enter the following commands:

```
$PROMPT> cd bin
$PROMPT> ./emctl stop oms
```


Table 2–3 Starting, Stopping, and Checking the Status of the Management Service

Command	Purpose
emctl start oms	<p>Starts the Oracle Application Server components required to run the Management Service J2EE application. Specifically, this command starts OPMN, the Oracle HTTP Server, and the OC4J_EM instance where the Management Service is deployed.</p> <p>Note: The <code>emctl start oms</code> command does not start Oracle Application Server Web Cache. For more information, see "Starting and Stopping Oracle Application Server Web Cache" on page 2-5.</p>
emctl stop oms	<p>Stops the Management Service.</p> <p>Note that this command does not stop the other processes that are managed by the Oracle Process Manager and Notification Server (OPMN) utility.</p> <p>To stop the other Oracle Application Server components, such as the Oracle HTTP Server and Oracle Application Server Web Cache, see "Starting and Stopping Oracle Enterprise Manager 10g Grid Control" on page 2-10.</p>
emctl status oms	Displays a message indicating whether or not the Management Service is running.

2.2.1.3 Starting and Stopping Oracle Application Server Web Cache

By default, when you install Oracle Enterprise Manager 10g, the Grid Control Console is configured to use Oracle Application Server Web Cache.

See Also: *Oracle Application Server Web Cache Administrator's Guide* for more information about Oracle Application Server Web Cache

Oracle Application Server Web Cache not only improves the performance of the Grid Control Console, but also makes it possible to measure the end-user performance of the Enterprise Manager Web application.

See Also: [Chapter 6, "Configuring Services"](#) for more information about End-User Performance Monitoring and the Enterprise Manager Web Application

To view the Grid Control Console using Oracle Application Server Web Cache, you access the Grid Control Console using the standard port number assigned during the Oracle Enterprise Manager 10g installation procedure. You can find this default port number (usually 7777) in the `setupinfo.txt` file, which is copied to the following directory during the Enterprise Manager installation procedure:

```
AS_HOME/Apache/Apache
```

If Oracle Application Server Web Cache is not running, you will receive an error message, such as the following, if you try to access the Grid Control Console using the default port number:

```
HTTP 500 - Internal server error
```

To start Oracle Application Server Web Cache:

1. Change directory to the `ORACLE_HOME/opmn/bin` directory in the Management Service home.

- Use the appropriate command described in [Table 2–4](#).

For example, to stop Oracle Application Server Web Cache, enter the following commands:

```
$PROMPT> cd opmn/bin
$PROMPT> ./opmnctl stopproc ias-component=WebCache
```

Table 2–4 Starting, Stopping, and Checking the Status of Oracle Application Server Web Cache

Command	Purpose
<code>opmnctl startproc ias-component=WebCache</code>	Starts Oracle Application Server Web Cache.
<code>opmnctl stopproc ias-component=WebCache</code>	Stops Oracle Application Server Web Cache.
<code>opmnctl status</code>	Displays a message showing the status of all the application server components managed by OPMN, including Oracle Application Server Web Cache.

2.2.2 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 2.2.1](#) 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 2–5](#) describes the Windows services that you use to control the Oracle Management Service.

Table 2–5 Summary of Services Installed and Configured When You Install the Oracle Management Service on Windows

Component	Service Name Format	Description
Application Server Control	Oracle<oracle_home>ASControl For example: OracleOraHome1ASControl	Use this Service to start and stop the Application Server Control for the Oracle Application Server instance that was installed and configured to deploy the Management Service J2EE application.

Table 2–5 (Cont.) Summary of Services Installed and Configured When You Install the Oracle Management Service on Windows

Component	Service Name Format	Description
Oracle Process Management and Notification (OPMN)	Oracle<oracle_home>ProcessManager For example: OracleOraHome1ProcessManager	Use this service to start and stop all the components of the Oracle Application Server instance that were installed and configured to deploy the Management Service J2EE application. Use this service to start and stop the Management Service and all its related components, including OC4J, Oracle HTTP Server, and OracleAS Web Cache, which by default must be running in order for you to access the Grid Control Console.

2.3 Controlling the Application Server Control

The Application Server Control is a component of Oracle Enterprise Manager 10g that is installed as part of any Oracle Application Server installation. The following sections describe how to start and stop the Application Server Control:

- [Starting and Stopping the Application Server Control on UNIX](#)
- [Starting and Stopping the Application Server Control on Windows](#)

See Also: *Oracle Application Server 10g Administrator's Guide* for more information about:

- Using `emctl` to control the Application Server Control Console
- Starting and stopping the Application Server Control Console on Windows
- Displaying disabled components of the Application Server

2.3.1 Starting and Stopping the Application Server Control on UNIX

To control the Application Server Control Console on UNIX systems, you use the `emctl` command line utility that is available in the `IAS_HOME/bin` directory after you install Oracle Application Server.

To start the Application Server Control Console, change directory to the `IAS_HOME/bin` directory and then enter the following command:

```
$PROMPT> ./emctl start iasconsole
```

To stop the Application Server Control Console, enter the following command:

```
$PROMPT> ./emctl stop iasconsole
```

2.3.2 Starting and Stopping the Application Server Control on Windows

To start or stop the Application Server 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 Application Server 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, followed by the word "ASControl." For example, if you specified AS10g as the Oracle Home, the Service name would be:

```
OracleAS10gASControl
```

3. After you locate the service, you can use the Services control panel to start or stop the Application Server Control service.

By default, the Application Server Control service is configured to start automatically when the system starts.

You can also start the Oracle Application Server Control console (iasconsole) on Windows using `emctl start iasconsole` as described in [Section 2.3.1](#).

2.4 Controlling the Database Control on UNIX

The Oracle Enterprise Manager 10g Database Control Console is a component of Oracle Enterprise Manager 10g that is installed as part of any Oracle Database 10g installation.

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 10g.

2.4.1 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
```

2.4.2 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
```

2.4.3 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 DBd10g as the Oracle Home, the Service name would be:

```
OracleDB10gDBControl
```

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.

You can also start the Database Control on Windows using `emctl start iasconsole` as described in [Section 2.4.2](#).

2.5 Guidelines for Starting Multiple Enterprise Manager Components on a Single Host

Oracle Enterprise Manager 10g components are used to manage a variety of Oracle software products. For example, each time you install Oracle Application Server 10g (9.0.4) instance, you also install an Application Server Control. Similarly, each time you install Oracle Database 10g, you install a Database Control. In addition, if you want to centrally manage your system with Database Control, the Management Agent is also installed on each host you monitor.

In most cases, in a production environment, you will want to distribute your database and application server instances among multiple hosts to improve performance and availability of your software resources. However, in rare cases where you must install multiple application servers or databases on the same host, consider the following guidelines.

When you start Application Server 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.

For example, suppose you have installed OracleAS Infrastructure 10g, the J2EE and Web Cache application server installation type, and the Management Agent on the same host. 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 `opmnctl startall` command to start all the OPMN-managed processes in the OracleAS Infrastructure 10g home directory.
2. Wait 15 seconds.

3. Use the `emctl start iasconsole` command to start the Application Server Control in the OracleAS Infrastructure 10g home directory.
4. Wait 15 seconds.
5. Use the `opmnctl startall` command to start all the OPMN-managed processes in the J2EE and Web Cache home directory.
6. Wait 15 seconds.
7. Use the `emctl start iasconsole` command to start the Application Server Control in the J2EE and Web Cache home directory.
8. Wait 15 seconds.
9. 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.

2.6 Starting and Stopping Oracle Enterprise Manager 10g Grid Control

As described in the previous sections, you use separate commands to control the Oracle Management Service, Oracle Management Agent, and the Oracle Application Server 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 10g 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.

2.6.1 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 2-4

3. Start OracleAS Web Cache:

```
$PROPMT> $ORACLE_HOME/opmn/bin/opmnctl startproc ias-component=WebCache
```

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 2-1

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.

5. Optionally, start the Application Server Control Console, which is used to manage the Oracle Application Server instance that is used to deploy the Management Service:

```
$PROMPT> $ORACLE_HOME/bin/emctl start iasconsole
```

See Also: "[Controlling the Application Server Control](#)" on page 2-7

2.6.2 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 2-4

2. If necessary, stop the Application Server Control Console, which is used to manage the Oracle Application Server instance used to deploy the Management Service:

```
$PROMPT> $ORACLE_HOME/bin/emctl stop iasconsole
```

See Also: ["Controlling the Application Server Control"](#) on page 2-7

3. Stop all the Oracle Application Server components, such as the Oracle HTTP Server the OracleAS Web Cache:

```
$PROMPT> $ORACLE_HOME/opmn/bin/opmnctl stopall
```

See Also: *Oracle Application Server 10g Administrator's Guide*

4. 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 2-1

Note: Be sure to run the `emctl stop agent` command in the Oracle Management Agent home directory and not in the Oracle Application Server home directory.

5. 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
```

2.7 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](#)

2.7.1 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 2-6](#).

Table 2-6 *Manually Reloading and Uploading Management Data*

Command	Purpose
<code>emctl upload</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</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.

2.7.2 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.

For more information about the monitoring credentials for the Management Repository, see ["Changing the SYSMAN Password"](#) on page 8-5.

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 10g, 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 10g 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 10g Grid Control Console or by using the Enterprise Manager command line utility (`emctl`).

2.7.2.1 Using the Grid Control Console to Modify the Monitoring Credentials

To modify the password for the DBSNMP account in the Oracle Enterprise Manager 10g 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**.

2.7.2.2 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](#)" on page 2-15.

[Example 2-2](#) shows an example of the prompts and the output you receive from the command.

Example 2-2 Modifying the Database Monitoring Credentials

```
$PROMPT>./emctl config agent credentials emrep10.acme.com:oracle_database
Oracle Enterprise Manager 10g Release 10.1.0.2.0
Copyright (c) 2002, 2003 Oracle Corporation. All rights reserved.
Name = emrep10.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
```

2.7.3 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 [AGENT_HOME]
```

[Example 2-3](#) shows the typical output of the command.

Example 2-3 Listing the Targets on a Managed Host

```
./emctl config agent listtargets
Oracle Enterprise Manager 10g Release 10.1.0.2.0
Copyright (c) 2002, 2003 Oracle Corporation. All rights reserved.
[usunnab08.us.oracle.com, host]
[LISTENER_usunnab08.us.oracle.com, oracle_listener]
[EnterpriseManager.usunnab08.us.oracle.com_HTTP Server, oracle_apache]
[EnterpriseManager.usunnab08.us.oracle.com_home, oc4j]
[EnterpriseManager.usunnab08.us.oracle.com_Web Cache, oracle_webcache]
[EnterpriseManager.usunnab08.us.oracle.com, oracle_ias]
[EnterpriseManager.usunnab08.us.oracle.com_OC4J_EM, oc4j]
[EnterpriseManager.usunnab08.us.oracle.com_OC4J_Demos, oc4j]
[EM_Repository, oracle_emrep]
[usunnab08.us.oracle.com:1813, oracle_emd]
[EM Website, website]
[emrepl0.us.oracle.com, oracle_database]
```

2.7.4 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 10g 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 2-7](#).

Note: When you start a blackout, you must identify the target or targets affected by the blackout. To obtain the correct target name and target type for a target, see "[Listing the Targets on a Managed Host](#)" on page 2-15.

Table 2-7 Summary of Blackout Commands

Blackout Action	Command
Set an immediate blackout on a particular target or list of targets	<pre>emctl start blackout <Blackoutname> [<Target_name>[:<Target_Type>]]... [-d <Duration>]</pre> <p>Be sure to use a unique name for the blackout so you can refer to it later when you want to stop or check the status of the blackout.</p> <p>The <code>-d</code> option is used to specify the duration of the blackout. Duration is specified in [days] hh:mm where:</p> <ul style="list-style-type: none"> ■ days indicates number of days, which is optional ■ hh indicates number of hours ■ mm indicates number of minutes <p>If you do not specify a target or list of targets, Enterprise Manager will blackout the local host target. All monitored targets on the host are not blacked out unless a list is specified or you use the <code>-nodelevel</code> argument.</p> <p>If two targets of different target types share the same name, you must identify the target with its target type.</p>
Stop an immediate blackout	<pre>emctl stop blackout <Blackoutname></pre>
Set an immediate blackout for all targets on a host	<pre>emctl start blackout <Blackoutname> [-nodeLevel] [-d <Duration>]</pre> <p>The <code>-nodeLevel</code> option is used to specify a blackout for all the targets on the host; in other words, all the targets that the Management Agent is monitoring, including the Management Agent host itself. The <code>-nodeLevel</code> option must follow the blackout name. If you specify any targets after the <code>-nodeLevel</code> option, the list is ignored.</p>
Check the status of a blackout	<pre>emctl status blackout [<Target_name>[:<Target_Type>]]...</pre>

Use the following examples to learn more about controlling blackouts from the Enterprise Manager command line:

- To start a blackout called "bk1" for databases "db1" and "db2," and for Oracle Listener "ldb2," enter the following command:

```
$PROMPT> emctl start blackout bk1 db1 db2 ldb2:oracle_listener -d 5 02:30
```

The blackout starts immediately and will last for 5 days 2 hours and 30 minutes.

- To check the status of all the blackouts on a managed host:

```
$PROMPT> emctl status blackout
```

- To stop blackout "bk2" immediately:

```
$PROMPT> emctl stop blackout bk2
```

- To start an immediate blackout called "bk3" for all targets on the host:

```
$PROMPT> emctl start blackout bk3 -nodeLevel
```

- To start an immediate blackout called "bk3" for database "db1" for 30 minutes:

```
$PROMPT> emctl start blackout bk3 db1 -d 30
```

- To start an immediate blackout called "bk3" for database "db2" for five hours:

```
$PROMPT> emctl start blackout bk db2 -d 5:00
```

2.7.5 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.

There were bugs in Enterprise Manager Releases 10.1.0.2 and 10.1.0.3 RAC Management Agent installs that caused the Management Agent to be configured with a UTC timezone.

You can correct 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.

See Also: [Section 10.1.5, "Setting the Management Agent Time Zone"](#) on page 10-4

2.7.6 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 CollectinItem name oracle_security is what you would use in the emctl command to reevaluate the Open Ports metric.

Grid Control Common Configurations

Oracle Enterprise Manager 10g Grid Control is based on a flexible architecture, which allows you to deploy the Grid Control components in the most efficient and practical manner for your organization. This chapter describes some common configurations that demonstrate how you can deploy the Grid Control architecture in various computing environments.

This chapter presents the common configurations in a logical progression, starting with the simplest configuration and ending with a complex configuration that involves the deployment of high availability components, such as server load balancers, Oracle Real Application Clusters, and Oracle Data Guard.

This chapter contains the following sections:

- [About Common Configurations](#)
- [Summary of the Grid Control Architecture and Components](#)
- [Deploying Grid Control Components on a Single Host](#)
- [Managing Multiple Hosts and Deploying a Remote Management Repository](#)
- [Using Multiple Management Service Installations](#)
- [High Availability Configurations](#)

3.1 About Common Configurations

The common 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, the examples in this chapter assume you are using the OracleAS Web Cache port to access the Grid Control Console. By default, when you first install Grid Control, you display the Grid Control Console by navigating to the default OracleAS Web Cache port. In fact, you can modify your own configuration so administrators bypass OracleAS Web Cache and use a port that connects them directly to the Oracle HTTP Server.

For another example, in a production environment you will likely want to implement firewalls and other security considerations. The common configurations described in this chapter are not meant to show how firewalls and security policies should be implemented in your environment.

See Also: [Chapter 4, "Enterprise Manager Security"](#) for information about securing the connections between Grid Control components

[Chapter 5, "Configuring Enterprise Manager for Firewalls"](#) for information about configuring firewalls between Grid Control components

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.

3.2 Summary of the Grid Control Architecture and Components

The Grid Control architecture consists of the following software components:


- The Oracle Management Agent
- The Oracle Management Service
- The Oracle Management Repository
- The Oracle Enterprise Manager 10g Grid Control Console

See Also: *Oracle Enterprise Manager Concepts* for more information about each of the Grid Control components

The remaining sections of this chapter describe how you can deploy these components in a variety of combinations and across a single host or multiple hosts.

3.3 Deploying Grid Control Components on a Single Host

[Figure 3–1](#) 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 Grid Control installation procedure to install the **Enterprise Manager 10g Grid Control Using a New Database** installation type.

Figure 3–1 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 default OracleAS Web Cache port (for example, port 7777 on UNIX systems and port 80 on Windows systems) to connect to the Oracle HTTP Server. The Management Service retrieves data from the Management Repository as it is requested by the administrator using the Grid Control Console.

See Also: *Oracle Application Server Web Cache Administrator's Guide* for more information about the benefits of using OracleAS Web Cache

2. The Management Agent loads its data (which includes management data about all of 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 and bypasses OracleAS Web Cache. The default port for the upload URL is 4889 (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: ["Understanding the Enterprise Manager Directory Structure"](#) on page 1-1 for more information about the AGENT_HOME directory

3. The Management Service uses JDBC connections to load data into the repository database and to retrieve information from the repository so it can be displayed in the Grid Control Console. The repository connection information is defined in the following configuration file in the Management Service home directory:

```
ORACLE_HOME/sysman/config/emoms.properties (UNIX)
ORACLE_HOME\sysman\config\emoms.properties (Windows)
```

See Also: ["Reconfiguring the Oracle Management Service"](#) on page 10-8 for more information on modifying the repository connection information in the emoms.properties file

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=http://host1.acme.com:1831/emd/main/
```


In addition, by default, 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.

3.4 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. Such a configuration is shown in [Figure 3-2](#).

Figure 3–2 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 [Section 3.3](#).
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 bypasses OracleAS Web Cache and 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 repository database and to retrieve information from the repository so it can be displayed in the Grid Control Console. This remote connection is defined in the `emoms.properties` configuration file in the Management Service home directory.
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 3.3](#), the Management Agent includes a built-in HTTP listener so no Oracle HTTP Server is required on the Management Agent host.

3.5 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*.

The following sections provide more information about this configuration:

- [Determining When to Use Multiple Management Service Installations](#)
- [Understanding the Flow of Management Data When Using Multiple Management Services](#)

3.5.1 Determining When to Use Multiple Management Service Installations

Management Services not only exist as the receivers of upload information from Management Agents. They also retrieve data from the Management Repository. The Management Service renders this data in the form of HTML pages, which are requested by and displayed in the client Web browser. In addition, the Management Services perform background processing tasks, such as notification delivery and the dispatch of Enterprise Manager jobs.

As a result, the assignment of Management Agents to Management Services must be carefully managed and balanced. Improper distribution of load from Management Agents to Management Services may result in perceived:

- Sluggish user interface response
- Delays in delivering notification messages
- Backlog in monitoring information being uploaded to the Management Repository
- Delays in dispatching jobs

The following sections provide some tips for monitoring the load and response time of your Management Service installations:

- [Monitoring the Load on Your Management Service Installations](#)
- [Monitoring the Response Time of the Enterprise Manager Web Application Target](#)

3.5.1.1 Monitoring the Load on Your Management Service Installations

To keep the workload evenly distributed, you should always be aware of how many Management Agents are configured for each Management Service and monitor the load on each Management Service.

At any time, you can view a list of Management Agents and Management Services using the Setup tab of the Grid Control Console.

Use the charts on the Overview page of the Management Services and Repository tab to monitor:

- Loader backlog (files)

The Loader is part of the Management Service that pushes metric data into the repository at periodic intervals. If the Loader Backlog chart indicates that the backlog is high and Loader output is low, there is data pending load, which may indicate a system bottleneck or the need for another Management Service. The chart shows the total backlog of files totalled over all Oracle Management Services for the past 24 hours. Click the image to display loader backlog charts for each individual Management Service over the past 24 hours.

- Notification delivery backlog

The Notification Delivery Backlog chart displays the number of notifications to be delivered that could not be processed in the time allocated. Notifications are delivered by the Management Services. This number is summed across all Management Services and is sampled every 10 minutes. The graph displays the data for the last 24 hours. It is useful for determining a growing backlog of notifications. When this graph shows constant growth over the past 24 hours, then you may want to consider adding another Management Service, reducing the number of notification rules, and verifying that all rules and notification methods are useful and valid.

3.5.1.2 Monitoring the Response Time of the Enterprise Manager Web Application Target

The information on the Management Services and Repository tab can help you determine the load being placed on your Management Service installations. More importantly, you should also consider how the performance of your Management Service installations is affecting the performance of the Grid Control Console.

Use the **EM Website** Web Application target to review the response time of the Grid Control Console pages:

1. From the Grid Control Console, click the **Targets** tab and then click the **Web Applications** subtab.
2. Click **EM Website** in the list of Web Application targets.
3. In the Key Test Summary table, click **homepage**. The resulting page provides the response time of the Grid Control Console homepage URL.

See Also: The Enterprise Manager online help for more information about using the homepage URL and Application Performance Management (also known as Application Performance Monitoring) to determine the performance of your Web Applications

4. Click **Page Performance** to view the response time of some selected Grid Control Console pages.

Note: The Page Performance page provides data generated only by users who access the Grid Control Console by way of the OracleAS Web Cache port (usually, 7777).

5. Select **7 Days** or **31 Days** from the **View Data** menu to determine whether or not there are any trends in the performance of your Grid Control installation.


Consider adding additional Management Service installations if the response time of all pages is increasing over time or if the response time is unusually high for specific popular pages within the Grid Control Console.

Note: You can use Application Performance Management and Web Application targets to monitor your own Web applications. For more information, see [Chapter 6, "Configuring Services"](#)

3.5.2 Understanding the Flow of Management Data When Using Multiple Management Services

Figure 3–3 shows a typical environment where an additional Management Service has been added to improve the performance of the Grid Control environment.

Figure 3–3 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 OracleAS Web Cache and 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, bypassing OracleAS Web Cache.

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 `emoms.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 3.3](#), 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.

3.6 High Availability Configurations

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.

One way to protect your Grid Control components is to use high availability software deployment techniques, which usually involve the deployment of hardware server load balancers, Oracle Real Application Clusters, and Oracle Data Guard.

Note: The following sections do not provide a comprehensive set of instructions for configuring Grid Control for high availability. The examples here are shown only to provide examples of some common configurations of Grid Control components. These examples are designed to help you understand some of your options when you deploy Grid Control in your environment.

For a complete discussion of configuring Oracle products for high availability, refer to *Oracle High Availability Architecture and Best Practices*

Refer to the following sections for more information about common Grid Control configurations that take advantage of high availability hardware and software solutions:

- [Load Balancing Connections Between the Management Agent and the Management Service](#)
- [Load Balancing Connections Between the Grid Control Console and the Management Service](#)
- [Configuring the Management Repository for High Availability](#)

3.6.1 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, you should be aware that 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 re-attempts to send the information later.

To protect against the situation where a Management Service is unavailable, you can use a server load balancer between the Management Agents and the Management Services.

Note: To configure the Management Services for High Availability, you need a storage device that is shared by all the Management Services. The shared storage can be a NFS mounted disk accessible to all Management Services. For achieving a truly high available deployment, a sharable file system like Network Appliance™ Filer is recommended.

See "[Configuring the Management Services for High Availability](#)" on page 3-10 for steps on configuring the Management Services for High Availability.

However, if you decide to implement such a configuration, be sure to review the following sections carefully before proceeding:

- [Configuring the Management Services for High Availability](#)
- [Understanding the Flow of Data When Load Balancing the Upload of Management Data](#)
- [Configuring a Server Load Balancer for Management Agent Data Upload](#)

3.6.1.1 Configuring the Management Services for High Availability

The Management Service for Grid Control 10g Release 2 has a new 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 amongst themselves the workload of uploading files into the Management Repository. Should a Management Service go down for some reason, its workload is taken up by surviving Management Services.

Configuring the Shared Filesystem Loader

To configure the Management Service to use Shared Filesystem Loader, you must use the `emctl config oms loader` command.

1. Stop all the Management Services.
2. Configure a shared receive directory that is accessible by all Management Services.


3. Run `emctl config oms loader -shared yes -dir <loader directory>` individually on all Management Services hosts, where `<loader directory>` is the full path to the shared receive directory.
4. Restart all Management Services.

Caution: Shared Filesystem Loader mode should be configured on *all* the Management Services in your Grid Control deployment using the previous steps. Management Services will fail to start if all the Management Services are not configured to run in the same mode.

3.6.1.2 Understanding the Flow of Data When Load Balancing the Upload of Management Data

Figure 3–4 shows a typical scenario where a set of Management Agents upload their data to a server load balancer, which redirects the data to one of two Management Service installations.

Figure 3–4 Load Balancing Between the Management Agent and the Management Service



In this example, only the upload of Management Agent data is routed through the server load balancer. The Grid Control Console still connects directly to a single Management Service by way of the unique Management Service upload URL.

When you load balance the upload of Management Agent data to multiple Management Service installations, the data is directed along the following paths:

1. Administrators can use one of two URLs to access the Grid Control Console just as they do in the multiple Management Service configuration defined in [Section 3.5](#).

2. Each Management Agent uploads its data to a common server 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 server load balancer. The server load balancer routes the request to any one of a number of available servers that provide the requested service.
3. 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 amongst themselves and one of them loads the data in a background thread in the correct chronological order.
4. 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 3.5](#).
5. 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 3.5](#).

3.6.1.3 Configuring a Server Load Balancer for Management Agent Data Upload

This section describes some guidelines for configuring a server load balancer to balance the upload of data from Management Agents to multiple Management Service installations.

Specifically, you should use the administration tools that are packaged with your server load balancer to configure a virtual pool that consists of the hosts and the services that each host provides. In the case of the Management Services pool, specify the host name and agent upload port. To insure a highly available Management Service, you should have two or more Management Services defined within the virtual pool.

Modify the `REPOSITORY_URL` property in the `emd.properties` file located in the `sysman/config` directory of the Management Agent home directory. The host name and port specified must be that of the server load balancer virtual service.

See Also: ["Configuring the Management Agent to Use a New Management Service"](#) on page 10-1 for more information about modifying the `REPOSITORY_URL` property for a Management Agent

Declare the pool to use a load balancing policy, for example, Round Robin or Least Loaded. Do not configure persistence between Management Agents and Management Services.

This configuration allows the load balancer to distribute connections from Management Agents equally between Management Services. In the event a Management Service becomes unavailable, the load balancer should be configured to direct connections to the surviving Management Services.

To successfully implement this configuration, the load balancer can be configured to monitor the underlying Management Service. On some models, for example, you can configure a *monitor* on the server load balancer. The monitor defines the:

- HTTP request that is to be sent to a Management Service
- Expected result in the event of success
- Frequency of evaluation

For example, the load balancer can be configured to check the state of the Management Service every 5 seconds. On three successive failures, the load balancer can then mark

the component as unavailable and no longer route requests to it. The monitor should be configured to send the string `GET /em/upload` over HTTP and expect to get the response `Http XML File receiver`.

Note: The network bandwidth requirements on the Server Load Balancer need to be reviewed carefully. Monitor the traffic being handled by the load balancer using the administrative tools packaged with your load balancer. Ensure that the load balancer is capable of handling the traffic passing through it. For example, deployments with a large number of targets can easily exhaust a 100 Mbps Ethernet card. A Gigabit Ethernet card would be required in such cases.

See Also: Your Server Load Balancer documentation for more information on configuring virtual pools, load balancing policies, and monitoring network traffic

3.6.2 Load Balancing Connections Between the Grid Control Console and the Management Service


Using a server load balancer to manage the flow of data from the Management Agents is not the only way in which a load balancer can help you configure a highly available Grid Control environment. You can also use a load balancer to balance the load and to provide a failover solution for the Grid Control Console.

The following sections provide more information about this configuration:

- [Understanding the Flow of Data When Load Balancing the Grid Control Console](#)
- [Configuring a Server Load Balancer for the Grid Control Console](#)
- [Configuring Oracle HTTP Server When Using a Server Load Balancer for the Grid Control Console](#)

3.6.2.1 Understanding the Flow of Data When Load Balancing the Grid Control Console

[Figure 3–5](#) shows a typical configuration where a server load balancer is used between the Management Agents and multiple Management Services, as well as between the Grid Control Console and multiple Management Services.

Figure 3–5 Load Balancing Between the Grid Control Console and the Management Service

In this example, a single server load balancer is used for the upload of data from the Management Agents and for the connections between the Grid Control Console and the Management Service.

When you use a server load balancer for the Grid Control Console, the management data uses the following paths through the Grid Control architecture:

1. Administrators use one URL to access the Grid Control Console. This URL directs the browser to the server load balancer virtual service. The virtual service redirects the browser to one of the Management Service installations. Depending upon the host name and port selected by the server load balancer from the virtual pool of Management Service installations, the Grid Control Console obtains the management data by way of OracleAS Web Cache and the Oracle HTTP Server on one of the Management Service hosts.
2. Each Management Agent uploads its data to a common server load balancer URL (as described in [Section 3.6.1](#)) and data is written to the shared receive directory.
3. 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 3.5](#).
4. 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 3.5](#).

3.6.2.2 Configuring a Server Load Balancer for the Grid Control Console

Use the administration tools that are packaged with your server load balancer to configure a virtual pool that consists of the hosts and the services that each host provides. In the case of the Management Services pool, specify the host name and default OracleAS Web Cache port. To insure a highly available Management Service, you should have two or more Management Services defined within the virtual pool.

The load balancer parcels the work to any number of Management Service processes that it has in its virtual pool. This provides a method for constant communication to the Grid Control Console in the event of the failure of a Management Service.

The virtual pool for Grid Control Console should to be configured for session persistence. It is necessary that all requests from one user go to the same Management Service for the duration of a session. Use the persistence method provided by your load balancer. For example if you have enabled Enterprise Manager Framework Security and you are running the Management Service in a secure environment (using HTTPS and SSL), use SSL Session ID based persistence. If you have not enabled Enterprise Manager Framework Security and you are running in an environment that is not secure (using HTTP), you could use Client IP or Cookie based persistence.

3.6.2.3 Configuring Oracle HTTP Server When Using a Server Load Balancer for the Grid Control Console

The Management Service is implemented as a J2EE Web application, which is deployed on an instance of Oracle Application Server. Like many Web-based applications, the Management Service often redirects the client browser to a specific set of HTML pages, such as a logon screen and a specific application component or feature.

When the Oracle HTTP Server redirects a URL, it sends the URL, including the Oracle HTTP Server host name, back to the client browser. The browser then uses that URL, which includes the Oracle HTTP Server host name, to reconnect to the Oracle HTTP Server. As a result, the client browser attempts to connect directly to the Management Service host and bypasses the server load balancer.

To prevent the browser from bypassing the load balancer when a URL is redirected, edit the `ServerName` directive defined in the Oracle HTTP Server configuration file. This directive will be found in one of two places:

- If you have enabled Enterprise Manager Framework Security and you are running the Management Service in a secure environment (using HTTPS and SSL), the `ServerName` directive you must change is located in the following configuration file:

```
ORACLE_HOME/Apache/Apache/conf/ssl.conf
```

- If you have not enabled Enterprise Manager Framework Security and you are running in an environment that is not secure (using HTTP), the `ServerName` directive you must change is located in the following configuration file:

```
ORACLE_HOME/Apache/Apache/conf/httpd.conf
```

Change the `ServerName` directive so it matches the name of the server load balancer virtual service that you configured in [Section 3.6.2.2](#).

See Also: *Oracle HTTP Server Administrator's Guide*

3.6.3 Configuring the Management Repository for High Availability

When you configure Grid Control for high availability, there are several ways to configure the Management Repository to prevent the loss of management data stored in the database.


The following sections describe a typical configuration designed to safeguard your Management Repository:

- Understanding the Flow of Data When Configuring the Management Repository for High Availability
- Installing the Management Repository into a Real Applications Clusters (RAC) Database
- Specifying the Size of the Management Repository Tablespaces in a RAC Database
- Configuring the Management Service to Use Oracle Net Load Balancing and Failover

3.6.3.1 Understanding the Flow of Data When Configuring the Management Repository for High Availability

Figure 3–6 shows a typical Grid Control high availability configuration, where server load balancers are balancing the load on the multiple Management Service installations and the Management Repository is protected by Oracle Real Application Clusters and Oracle Data Guard.

Figure 3–6 Grid Control High Availability Configuration



When you install the Management Repository in a RAC database and incorporate Oracle Data Guard into the configuration, the management data uses the following paths through the Grid Control architecture:

1. Administrators use one URL to access the Grid Control Console. This URL directs the browser to the server load balancer virtual service as described in Section 3.6.2.
2. Each Management Agent uploads its data to a common server load balancer URL as described in Section 3.6.1.

Caution: Before deploying a server load balancer for the upload of Management Agent data, be sure to review [Section 3.6.1.3, "Configuring a Server Load Balancer for Management Agent Data Upload"](#)

- Each Management Service communicates by way of JDBC with a common Management Repository, which is installed in a Real Application Clusters instance. Each Management Service uses the same database connection information, defined in the `emoms.properties` file, to load data 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.

See Also: ["Configuring the Management Service to Use Oracle Net Load Balancing and Failover"](#) on page 3-18 for information about configuring the connection to a Management Repository that is installed in a RAC database

In addition, the Management Repository is also protected by Oracle Data Guard. Note that only physical Data Guard is supported for protecting the Management Repository.

See Also: *Oracle Data Guard Concepts and Administration*

- 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 3.5](#).

See Also: For information about Maximum Availability Architecture (MAA) refer to *Oracle Application Server 10g High Availability Guide*

3.6.3.2 Installing the Management Repository into a Real Applications Clusters (RAC) Database

To install the Management Repository into a RAC database, use the following procedure:

- Install the Oracle 10g Database Release 2 (10.2) software and create a RAC database.
- Begin installing Grid Control, using the **Enterprise Manager 10g Grid Control Using an Existing Database** installation option.
- When you are prompted for a database system identifier (SID) and port, specify the `SERVICE_NAME` for one of the RAC instances.
- After the Grid Control installation is complete, modify the Management Service connection string to take advantage of client failover in the event of a RAC host outage.

See Also: ["Configuring the Management Service to Use Oracle Net Load Balancing and Failover"](#) on page 3-18

- The Management Services rely on repository database listener's connect time load balancing to distribute connections between RAC instances. For the distribution to work optimally in Enterprise Manager Grid Control, ensure that `PREFER LEAST`

LOADED NODE <listener_name> property in listener.ora files is commented out or set to ON.

3.6.3.3 Specifying the Size of the Management Repository Tablespaces in a RAC Database

When you install the Management Repository into a RAC database instance, you cannot set the size of the required Enterprise Manager tablespaces. You can, however, specify the name and location of data files to be used by the Management Repository schema. The default sizes for the initial data file extents depend on using the AUTOEXTEND feature and as such are insufficient for a production installation. This is particularly problematic when storage for the RAC database is on a raw device.

If the RAC database being used for the Management Repository is configured with raw devices, there are two options for increasing the size of the repository.

- You can create multiple raw partitions, with the first one equal to the default size of the tablespace as defined by the installation process.
- Alternatively, you can create the tablespace using the default size, create a dummy object that will increase the size of the tablespace to the end of the raw partition, then drop that object.

Regardless, if raw devices are used, disable the default space management for these objects, which is to auto-extend.

An alternative for RAC installations is to use ASM managed storage. The location string will have to be specified manually (for example, +DATA/<database_name>/<datafile_name>). If ASM storage is used, there is no need to disable any space management storage settings.

3.6.3.4 Configuring the Management Service to Use Oracle Net Load Balancing and Failover

When you use a RAC cluster, a standby system, or both to provide high availability for the Management Repository, the Management Service can be configured to use an Oracle Net connect string that will take advantage of redundancy in the repository. Correctly configured, the Management Service process will continue to process data from Management Agents even during a database node outage.

To configure the Management Service to take advantage of this feature:

1. Use a text editor to open the following configuration file in the Management Service home directory:

```
ORACLE_HOME/sysman/config/emoms.properties
```

2. Locate the following entry in the `emoms.properties` file:

```
oracle.sysman.eml.mntr.emdRepConnectDescriptor=
```

3. Edit the entry so it includes references to the individual nodes within the RAC database.

The following example shows a connect string that supports a two-node RAC configuration. Note the backslash (\) before each equal sign (=), which is required when you are entering the connect string within the `emoms.properties` configuration file:

```
oracle.sysman.eml.mntr.emdRepConnectDescriptor=(DESCRIPTION\=(ADDRESS_
LIST\=(FAILOVER\=ON) (ADDRESS\=(PROTOCOL\=TCP) (HOST\=haem1.us.oracle.com) (PORT\
1521)) (ADDRESS\=(PROTOCOL\=TCP) (HOST\=haem2.us.oracle.com) (PORT\=1521)))
```

```
(CONNECT_DATA=(SERVICE_NAME=em10))
```

See Also: "Enabling Advanced Features of Oracle Net Services" in the *Oracle Database Net Services Administrator's Guide* for more information about using the `FAILOVER` parameter and other advanced features within a database connect string

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](#)
- [Configuring Security for Grid Control](#)
- [Configuring Enterprise Manager for Use with Oracle Application Server Single Sign-On](#)
- [Configuring Enterprise Manager for Use with Enterprise User Security](#)
- [Setting Up the Auditing System for Enterprise Manager](#)
- [Configuring the emkey](#)
- [Additional Security Considerations](#)
- [Other Security Features](#)

4.1 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 following sections describe the security features provided by Enterprise Manager.

4.1.1 Oracle Enterprise Manager Security Model

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. This includes access to the Oracle Enterprise Manager 10g Grid Control Console and the Oracle Enterprise Manager 10g Application Server Control Console.

- 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.

See Also: ["About Enterprise Manager Framework Security"](#) on page 4-4

4.1.2 Classes of Users and Their Privileges

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 10g. For example:

- The Grid Control Console provides support for creating and managing Enterprise Manager administrator accounts.

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.

- Oracle Application Server administrators use the Oracle Application Server administrator account (`ias_admin`) to log in to the Application Server Control Console.
- You use the `ias_admin` account to manage the components of a specific Oracle Application Server instance. You define the password for the `ias_admin` account during the Oracle Application Server installation procedure.

4.1.3 Resources Protected

By restricting access to privileged users and providing tools to secure communications between Oracle Enterprise Manager 10g components, Enterprise Manager protects critical information in the Oracle Management Repository.

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

The Management Repository also contains information about the Enterprise Manager administrators who have the privileges to access the management data.

4.1.4 Authorization and Access Enforcement

Authorization and access enforcement for Enterprise Manager is controlled as follows:

- When you use the Grid Control Console, you create and manage Enterprise Manager administrator accounts. The SYSMAN super administrator can assign specific privileges and roles to each of the additional administrators. These privileges and roles control the targets an administrator can manage and the specific types of tasks the administrator can perform.

See Also: ["About Administrators and Roles"](#) in the Enterprise Manager online Help

- When you use the Application Server Control Console, access to the Console is restricted to administrators who use the `ias_admin` administrator's account. The `ias_admin` account is set up automatically and you assign a password for the account during the Oracle Application Server installation procedure.

See Also: *Oracle Application Server 10g Administrator's Guide* for more information about the `ias_admin` account

See Also: "About Administrators and Roles" in the Enterprise Manager online Help

4.1.5 Leveraging Oracle Application Server Security Services

As a Web-based application, Enterprise Manager relies on industry-standard technologies to provide secure access to the Oracle Enterprise Manager 10g Grid Control Console and Application Server Control Console.

When you configure security for the Oracle Enterprise Manager 10g Grid Control Console, Enterprise Manager Framework Security provides secure communications between the components of your Enterprise Manager installation. However, you should also use the security services of your Oracle HTTP Server to be sure access to the Grid Control Console is secure.

See Also: ["Configuring Security for Grid Control"](#) on page 4-4 for more information about the Enterprise Manager Framework Security

Oracle HTTP Server Administrator's Guide for information about configuring security for your Oracle HTTP Server

Enterprise Manager deploys the Application Server Control Console and Database Control within a single, standalone Oracle Application Server Containers for J2EE (OC4J) instance. As a result, when you configure security for the Application Server Control Console, or for the Database Control, Enterprise Manager uses the standard security services of OC4J to protect your management data.

4.1.6 Leveraging Oracle Identity Management Infrastructure

Oracle Enterprise Manager 10g takes advantage of Oracle Identity Management in two ways:

- First, you can configure the Grid Control Console so it uses Oracle Application Server Single Sign-On. Administrators can then use their Single Sign-On credentials to log in to the Grid Control Console.

See Also: *Oracle Application Server Single Sign-On Administrator's Guide* for general information about Oracle Application Server Single Sign-On

["Configuring Enterprise Manager for Use with Oracle Application Server Single Sign-On"](#) on page 4-18

- Second, you can take advantage of the Enterprise User Security features of the Oracle database. Enterprise User Security provides single sign-on (SSO) or single password authentication for your database users.

See Also: "Managing Enterprise User Security" in the *Oracle Advanced Security Administrator's Guide*

"Configuring Enterprise Manager for Use with Enterprise User Security" on page 4-23

Note: You can configure Enterprise Manager to either use Oracle Application Server Single Sign-On or the Enterprise User Security features. You cannot use both options at the same time.

4.2 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 for the Management Repository Database](#)
- [Enabling Security with a Server Load Balancer](#)

4.2.1 About Enterprise Manager Framework Security


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

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

Enterprise Manager Framework Security works in concert with—but does not replace—the security features you should enable for your Oracle HTTP Server. Oracle HTTP Server is part of the Oracle Application Server instance that is used to deploy the Management Service J2EE Web application.

See Also: *Oracle Application Server 10g Security Guide*

Figure 4–1 shows how Enterprise Manager Framework Security provides security for the connections between the Enterprise Manager components. However, the secure HTTPS connection between your browser and the Grid Control Console should be configured like any other Web application by using the security features of your Oracle HTTP Server.

Figure 4–1 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*

4.2.2 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:

1. Use the `opmnctl stopall` command to stop the Management Service, the Oracle HTTP Server, and the other components of the Oracle Application Server that are used to deploy the Management Service.
2. Use `emctl secure oms` to enable security for the Management Service.
3. Restart the Management Service, the Oracle HTTP Server, OracleAS Web Cache, and the other application server components using the `opmnctl startall` 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 all data gathered from the Management Agents is uploaded over a secure HTTPS connection.

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

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

4.2.3 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 Oracle HTTP Server to enable an HTTPS channel between your Management Service and Management Agents, independent from any existing HTTPS configuration that may be present in your Oracle HTTP Server.
- 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. Change directory to the following directory in the Management Service home:

```
ORACLE_HOME/opmn/bin
```

2. Stop the Management Service, the Oracle HTTP Server, and the other application server components using the following command:

```
$PROMPT> ./opmnctl stopall
```

3. Change directory to the following directory in the Management Service home:

```
ORACLE_HOME/bin
```

4. Enter the following command:

```
$PROMPT> ./emctl secure oms
```

Enterprise Manager prompts you for the Enterprise Manager Root Password.

5. Enter the password for the SYSMAN administrator account used for the Management Repository.

Enterprise Manager prompts you to specify an Agent Registration Password, which is a new password that will be required for any Management Agents that attempt to connect to the Management Service.

6. Specify an Agent Registration Password for the Management Service.

Enterprise Manager prompts you to confirm the host name of the Management Service.

7. When the operation is complete, restart the Management Service, the Oracle HTTP Server, and OracleAS Web Cache:

```
$PROMPT> cd $ORACLE_HOME/opmn/bin
$PROMPT> ./opmnctl startall
```

8. 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:4888/
```

For example:

```
https://mgmthost1.acme.com:4888/
```

If the Management Service security has been enabled, your browser displays the Oracle Application Server Welcome 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.

See Also: ["Viewing a Summary of the Ports Assigned During the Application Server Installation"](#) on page 5-10

Caution: While the `emctl secure oms` command provides immediate HTTPS browser access to the Grid Control Console by using the secure Management Agent upload port, it does not enable security for the default OracleAS Web Cache or Oracle HTTP Server ports that your administrators use to display the Grid Control Console.

To enable security for users who access the Grid Control through OracleAS Web Cache and the default Oracle HTTP Server ports, refer to *Oracle Application Server 10g Security Guide*.

Example 4-1 Sample Output of the `emctl secure oms` Command

```
$PROMPT> ./emctl secure oms
Oracle Enterprise Manager 10g Release 10.2.0.0.0 Copyright (c) 1996, 2005 Oracle
Corporation. All rights reserved.
Enter Enterprise Manager Root Password :
Enter Agent Registration password :
OPMN processes already stopped... Done.
Securing central oms... Started.
Checking Repository... Done.
Checking Em Key... Done.
Checking Repository for an existing Enterprise Manager Root Key... Done.
```

```

Fetching Root Certificate from the Repository... Done.
Generating Registration Password Verifier in the Repository... Done.
Generating Oracle Wallet Password for Enterprise Manager OMS... Done.
Generating Oracle Wallet for Enterprise Manager OMS... Done.
Generating Oracle Wallet for iAS HTTP Server... Done.
Updating HTTPS port in emoms.properties file... Done.
Generating HTTPS Virtual Host for Enterprise Manager OMS... Done.
Securing central oms... Ended.

```

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 4-2 Sample Output of the `emctl secure oms` Command (II)

```

$PROMPT> emctl secure oms -sysman_pwd <sysman password> -reg_pwd <registration
password>[-host <hostname>][-reset][-secure_port <secure_port>][-root_dc <root_
dc>][-root_country <root_country>][-root_state <root_state>][-root_loc <root_
loc>][-root_org <root_org>][-root_unit <root_unit>][-root_email <root_email>]

```

The parameters are explained below:

- `sysman_password` - Oracle Management Repository user password.
- `registration_password` - The Management Agent registration password.
- `hostname` - 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` - If the Oracle Management Service is secured with this option, a new root certificate is generated. All the agents and the Oracle Management Services need to be resecured for use with the new root certificate.
- `secure_port` - The port to be used for secure communication. The default value is **4888**.
- `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>**.

4.2.3.1 Checking the Security Status

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

Example 4-3 Sample Output of the `emctl secure status oms` Command

```

$prompt> emctl secure status oms
Oracle Enterprise Manager 10g Release 10.2.0.0.0 Copyright (c) 1996, 2005 Oracle
Corporation. All rights reserved.
Checking the security status of the OMS at location set in /ade/rpinnama_emcore_
main3/oracle/sysman/config/emoms.properties... Done.
OMS is secure on HTTPS Port 4888

```

4.2.4 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, 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:

```

$PROMPT> ./emctl stop agent

```

4. Enter the following command:

```

$PROMPT> ./emctl secure agent (UNIX)
$PROMPT> 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:

```
$PROMPT> ./emctl secure agent agent_registration_pwd (UNIX)
$PROMPT> emctl secure agent agent_registration_pwd (Windows)
```

[Example 4-4](#) shows sample output of the `emctl secure agent` utility.

5. Restart the Management Agent:

```
$PROMPT> ./emctl start agent
```

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

In the General section of the Management Agent home page ([Figure 4-2](#)), 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 4-4 Sample Output of the `emctl secure agent` Utility

```
$PROMPT> ./emctl secure agent
Oracle Enterprise Manager 10g Release 10.2.0.0.0. Copyright (c) 1996, 2005 Oracle
Corporation. All rights reserved.
Enter Agent Registration password :
Agent is already stopped... Done.
Securing agent... Started.
Requesting an HTTPS Upload URL from the OMS... Done.
Requesting an Oracle Wallet and Agent Key from the OMS... Done.
Check if HTTPS Upload URL is accessible from the agent... Done.
Configuring Agent for HTTPS in CENTRAL_AGENT mode... Done.
EMD_URL set in /private/oracle/agent/sysman/config/emd.properties
Securing agent... Successful.
```

Example 4-5 Sample Output of the `emctl secure status agent` Command

```
[oracle@stang14 bin]$ ./emctl secure status agent
Oracle Enterprise Manager 10g Release 10.2.0.0.0.
Copyright (c) 1996, 2005 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 4-2 Secure Upload Field on the Management Agent Home Page

General	
	Status Up
	Host usunnaa05.us.oracle.com
	Management Service usunna08.us.oracle.com:4888
	Secure Upload Yes
	Version 4.1.0.1.0
	Oracle Home /private/oracle/AGENT_SH5
	Data Pending Upload (MB) 0.000000
	Last Successful Upload Jan 14, 2003 12:53:38 PM

4.2.5 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. For more information on multiple Management Service installations, refer to [Using Multiple Management Service Installations](#) on page 3-6.

All the registration passwords assigned to the current Management Repository are listed on the Registration Passwords page in the Oracle Enterprise Manager 10g Grid Control Console.

See Also: ["Managing Agent Registration Passwords"](#) on page 4-13

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.

4.2.6 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. Any Oracle Management Agent can access the Grid Control Console and Management Service using HTTP or HTTPS connections.

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.

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

1. Stop the Management Service, the Oracle HTTP Server, and the other application server components:

```
$PROMPT> cd $ORACLE_HOME/opmn/bin
```

```
$PROMPT> ./opmnctl stopall
```

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:

```
$PROMPT> emctl secure lock
```

4. Restart the Management Service, the Oracle HTTP Server, and the other application server components:

```
$PROMPT> cd $ORACLE_HOME/opmn/bin
```

```
$PROMPT> ./opmnctl startall
```

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 don't have permission to access /em/upload on this server
```

6. Verify that you can access the Management Agent 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 remove the restriction for HTTPS uploads from the Management Agents, repeat the preceding procedure, but replace the `emctl secure lock` command with the following command:

```
$PROMPT> emctl secure unlock
```

Caution: The `emctl secure lock` command does not prevent users from accessing the Oracle Enterprise Manager 10g Grid Control Console over HTTP. It restricts non-secure access only for Management Agents that attempt to upload data to the Management Service using the upload URL, which is usually:

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

Example 4-6 Sample Output of the `emctl secure lock` Command

```
$prompt> emctl secure lock
```

```
Oracle Enterprise Manager 10g Release 10.2.0.0.0 Copyright (c) 1996, 2005 Oracle Corporation. All rights reserved.
```

```
Checking the security status of the OMS... Done.
```

```
Updating HTTPS Virtual Host for Enterprise Manager OMS... Done.
```


OMS Locked. Agents must be Secure and upload over HTTPS Port 4888.

Example 4-7 Sample Output of the `emctl secure unlock` Command

```
$prompt> emctl secure unlock
Oracle Enterprise Manager 10g Release 10.2.0.0.0 Copyright (c) 1996, 2005 Oracle
Corporation. All rights reserved.
Checking the security status of the OMS... Done.
Updating HTTPS Virtual Host for Enterprise Manager OMS... Done.
OMS Unlocked. Non Secure Agents may upload using HTTP.
```

To restrict HTTP access to the Oracle Enterprise Manager 10g Grid Control Console, configure your Oracle HTTP Server and OracleAS Web Cache as described in the Oracle Application Server documentation.

See Also: *Oracle HTTP Server Administrator's Guide*

4.2.7 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 proper Oracle Management Service.

You create the registration password when you use `emctl secure oms` to configure security for the Oracle Management Service installation.


4.2.7.1 Using the Grid Control Console to Manage Agent Registration Passwords

After you enable security for your Enterprise Manager components, 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 4-3). After you enable security for the Management Service, 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 4–3 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.

4.2.7.2 Using `emctl` to Change the Agent Registration Password

To change an existing Agent Registration Password, use the following `emctl` command:

```
$PROMPT> emctl secure setpwd sysman_password new_Install_Password
```

Note that the `emctl secure setpwd` command requires that you provide the password of the Enterprise Manager super administrator user, `sysman`, to authorize the resetting 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.

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

```
$PROMPT> emctl secure oms -host
```

Enterprise Manager prompts you for the Agent Registration Password and then for the preferred host name ([Example 4-1](#)). Enter the host name for the Server Load Balancer.

4.2.9 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](#)

4.2.9.1 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"](#) on page 4-18

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*.

4.2.9.2 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:

```
$PROMPT> ORACLE_HOME/bin/emctl stop oms
```
2. Locate the following configuration file in the Management Service home directory:

```
ORACLE_HOME/sysman/config/emoms.properties
```
3. Using a text editor, add the entries described in [Table 4–1](#) to the `emoms.properties` file.

The entries described in the table correspond to valid parameters you can set when you configure network data encryption for the Oracle Database.

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

4. Save your changes and exit the text editor.
5. Restart the Management Service.

```
$PROMPT> ORACLE_HOME/bin/emctl start oms
```

See Also: "[Starting and Stopping Oracle Enterprise Manager 10g Grid Control](#)" on page 2-10

Table 4–1 Oracle Advanced Security Properties in the Enterprise Manager Properties File

Property	Description
<code>oracle.sysman.emRep.dbConn.enableEncryption</code>	<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 FALSE.</p> <p>For example:</p> <pre>oracle.sysman.emRep.dbConn. enableEncryption=true</pre>

Table 4–1 (Cont.) Oracle Advanced Security Properties in the Enterprise Manager Properties File

Property	Description
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>

4.2.9.3 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 10g 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.

4.2.9.4 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"
```

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.
4. Restart the Management Agent.

See Also: ["Controlling the Oracle Management Agent"](#) on page 2-1

4.3 Configuring Enterprise Manager for Use with Oracle Application Server Single Sign-On

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 authorize 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 to either use Oracle Application Server Single Sign-On or the Enterprise User Security features. You cannot use both options at the same time.

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

- [Configuring Enterprise Manager to Use the Single Sign-On Logon Page](#)
- [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](#)

4.3.1 Configuring Enterprise Manager to Use the Single Sign-On Logon Page

To configure the Grid Control Console for use with Oracle Application Server Single Sign-On:

1. Set the ORACLE_HOME environment variables to the Management Service home directory.

For example:

```
$PROMPT> setenv ORACLE_HOME /dev01/oracle/em10g_GridControl
```

2. Change directory to the bin directory of the Management Service Oracle home:

```
$PROMPT> cd $ORACLE_HOME/opmn/bin
```

3. Stop the Management Service, the Oracle HTTP Server, and the other components of the application server:

```
$PROMPT> ./opmnctl stopall
```

4. Change directory to the bin directory of the Management Service Oracle home:

```
$PROMPT> cd $ORACLE_HOME/bin
```

5. Enter the following command at the operating system prompt:

```
$PROMPT> ./emctl config oms sso -host ssoHost -port ssoPort -sid ssoSid -pass ssoPassword -das http://ssohost:port/
```

For example:

```
$PROMPT> ./emctl config oms sso -host ssoHost1.acme.com -port 1521 -sid asdb -pass Ch22x5xt -das http://ssohost1.acme.com:7777
```

[Table 4–2](#) describes the arguments on the `emctl config oms sso` command line.

[Example 4–8](#) shows the typical output generated by the `emctl config oms sso` command.

- Restart the Management Service, Oracle HTTP Server, and the other application server components:

```
$PROMPT> cd $ORACLE_HOME/opmn/bin
$PROMPT> ./opmnctl startall
```

- Go the Grid Control Console URL.

For example:

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

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

Table 4–2 Arguments for the emctl sso Command

Argument	Description
-host	The name of the host computer where the Oracle Application Server Single Sign-On server resides. Be sure to use the fully-qualified host name.
-port	The port for the Oracle Application Server Single Sign-On database, for example, 1521.
-sid	The system identifier (SID) for the Oracle Application Server Single Sign-On database.
-pass	The password for the Oracle Application Server Single Sign-On schema (orasso). The orasso schema password is randomized when the Oracle Application Server infrastructure is installed. To obtain the password, see "Obtaining the Single Sign-On Schema Password" in the <i>Oracle Application Server Single Sign-On Administrator's Guide</i> .
-das	The URL containing 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

Example 4–8 Sample Output of the emctl config oms sso Command

```
smpstest@stamt03 bin]$ ./emctl config oms sso -host
isunraj29.us.oracle.com -port 1521 -sid orcl -pass W5RB9YD3 -das
http://isunraj29.us.oracle.com:7777 -u oracle

Oracle Enterprise Manager 10g Release 10.2.0.0.0 Copyright (c) 1996,
2005 Oracle Corporation. All rights reserved.
/scratch/smpstest/mm9/oms10g/Apache/Apache/conf/httpd.conf has been modified.
/scratch/smpstest/mm9/oms10g/sysman/config/emoms.properties has been
modified.
Registering to SSO server, please wait...
Parameters passed to SSO registration tool :
param0:-oracle_home_path param1:/scratch/smpstest/mm9/oms10g param2:-host
param3:
isunraj29.us.oracle.com param4:-port param5:1521 param6:-sid param7:orcl
param8:
-schema param9:orasso param10:-pass param11:**** param12:-site_name
param13:stam
t03.us.oracle.com:4889 param14:-success_url
param15:http://stamt03.us.oracle.com
:4889/osso_login_success param16:-logout_url
param17:http://stamt03.us.oracle.co
```



```

m:4889/osso_logout_success param18:-cancel_url
param19:http://stamt03.us.oracle.
com:4889/ param20:-home_url param21:http://stamt03.us.oracle.com:4889/
param22:-
config_mod_osso param23:TRUE param24:-u param25:oracle
param26:-sso_server_versi
on param27:v1.2 -DinstallType=
-DoldOracleHome=
-DoldOHSUser=root
Check /scratch/smpstest/mm9/oms10g/sso/log/ssoreg.log for details of this
registration
SSO registration tool finished successfully.
Done!

```

4.3.2 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:

1. Go the 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 4-4).

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 (Figure 4-5).

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 4–4 Modified Enterprise Manager Logon Page When Configuring SSO


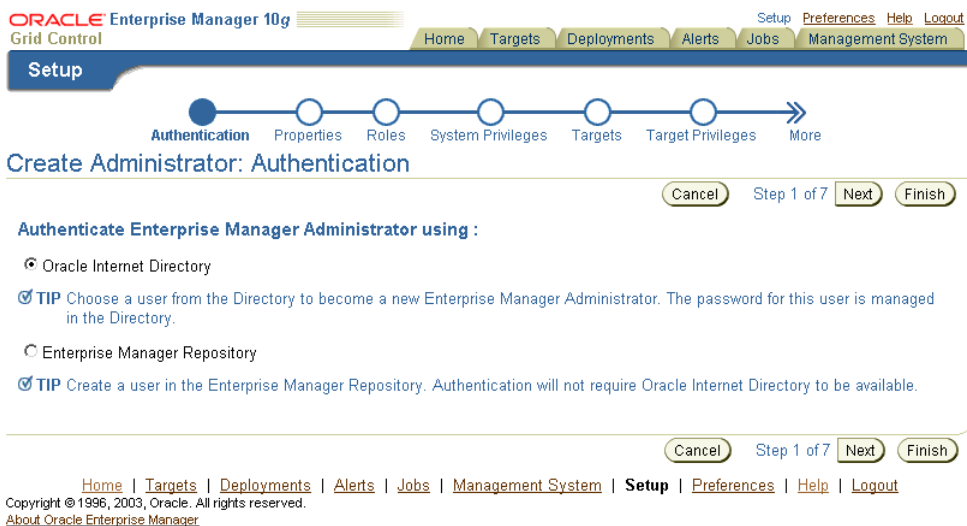


Figure 4–5 Create Administrator Page When SSO Support Is Enabled



4.3.3 Grid Control as a Single Sign-On Partner Application

The `emctl config oms sso` command adds the Oracle Enterprise Manager 10g 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`

4.3.4 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`

4.4 Configuring Enterprise Manager for Use with Enterprise User Security

Enterprise User Security enables you to create and store Oracle9i 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 Oracle9i 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 Oracle9i 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 10g 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.
2. Stop the Oracle Management Service.

See Also: "[Controlling the Oracle Management Service](#)" on page 2-4

3. Change directory to the `IAS_HOME/sysman/config` directory and open the `emoms.properties` file with your favorite text editor.
4. Add the following entry in the `emoms.properties` file:


```
oracle.sysman.emSDK.sec.DirectoryAuthenticationType=EnterpriseUser
```
5. Save and close the `emoms.properties` file.
6. Start the Management Service.

The next time you use the Oracle Enterprise Manager 10g 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.

4.5 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 employ in order 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.

4.5.1 Audit Data

The following data is audited for all Enterprise Manager operations:

Table 4–3 Common Audit Data

Field Name	Description
User Name	The name of the current Enterprise Manager user.
User Type	This can be any of the following: <ul style="list-style-type: none"> ■ Enterprise Manager User ■ Single Sign On User ■ Enterprise User ■ System User
Client Host Name	The name of the user's host machine.
IP Address	The IP address of the user's host machine.
Operation Code	The type of operation. To see a list of operation codes, refer to Operation Codes .
Operation Description	The operation being audited.
Operation Payload	The payload for the selected operation. For example, if the <code>grant_target_priv</code> operation is to be audited, apart from the common data, the <code>user_name</code> , <code>target_name</code> , <code>target_type</code> , and <code>target_owner</code> will be audited.
Object Name	The operation being performed on an object. Each operation has an operation code and an object associated with it. For example, the <code>create_user</code> operation is associated with <code>user_name</code> object, the <code>submit_job</code> operation has a job name associated with it.
Object Type	Each operation code has an object type associated with it. For example, the <code>create_user</code> operation has the <code>user_type</code> object associated with it, the <code>submit_job</code> operation has a job type (OS command, SQL Script) associated with it.
Object Owner	The owner of the object - job owner, operation owner.
Timestamp	The date and time on which the operation took place.
Client Type	The type of browser (UI) or Terminal (Backend).
Client Session	The nature of the session (HTTP Session, DB Session)
OMS Host Name	The host name of the Oracle Management Service.
OMS Time Zone	The time zone of the Oracle Management Service.
Client Session ID	This can be either the HTTP Session ID or the DBMS Session ID.
Login Time	The time at which the user logged into Enterprise Manager.
Logout Time	The time at which the user logged out of Enterprise Manager.
Login Status	The login status indicating whether the login was successful, failed, or timed out.
	Note: The login and logout operations are always audited even if the operation code is turned off.

4.5.2 Operation Codes

Apart from the common audit data, data specific to each operation is also audited. The following table lists the names of operation and their corresponding codes, and additional payloads audited for each operation.

Table 4–4 Operation Specific Data

Operation Name	Operation Code	Additional Payload
change_password	1	user_name
create_user	2	user_name, time_stamp, user_type
delete_user	3	user_name, time_stamp, user_type
logon / logoff	4 and 5	
grant_role	6	user_name
grant_target_priv	7	user_name, target_name, target_type, target_owner
revoke_role	8	user_name
revoke_target_priv	9	user_name, target_name, target_type, target_owner
submit_job	10	
edit_job	11	
delete_job	12	operation_type
modify_user	14	
grant_system_priv	15	user_name
grant_job_priv	16	user_name, job_name, job_type, job_owner
revoke_system_priv	17	user_name
revoke_job_priv	18	user_name, job_name, job_type, job_owner
remote_op	19	step_id, step_status, args, input, remote_command, target_name, target_type, user_name, output
get_file	20	step_id, step_status, dest_file, dest_type, source_file, target_name, target_type, user_name, output
put_file	21	step_id, step_status, dest_file, source_file, source_type, target_name, target_type, user_name, output
file_transfer	22	step_id, step_status, dest_file, dest_target_name, dest_target_type, dest_args, dest_command, dest_user_name, source_file, source_target_name, source_target_type, source_args, source_commands, source_user_name, output
create_role	23	
delete_role	24	
modify_role	25	

4.5.3 Audit APIs

The following APIs allow the administrator (SYSMAN user) to set up the audit function for one or more operations:

- **mgmt_audit_admin.set_audit()** - This API sets the AUDIT_LEVEL and AUDIT_MODE parameters. The AUDIT_LEVEL parameter can be set to 0, 1 or 2 depending on your requirement.

- 0 - All operations in Enterprise Manager will be audited.
- 1 - Only selected operations will be audited. If you select this level, you must turn on the audit function for the operations that are to be audited.
- 2 - None of the operations will be audited.

This API also sets the `AUDIT_MODE` parameter to 0 to store the audited data in the Management Repository.

- `mgmt_audit_admin.set_audit_on()` - This API turns on the audit function for specific operation.
- `mgmt_audit_admin.set_audit_off()` - This API turns off the audit function for a specific operation.

4.5.4 Configuring the Enterprise Manager Audit System

To set up the audit system in Enterprise Manager:

1. Make sure that the Oracle Management Service is up and running.
2. The audit function is turned off by default. Log in to the Enterprise Manager Management Repository as the `sysman` user. To turn on the audit function, enter the following commands:

```
SQL> exec mgmt_audit_admin.set_audit(AUDIT_MODE, null, AUDIT_LEVEL);  
set AUDIT_MODE = 0  
set AUDIT_LEVEL = (0-all, 1-selected, 2-none)
```

3. If the `AUDIT_LEVEL` is set to 1, the audit function needs to be turned on / off for the specific operations that need to be audited by using the following commands:

```
SQL> exec mgmt_audit_admin.set_audit_on(op_code); (Turns on the  
audit function for the specified operation code.)
```

```
SQL> commit;
```

```
SQL> exec mgmt_audit_admin.set_audit_off (op_code); (Turns off  
the audit function for the specified operation code. )
```

```
SQL> commit;
```

For a list of operation codes, refer to [Operation Codes](#) on page 4-25.

4. After setting the `AUDIT_LEVEL`, you must restart the Oracle Management Service to ensure that this change has taken effect.
5. You can then login to Enterprise Manager and create a job or perform other user operations.

Notes:

- Only the SYSMAN user has execute permissions to the audit data. Other users can only view the data in the MGMT\$AUDIT_LOG view if they have the required privileges.

- The audit data can be viewed from MGMT\$AUDIT_LOG view with the following query

```
select * from mgmt$audit_log order by op_code, time_stamp;
```

- The audit data can be purged by using the following command:

```
SQL> exec mgmt_audit_admin.audit_purge(time);
```

The time specified here must be in SYSDATE (DD_MMM_YY) format.

Example: SQL> exec mgmt_audit_admin.audit_purge ('21-SEP-05');

4.6 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 4–9 `emctl start oms` Command

```
$prompt> emctl start oms
Oracle Enterprise Manager 10g Release 10.2.0.0.0
Copyright (c) 1996, 2005 Oracle Corporation. All rights reserved.
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.
```

4.6.1 Generating the emkey

The `emkey` is a random number that is generated during the installation of the Oracle Management Repository and is stored in a table. When the Oracle Management Service is installed, the `emkey` is copied from the Management Repository to the `emkey.ora` file and stored in the `ORACLE_HOME/sysman/config/` directory of each Oracle Management Service.

WARNING: After the **emkey** has been copied, you must remove it from the Management Repository as it is not considered secure. If it is not removed, data such as database passwords, server passwords and other sensitive information can be easily decrypted. To remove the emkey from the Management Repository, enter the following command:

```
$prompt> emctl config emkey - remove_from_repos
```

4.6.2 emctl Commands

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

- `emctl status key`
- `emctl config emkey -repos`
- `emctl config emkey -emkeyfile`
- `emctl config emkey -emkey`
- `emctl config emkey -remove_from_repos`
- `emctl config emkey -copy_to_repos`

The usage of these commands is given below:

```
$prompt> emctl status emkey [-sysman_pwd <sysman password>]
$prompt> emctl config emkey -repos [-emkeyfile <emkey.ora path>] [-force]
[-sysman_pwd <sysman password>]

$prompt> emctl config emkey -emkeyfile <emkey.ora path> [-force] [-sysman_pwd
<sysman password>]
$prompt> emctl config emkey -emkey [-emkeyfile <emkey.ora path>] [-force]
[-sysman_pwd <sysman password>]
$prompt> emctl config emkey -remove_from_repos [-sysman_pwd <sysman password>]
$prompt> emctl config emkey -copy_to_repos [-sysman_pwd <sysman password>]
```

4.6.2.1 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 Management Service but is still present in the Management Repository, the following message is displayed.

Example 4–10 `emctl status emkey` - Example 1

```
Oracle Enterprise Manager 10g Release 10.2.0.0.0
Copyright (c) 1996, 2005 Oracle Corporation. All rights reserved.
The Em Key is configured properly, but is not secure. Secure the Em Key by running
"emctl config emkey -remove_from_repos".
```

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

Example 4–11 `emctl status emkey` - Example 2

```
Oracle Enterprise Manager 10g Release 10.2.0.0.0
Copyright (c) 1996, 2005 Oracle Corporation. All rights reserved.
The Em Key is configured properly.
```


- When the `emkey.ora` file is corrupt or missing and is present in the Management Repository, the following message is displayed.

Example 4–12 `emctl status emkey` - Example 3

```
Oracle Enterprise Manager 10g Release 10.2.0.0.0
Copyright (c) 1996, 2005 Oracle Corporation. All rights reserved.
The Em Key exists in the Management Repository, but is not configured properly or
is corrupted in the file system.
Configure the Em Key by running "emctl config emkey -repos".
```

- When the `emkey.ora` file is corrupt or missing and is not present in the Management Repository, the following message is displayed.

Example 4–13 `emctl status emkey` - Example 4

```
Oracle Enterprise Manager 10g Release 10.2.0.0.0
Copyright (c) 1996, 2005 Oracle Corporation. All rights reserved.
The Em Key is not configured properly or is corrupted in the file system and does
not exist in the Management Repository. To correct the problem:
1) Copy the emkey.ora file from another OMS or backup machine to the
OH/sysman/config directory.
2) Configure the emkey.ora file by running "emctl config emkey -emkeyfile
<emkey.ora file location>".
```

4.6.2.2 `emctl config emkey -repos`

This command copies the `emkey` from the Management Repository to the `emkey.ora` file.

Example 4–14 Sample Output of the `emctl config emkey -repos` Command

```
$ emctl config emkey -repos -emkeyfile /tmp/emkey.ora.0 -force
Oracle Enterprise Manager 10g Release 10.2.0.0.0
Copyright (c) 1996, 2005 Oracle Corporation. All rights reserved.
Please enter repository password:
The Em Key has been configured successfully.
```

In this example, the `emkey` is copied from the Management Repository to the `/tmp/emkey.ora.0` file. The command configures the `oracle.sysman.emkeyfile` property in the `emoms.properties` to point to this file.

Note: The `-force` option is required only if the `emkey` file is already configured.

If the `-emkeyfile` option is not provided in the Management Repository, the `emkey` is overwritten to the already configured `emkey.ora` file.

4.6.2.3 `emctl config emkey -emkeyfile`

This command can be used to configure a new `emkey.ora` file.

Example 4–15 Sample Output of `emctl config emkey -emkeyfile` Command

```
$ emctl config emkey -emkeyfile /tmp/emkey.ora.1 -force
Oracle Enterprise Manager 10g Release 10.2.0.0.0
```

```
Copyright (c) 1996, 2005 Oracle Corporation. All rights reserved.  
Please enter repository password:  
The Em Key has been configured successfully.
```

This command configures the `/tmp/emkey.ora.1` file as the new `emkey.ora` file. It also modifies the `oracle.sysman.emkeyfile` property in `emoms.properties` to point to this file. The `-force` option is required only if the `emkey.ora` file has already been configured.

4.6.2.4 emctl config emkey -emkey

This command is used to configure a new `emkey`.

Example 4–16 Sample Output of emctl config emkey -emkey Command

```
$ emctl config emkey -emkey -emkeyfile /tmp/emkey.ora.2 -force  
Oracle Enterprise Manager 10g Release 10.2.0.0.0  
Copyright (c) 1996, 2005 Oracle Corporation. All rights reserved.  
Please enter repository password:  
Please enter the em key:  
The Em Key has been configured successfully.
```

This command writes the `emkey` provided as standard input into the `/tmp/emkey.ora.2` file and configures it. The `-force` option is required only if the `emkey.ora` file has already been configured. If the `-emkeyfile` option is not provided, the `emkey` is overwritten to the already configured `emkey.ora` file.

4.6.2.5 emctl config emkey -remove_from_repos

This command removes the `emkey` from the Management Repository.

Example 4–17 Sample Output of emctl config emkey -remove_from_repos Command

```
$ emctl config emkey -remove_from_repos  
Oracle Enterprise Manager 10g Release 10.2.0.0.0  
Copyright (c) 1996, 2005 Oracle Corporation. All rights reserved.  
Please enter repository password:  
The Em Key has been removed from the Management Repository.  
Make a backup copy of OH/sysman/config/emkey.ora file and store it on another  
machine.  
WARNING: Encrypted data in Enterprise Manager will become unusable if the  
emkey.ora file is lost or corrupted.
```

4.6.2.6 emctl config emkey -copy_to_repos

This command copies the `emkey` back to the Management Repository.

Example 4–18 Sample Output of emctl config emkey_copy_to_repos Command

```
$ emctl config emkey -copy_to_repos  
Oracle Enterprise Manager 10g Release 10.2.0.0.0  
Copyright (c) 1996, 2005 Oracle Corporation. All rights reserved.  
Please enter repository password:  
The Em Key has been copied to the Management Repository. This operation will cause  
the Em Key to become unsecure.
```

Note: This command is used during the additional Oracle Management Service install (See [Section 4.6.3](#)). When you use this command, the emkey will be present in the Management Repository, which is not considered secure. You can secure it after the additional Oracle Management Service install by running the command:

```
emctl config emkey -remove_from_repos
```

4.6.3 Install and Upgrade Scenarios

This section explains the install and upgrade scenarios for emkey.

4.6.3.1 Installing the Management Repository

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

4.6.3.2 Installing the First Oracle Management Service

When the Oracle Management Service is installed, the installer copies the emkey from the Management Repository and stores it in the `emkey.ora` file.

Note: After installation, the emkey will be present in the Management Repository. This is not considered secure. The user can secure the emkey by running the emctl command `emctl config emkey -remove_from_repos`

4.6.3.3 Installing Additional Oracle Management Service

Similar to the first Oracle Management Service install, the installer will copy the emkey from the Management Repository to the `emkey.ora` file of the additional Oracle Management Service.

Note: After the first Oracle Management Service install, you may have removed the emkey from the Management Repository using the emctl command.

Before the additional Oracle Management Service is installed, run the following command from the first Oracle Management Service home to copy the emkey to the Management Repository.

```
emctl config emkey -copy_to_repos
```

If the additional Oracle Management Service install is done without the emkey in the Management Repository, the installer will prompt the user to run the command mentioned above.

4.6.3.4 Upgrading from 10.1 to 10.2

The Management Repository is upgraded as usual. When the Oracle Management Service is upgraded, the upgrade script copies the emkey from the Management Repository to the `emkey.ora` file of each Oracle Management Service.

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
```

4.6.3.5 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.ora in the Oracle Management Service home directory. Enter the `emctl config emkey -repos -force` command to overwrite the new emkey to the emkey.ora file.

4.7 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:

- [Responding to Browser-Specific Security Certificate Alerts](#)
- [Configuring Beacons to Monitor Web Applications Over HTTPS](#)

4.7.1 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 Oracle HTTP Server properly.

See Also: *Oracle Application Server 10g Security Guide*

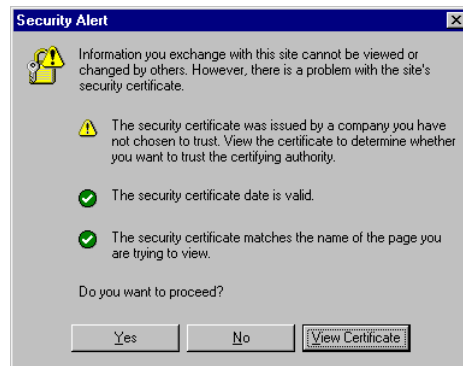
This section contains the following topics:

- [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](#)

4.7.1.1 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 Oracle HTTP Server, you will likely receive a Security Alert dialog box similar to the one shown in [Figure 4–6](#) 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 4-6 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 4-7](#).
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 4–7 Certificate Path Tab on the Internet Explorer Certificate Dialog Box

4.7.1.2 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 Oracle HTTP Server, you will likely receive a New Site Certificate dialog box similar to the one shown in [Figure 4–8](#) 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 4–8 Netscape Navigator New Site Certificate Dialog Box

4.7.1.3 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 4–9](#) 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 4–9 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.

4.7.2 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 *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 4-7](#).
- 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 4-19](#).

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 4-19 Sample Content of an Exported Certificate

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

4.8 Other Security Features

This section describes Enterprise Manager security features.

4.8.1 Using ORACLE_HOME Credentials

Oracle Enterprise Manager 10g Release 2 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 10g Release 2, one can explicitly set the ORACLE_HOME credential 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 (EM CLI) 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. [Table 4–5](#) describes the input values to the `emcli set_credential` command.

Table 4–5 *emcli set_credential Parameters*

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

4.8.2 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.

4.8.3 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.

4.8.4 Using the sudo Command

sudo allows a permitted user to execute a command as the superuser or another user, as specified in the sudoers file. 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 use sudo. 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`.

Configuring Enterprise Manager for Firewalls

Firewalls protect a company's Information Technology (IT) infrastructure by providing the ability to restrict network traffic by examining each network packet and determining the appropriate course of action.

Firewall configuration typically involves restricting the ports that are available to one side of the firewall, for example the Internet. It can also be set up to restrict the type of traffic that can pass through a particular port such as HTTP. If a client attempts to connect to a restricted port (a port not covered by a security "rule") or uses a protocol that is incorrect, then the client will be disconnected immediately by the firewall. Firewalls can also be used within a company Intranet to restrict user access to specific servers.

You can deploy the components of Oracle Enterprise Manager on different hosts throughout your enterprise. These hosts can be separated by firewalls. This chapter describes how firewalls can be configured to allow communication between the Enterprise Manager components.

See Also: [Chapter 3](#) for more information about some of the ways you can configure the Grid Control components on your network

This chapter contains the following topics:

- [Considerations Before Configuring Your Firewall](#)
- [Firewall Configurations for Enterprise Management Components](#)
- [Viewing a Summary of the Ports Assigned During the Application Server Installation](#)

5.1 Considerations Before Configuring Your Firewall

Firewall configuration should be the last phase of Enterprise Manager deployment. Before you configure your firewalls, make sure you are able to log in to the Grid Control Console and that your Management Agents are up and monitoring targets.

If you are deploying Enterprise Manager in an environment where firewalls are already installed, open the default Enterprise Manager communication ports for all traffic until you have completed the installation and configuration processes and are certain that you are able to log in to the Oracle Enterprise Manager 10g Grid Control Console and that your Oracle Management Agents are up and monitoring targets.

The default communication ports for Enterprise Manager are assigned during the installation. If you modify the default ports, be sure to use the new port assignments when you configure the firewalls.

See Also: [Chapter 10, "Reconfiguring the Management Agent and Management Service"](#) for information about locating and changing the default ports for the Oracle Management Service and the Oracle Management Agent

If you are enabling Enterprise Manager Framework Security for the Management Service, the final step in that configuration process is to restrict uploads from the Management Agents to secure channels only. Before completing that step, configure your firewalls to allow both HTTP and HTTPS traffic between the Management Agent and Management Repository and test to be sure that you can log in to Enterprise Manager and that data is being uploaded to the Management Repository.

After you have confirmed that the Management Service and Management Agents can communicate with both protocols enabled, complete the transition to secure mode and change your firewall configuration as necessary. If you incrementally configure your firewalls, it will be easier to troubleshoot any configuration problems.

5.2 Firewall Configurations for Enterprise Management Components

Your main task in enabling Enterprise Manager to work in a firewall-protected environment is to take advantage of proxy servers whenever possible, to make sure only the necessary ports are open for secure communications, and to make sure that only data necessary for running your business is allowed to pass through the firewall.

The following sections describe the ports and types of data required by Enterprise Manager in a secure, firewall-protected environment:

- [Firewalls Between Your Browser and the Grid Control Console](#)
- [Configuring the Management Agent on a Host Protected by a Firewall](#)
- [Configuring the Management Service on a Host Protected by a Firewall](#)
- [Firewalls Between the Management Service and the Management Repository](#)
- [Firewalls Between the Grid Control and a Managed Database Target](#)
- [Firewalls Used with Multiple Management Services](#)
- [Configuring Firewalls to Allow ICMP and UDP Traffic for Beacons](#)
- [Configuring Firewalls When Managing Oracle Application Server](#)

5.2.1 Firewalls Between Your Browser and the Grid Control Console

Connections from your browser to the Oracle Enterprise Manager 10g Grid Control Console are performed over the default port used for your Oracle HTTP Server.

For example, the default, non-secure port for the Oracle HTTP Server is usually port 7777. If you are accessing the Grid Control Console using the following URL and port, then you must configure the firewall to allow the Grid Control Console to receive HTTP traffic over port 7777:

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


On the other hand, if you have enabled security for your Oracle HTTP Server, you are likely using the default secure port for the server, which is usually port 4443. If you are accessing the Grid Control Console using the following URL and port, then you must configure the firewall to allow the Grid Control Console to receive HTTP traffic over port 4443:

<https://mgmthost.acme.com:4443/em>

See also: *Oracle Application Server 10g Security Guide*

Figure 5-1 shows the typical configuration of a firewall between your browser and the Grid Control Console Web-based console that is rendered by the Management Service.

Figure 5-1 Firewall Between Your Browser and the Grid Control Console



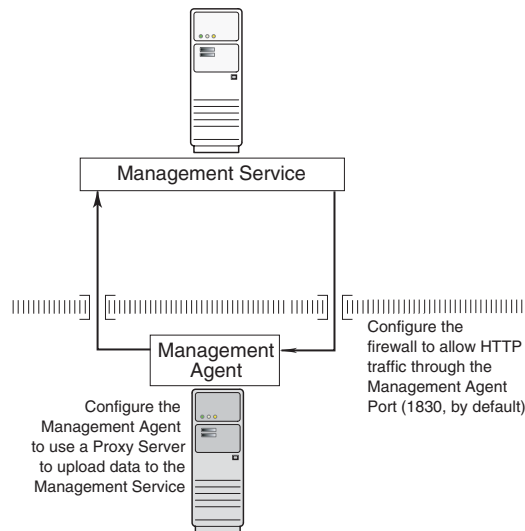
5.2.2 Configuring the Management Agent on a Host Protected by a Firewall

If your Management Agent is installed on a host that is protected by a firewall and the Management Service is on the other side of the firewall, you must perform the following tasks:

- Configure the Management Agent to use a proxy server for its uploads to the Management Service.
- Configure the firewall to allow incoming HTTP traffic from the Management Service service on the Management Agent port, which is 1830 by default, regardless of whether or not Enterprise Manager Framework Security has been enabled.

Figure 5-2 illustrates the connections the Management Agent must make when it is protected by a firewall.

Figure 5-2 Configuration Tasks When the Management Agent is Behind a Firewall



5.2.2.1 Configuring the Management Agent to Use a Proxy Server

You can configure the Management Agent to use a proxy server for its communications with a Management Service outside the firewall, or to manage a target outside the firewall.

1. Use a text editor to open the following Management Agent configuration file:

```
AGENT_HOME/sysman/config/emd.properties (UNIX)
AGENT_HOME\sysman\config\emd.properties (Windows)
```

2. Locate the following entry in the `emd.properties` file:

```
# If it is necessary to go through an http proxy server to get to the
# repository, uncomment the following lines
#REPOSITORY_PROXYHOST=
#REPOSITORY_PROXYPORT=
```

3. To enable support for authenticating the proxy server, the following additional properties need to be specified.

```
#REPOSITORY_PROXYREALM=
#REPOSITORY_PROXYUSER=
#REPOSITORY_PROXYPWD=
```

4. Edit the following properties by removing the pound sign (#) at the start of each line and entering a value as follows:

```
# If it is necessary to go through an http proxy server to get to the
# repository, uncomment the following lines
REPOSITORY_PROXYHOST=proxyhostname.domain
REPOSITORY_PROXYPORT=proxy_port
REPOSITORY_PROXYREALM=realm
REPOSITORY_PROXYUSER=proxyuser
REPOSITORY_PROXYPWD=proxypassword
```

For example:

```
REPOSITORY_PROXYHOST=proxy42.acme.com
REPOSITORY_PROXYPORT=80
REPOSITORY_PROXYREALM=
REPOSITORY_PROXYUSER=
REPOSITORY_PROXYPWD=
```

5. Save your changes and close the `emd.properties` file.
6. Stop and start the Management Agent.

See Also: ["Controlling the Oracle Management Agent"](#) on page 2-1

Note: The proxy password will be rewritten when you restart the Management Agent.

5.2.2.2 Configuring the Firewall to Allow Incoming Communication From the Management Service

While the Management Agents in your environment must upload data from your managed hosts to the Management Service, the Management Service must also communicate with the Management Agents. As a result, if the Management Agent is

protected by a firewall, the Management Service must be able to contact the Management Agent through the firewall on the Management Agent port.

By default, the Enterprise Manager installation procedure assigns port 1830 to the Management Agent. However, if that port is occupied, the installation may assign an alternate port number.

Note: The port number for the Management Agent does not change when you enable Enterprise Manager Framework Security. For more information, see "[Configuring Security for Grid Control](#)" on page 4-4

In addition, administrators can change the Management Agent port after the installation.

See Also: "[Chapter 10, "Reconfiguring the Management Agent and Management Service"](#)" for information about locating and changing the default ports for the Oracle Management Service and the Oracle Management Agent.

After you determine the port number assigned to the Management Agent, you must then configure the firewall to allow incoming HTTP or HTTPS traffic (depending upon whether or not you have enabled Enterprise Manager Framework Security) on that port.

See Also: Your firewall documentation for more information about opening specific ports for HTTP or HTTPS traffic.

["Configuring Security for Grid Control"](#) on page 4-4 for information about Enterprise Manager Framework Security

5.2.3 Configuring the Management Service on a Host Protected by a Firewall

If your Management Service is installed on a host that is protected by a firewall and the Management Agents that provide management data are on the other side of the firewall, you must perform the following tasks:

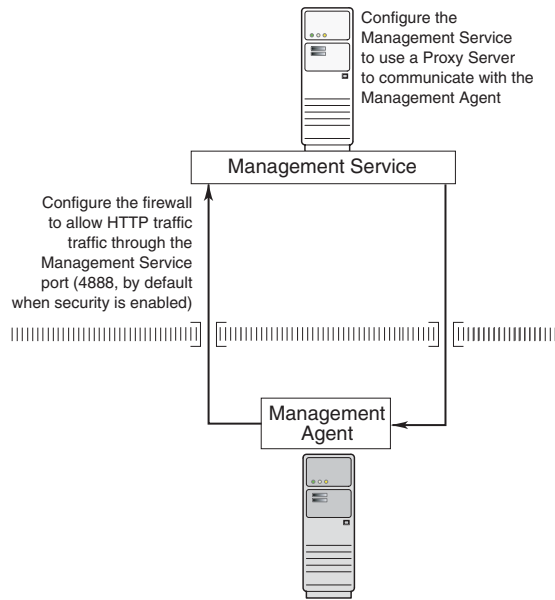
- Configure the Management Service to use a proxy server for its communications to the Management Agents.
- Configure the firewall to allow incoming HTTP traffic from the Management Agents on the Management Repository upload port.

If you have enabled Enterprise Manager Framework Security, the upload URL uses port 4888; if you have *not* enabled Enterprise Manager Framework Security, the upload port is 4889.

See also: "[Enabling Security for the Oracle Management Service](#)" on page 4-6

[Figure 5-3](#) illustrates the connections the Management Service must make when it is protected by a firewall.

Figure 5-3 Configuration Tasks When the Management Service is Behind a Firewall



5.2.3.1 Configuring the Management Service to Use a Proxy Server

This section describes how to configure the Management Service to use a proxy server for its communications with Management Agents outside the firewall.

Note: The proxy configuration properties described in this section are the same Management Service properties you must modify if your network is protected by a firewall and you want Enterprise Manager to search automatically for critical patches and patch sets. For more information, see "Specifying Patching Credentials" in the Enterprise Manager online Help.

To configure the Management Service to use a proxy server:

1. Use a text editor to open the following configuration file in the Management Service home directory:

```
ORACLE_HOME/sysman/config/emoms.properties
```

2. Add the following entries to emoms.properties file:

```
proxyHost=proxyhost.domain
proxyPort=proxy_port
dontProxyFor=.domain1, .domain2, .domain3, ...
proxyRealm=realm
proxyUser=proxyuser
proxyPwd=proxypassword
```

For example:

```
proxyHost=proxy42.acme.com
proxyHost=80
dontProxyFor=.acme.com, .acme.us.com
proxyRealm
proxyUser
proxyPwd
```

The `dontProxyFor` property identifies specific URL domains for which the proxy will not be used. The `proxyRealm` property indicates the protected space that requires authentication.

See Also: "[About the dontProxyfor Property](#)" on page 5-7 for guidelines on when to use the `dontProxyFor` property

3. Save your changes and close the `emoms.properties` file.
4. Stop and start the Management Service:

```
$PROMPT> ORACLE_HOME/bin/emctl stop oms
$PROMPT> ORACLE_HOME/bin/emctl start oms
```

Note: The proxy password will be rewritten when you restart the Management Service.

5.2.3.2 About the dontProxyfor Property

When you configure the Management Service to use a proxy server, it is important to understand the purpose of the `dontProxyFor` property, which identifies specific URL domains for which the proxy will not be used.

For example, suppose the following were true:

- You have installed the Management Service and several Management Agents on hosts that are inside the company firewall. These hosts are in the internal `.acme.com` and `.acme.us.com` domains.
- You have installed several additional Management Agents on hosts that are outside the firewall. These hosts are installed in the `.acme.uk` domain.
- You have configured Enterprise Manager to automatically check for critical software patches on the *OracleMetaLink* Internet site.

In this scenario, you want the Management Service to connect directly to the Management Agents inside the firewall without using the proxy server. On the other hand, you want the Management Service to use the proxy server to contact the Management Agents outside the firewall, as well as the *OracleMetaLink* Internet site, which resides at the following URL:

```
http://metalink.oracle.com
```

The following entry in the `emoms.properties` file will prevent the Management Service from using the proxy server for connections to the Management Agents inside the firewall. Connections to *OracleMetaLink* and to Management Agents outside the firewall will be routed through the proxy server:

```
proxyHost=proxy42.acme.com
proxyHost=80
dontProxyFor=.acme.com, .acme.us.com
```

5.2.3.3 Configuring the Firewall to Allow Incoming Management Data From the Management Agents

While the Management Agents in your environment must contact the Management Agents on your managed hosts, the Management Service must also be able to receive upload data from the Management Agents. If the Management Service is behind a firewall, you must configure the firewall to allow the Management Agents to upload data on the upload port.

By default, the Enterprise Manager installation procedure assigns port 4889 as the Repository upload port. However, if that port is occupied, the installation will assign an alternate port number.

In addition, when you enable Enterprise Manager Framework Security, the upload port is automatically changed to the secure 4888 HTTPS port.

See Also: ["Configuring Security for Grid Control"](#) on page 4-4 for information about Enterprise Manager Framework Security

Administrators can also change the upload port after the installation.

See Also: [Chapter 10, "Reconfiguring the Management Agent and Management Service"](#) for information about locating and changing the default ports for the Oracle Management Service and the Oracle Management Agent.

After you determine the port number assigned to the Management Service upload port, you must then configure the firewall to allow incoming HTTP or HTTPS traffic (depending upon whether or not you have enabled Enterprise Manager Framework Security) on that port.

See Also: Your firewall documentation for more information about opening specific ports for HTTP or HTTPS traffic


5.2.4 Firewalls Between the Management Service and the Management Repository

Secure connections between the Management Service and the Management Repository are performed using features of Oracle Advanced Security. As a result, if the Management Service and the Management Repository are separated by a firewall, you must configure the firewall to allow Oracle Net firewall proxy access.

See Also: ["Configuring Secure Sockets Layer Authentication"](#) in the *Oracle Database Advanced Security Administrator's Guide*

[Figure 5-4](#) shows a typical configuration of a firewall between the Management Service and the Management Repository.

Figure 5-4 Firewall Between the Management Service and the Management Repository



5.2.5 Firewalls Between the Grid Control and a Managed Database Target

When you are using the Grid Control Console to manage a database, you must log in to the database from the Grid Control Console in order to perform certain monitoring and administration tasks. If you are logging in to a database on the other side of a firewall, you will need to configure the firewall to allow Oracle Net firewall proxy access.


Specifically, to perform any administrative activities on the managed database, you must be sure that the firewall is configured to allow the Oracle Management Service to communicate with the database through the Oracle Listener port.

You can obtain the Listener port by reviewing the Listener home page in the Grid Control Console.

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

Figure 5-5 shows a typical configuration of a firewall between Grid Control and the Management Repository.

Figure 5-5 Firewall Between Grid Control and a Managed Database Target



5.2.6 Firewalls Used with Multiple Management Services

Enterprise Manager supports the use of multiple Management Services that communicate with a common Management Repository. For example, using more than one Management Service can be helpful for load balancing as you expand your central management capabilities across a growing e-business enterprise.

When you deploy multiple Management Services in an environment protected by firewalls, be sure to consider the following:

- Each Management Agent is configured to upload data to one Management Service. As a result, if there is a firewall between the Management Agent and its Management Service, you must configure the firewall to allow the Management Agent to upload data to the Management Service using the upload URL.

See Also: ["Configuring the Management Agent on a Host Protected by a Firewall"](#) on page 5-3

["Configuring the Management Service on a Host Protected by a Firewall"](#) on page 5-5

- In addition, each Management Service must be able to contact any Management Agent in your enterprise so it can check for the availability of the Management Agent. As a result, you must be sure that your firewall is configured so that each Management Service you deploy can communicate over HTTP or HTTPS with any Management Agent in your enterprise.

Otherwise, a Management Service without access to a particular Management Agent may report incorrect information about whether or not the Management Agent is up and running.

See Also: ["About Availability"](#) in the Enterprise Manager online Help for information about how Enterprise Manager determines host and Management Agent availability

5.2.7 Configuring Firewalls to Allow ICMP and UDP Traffic for Beacons

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

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

Enterprise Manager uses the industry-standard Internet Control Message Protocol (ICMP) and User Datagram Protocol (UDP) to transfer data between Beacon and the network components you are monitoring. There may be situations where your Web application components and the Beacons you use to monitor those components are separated by a firewall. In those cases, you must configure your firewall to allow ICMP, UDP, and HTTP traffic.

See Also: ["Configuring Beacons to Monitor Web Applications Over HTTPS"](#) on page 4-35

5.2.8 Configuring Firewalls When Managing Oracle Application Server

If you are using Grid Control to manage instances of Oracle Application Server, there may be other ports that you need to access through a firewall, depending upon your configurations.

For example, when you are monitoring the performance of your Oracle Application Server instance from the Grid Control Console, you can click **Administer** on the Application Server Home page to display the Application Server Control Console. If the Oracle Application Server target you are monitoring is separated from the Grid Control Console by a firewall, you will need to configure the firewall to allow an HTTP or HTTPS connection through Application Server Control Console port (usually, 1810).

See Also: *Oracle Application Server Administrator's Guide* for more information about configuring ports for Oracle Application Server

5.3 Viewing a Summary of the Ports Assigned During the Application Server Installation

As described in the previous sections of this chapter, it is important to understand and identify the ports used by each of the Oracle Enterprise Manager 10g components before you configure your firewalls.

When you install the Oracle Application Server 10g or the Oracle Enterprise Manager 10g Grid Control, you can view a list of the ports assigned during the application server installation by viewing the contents of the following file

ORACLE_HOME/install/portlist.ini

Note: The `portlist.ini` file lists the port numbers assigned during the installation. This file is not updated if port numbers are changed after the installation.

In addition, you can use the Application Server Control Console to view a list of all the ports in use by the application server:

1. Navigate to the Application Server home page in the Application Server Control Console.
2. Click **Ports**.

See Also: "Viewing and Modifying Application Server Port Assignments" in the Enterprise Manager online Help

5.4 Checking and Configuring Firewall Settings for HTTP/HTTPS

For secure agent install, ensure that the firewall settings are disabled for HTTP/HTTPS communication for Windows XP:

1. Go to **Start**, and then select **Control Panel**.
2. In Control Panel, click **Windows Firewall**.
3. In the **Exceptions** tab in the **Windows Firewall** dialog box, click **Add Port**.
4. In the **Add a Port** dialog box, specify the name and number of the port.
5. Click **Change scope** to specify the computers for which the port is unblocked.

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 Transactions](#)
- [Monitoring Settings](#)
- [Configuring Aggregate Services](#)
- [Configuring End-User Performance Monitoring](#)
- [Configuring OC4J for Request Performance Diagnostics](#)
- [Setting Up Monitoring Templates](#)
- [Configuring Service Levels](#)
- [Configuring a Service Using the Command Line Interface](#)

6.1 Summary of Service Management Tasks

This table provides a summary list of all the service management features and their requirements.

Table 6–1 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

Table 6–1 (Cont.) Summary of Service Management Tasks

Feature	Description	Requirements	Refer to
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
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

6.2 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.

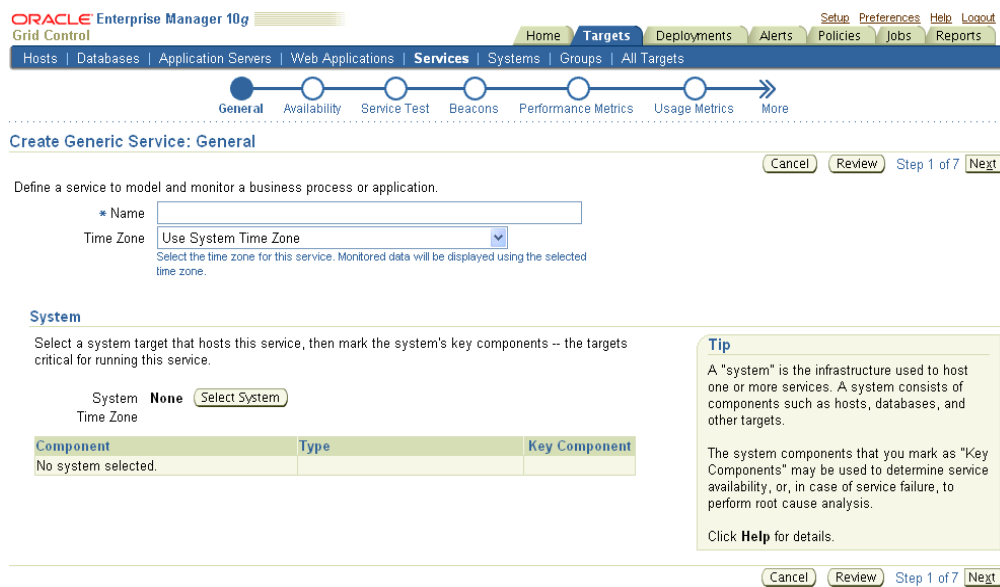
6.3 Creating a Service

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 6–1 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, and Aggregate Service. 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. 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.

6.4 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 6–2 Monitoring Configuration Page

The following options are available:

- [Availability Definition](#)
- [Performance Metrics](#)
- [Usage 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](#)

6.4.1 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 6-10.

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.

6.4.2 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 6–2 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 6–3 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.

6.4.3 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 6–4 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 Web application.
Minimum	The minimum value of the metric across all components will be used as the value of this usage metric for the Web application.
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.

6.4.4 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.

6.4.4.1 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:
 - For Web transactions, refer to instructions in the "[Configuring Beacons to Monitor Web Applications Over HTTPS](#)" on page 4-35.
 - To use the SSL option with the Port Checker test, you may need to add additional certificates to the agent's monitoring wallet. For details on adding certificates, refer to "[Adding Trust Points to the Management Agent Configuration](#)" on page 10-7.
- **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
```

6.4.5 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.

6.4.5.1 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.

6.5 Recording 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.

6.6 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.

6.7 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, and usage for the aggregate service depend on the availability, performance, 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 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 usage metrics to use, you can set the thresholds used 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 and usage metrics may be affected if they are based on a deleted subservice's metrics.

6.8 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](#)
- [Setting Up the Forms Application for End-User Performance Monitoring](#)

6.8.1 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.

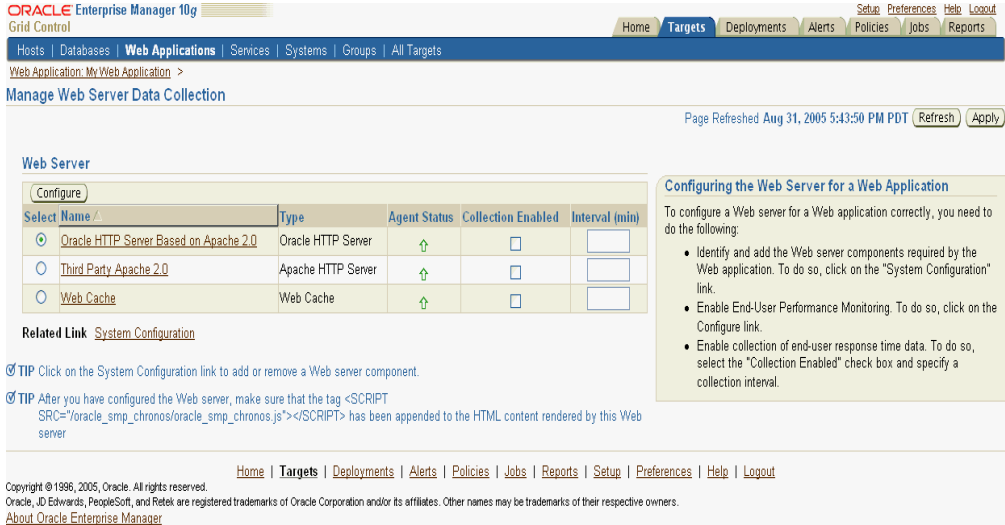
Note: If the Oracle HTTP Server is installed before the Management Agent has been installed, and is up and running during agent installation, then the target will be discovered automatically. Otherwise you need to manually create the Oracle HTTP Server target and specify the following properties: Machine 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 6-13.

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 6-3 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 6-24. 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 6-24.

6.8.1.1 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 Server2.0 server. You can then follow the steps to configure the Apache server and enable collection of end-user performance data.

6.8.2 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](#)

6.8.2.1 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 6-24. If you do not see data after end-user performance has been enabled, refer to "[Verifying and Troubleshooting End-User Performance Monitoring](#)" on page 6-24.

6.8.2.2 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 user name 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 6-24. If you do not see data after end-user performance has been enabled, refer to "[Verifying and Troubleshooting End-User Performance Monitoring](#)" on page 6-24.

6.8.2.3 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.

6.8.2.3.1 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.

6.8.2.3.2 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
```

6.8.2.3.3 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>`
-

6.8.2.3.4 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 2-15.

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

6.8.2.4 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](#)

6.8.2.4.1 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.

6.8.2.4.2 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.

6.8.2.4.3 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 6-24.

6.8.3 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.

6.8.4 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, go to Step g.
 - 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 Step h.
 - 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.

6.8.5 Setting Up the Forms Application for End-User Performance Monitoring

Enterprise Manager allows you to monitor response time data of end-users accessing an Oracle Forms application. The Forms operations that can be monitored include `commit`, `query`, `runform`, `callform`, `newform`, and `openform`. For each of these operations, the total response time, server time and database time for a particular URL are measured.

You can use either of the following Forms application versions:

- Forms 10.1.2 (All Forms operations can be monitored)
- Forms 6 Patch 16 (Only the `commit` operation can be monitored)

End-user performance monitoring for an Oracle Forms application requires an OracleAS Web Cache. OracleAS Web Cache is a component of the Oracle Application Server that runs the Forms application. To enable end-user performance monitoring, the Forms server needs to be configured and the logging format of the Oracle Web Cache needs to be changed to **End-User Performance Monitoring format**.

6.8.5.1 Configuring Forms Server for End-User Performance Monitoring

To set up the Forms application for End-User Performance Monitoring

1. Navigate to the Home page in Application Server Control.
2. Select the Forms system component and click **Configuration**.
3. Ensure the Forms Web Configuration (`formsweb.cfg`) View is selected.
4. Select the appropriate section to enable End-User Performance Monitoring, or click **Create New Section** to enable End-User Performance Monitoring for a new section and enter the name of the new section.
5. Click **Edit** to add or modify the following parameters in the `formsweb.cfg` file. If these parameters have not been defined, click **Add New Parameters** and enter the parameter names and their corresponding values.
6. Set the `endUserMonitoringEnabled` parameter to true.
7. Set the path of the `endUserMonitoringURL` to `http://<hostname>:<portnumber>/oracle_smp_chronos/oracle_smp_chronos_sdk.gif`. The hostname and port number are for the Web Cache that is serving the Forms application.

6.8.5.2 Changing the Logging Format for OracleAS Web Cache

To change the Logging format in the 10g version:

1. Navigate to the Web Cache Home page in Application Server Control.
2. Click **Administration**.

3. Click **Logging** and change the log file format to End-User Performance Monitoring format.
4. Click **OK** and restart the Web Cache.

For earlier versions of Web Cache, use the Web Cache Manager page to change the logging format of the Web Cache to the End-User Performance Monitoring format.

Note: After you have configured the Forms server and the Web Cache, follow the steps listed in "[Configuring End-User Performance Monitoring](#)" on page 6-13 to create the Forms Web application. You must then navigate to the Monitoring Configuration page of the Web application and specify FORMS in the Application Type field.

6.9 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](#)

6.9.1 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.

6.9.2 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.

6.9.3 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 OC4J 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.

6.9.4 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 6-28.
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.

6.10 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.

6.10.1 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.

6.11 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.
- **Actual Service Level:** This is calculated as percentage of time during business hours that your service meets both availability and performance criteria specified.
- **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**.

6.11.1 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, business hours, availability and performance criteria and click **OK**.

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.

6.11.2 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, 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 and service level information for a set of services.
- **Services Status Summary:** Information on one or more services' current status, performance, usage and component statuses.

Refer to the Online Help for more details on the report elements.

6.12 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.

Locating and Configuring Enterprise Manager Log Files

When you install the Oracle Management Agent or the Oracle Management Service, Enterprise Manager automatically configures the system to save certain informational, warning, and error information to a set of log files.

Log files can help you troubleshoot potential problems with an Enterprise Manager installation. They provide detailed information about the actions performed by Enterprise Manager and whether or not any warnings or errors occurred.

This chapter not only helps you locate and review the contents of Enterprise Manager log files, but also includes instructions for configuring the log files to provide more detailed information to help in troubleshooting or to provide less detailed information to save disk space.

This chapter contains the following sections:

- [Locating and Configuring Management Agent Log and Trace Files](#)
- [Locating and Configuring Management Service Log and Trace Files](#)

7.1 Locating and Configuring Management Agent Log and Trace Files

The following sections provide information on the log and trace files for the Oracle Management Agent:

- [About the Management Agent Log and Trace Files](#)
- [Locating the Management Agent Log and Trace Files](#)
- [About Management Agent Rollover Files](#)
- [Controlling the Size and Number of Management Agent Log and Trace Files](#)
- [Controlling the Size and Number of Fetchlet Log and Trace Files](#)
- [Controlling the Contents of the Fetchlet Trace File](#)

7.1.1 About the Management Agent Log and Trace Files

Oracle Management Agent log and trace files store important information that support personnel can later use to troubleshoot problems. The Management Agent uses three types of log files:

- The Management Agent log file (`emagent.log`)

The Agent saves information to the log file when the Agent performs an action (such as starting, stopping, or connecting to a Management Service) or when the

Agent generates an error (for example, when the Agent cannot connect to the Management Service).

- The Management Agent trace file (`emagent.trc`)

The Management Agent trace file provides an advanced method of troubleshooting that can provide support personnel with even more information about what actions the Agent was performing when a particular problem occurred.

- The Management Agent startup log file (`emagent.nohup`)

The Management Agent saves information to the startup log file when there is a problem starting the agent. This file is updated by the Management Agent Watchdog Process. When the Watchdog Process finds any problems, it logs to this file.

See Also: ["About the Management Agent Watchdog Process"](#) on page 10-4

In addition, Enterprise Manager also provides a log file and a trace file for the fetchlets, which are software programs used by the Management Agent for certain data-gathering tasks:

- `emagentfetchlet.log`
- `emagentfetchlet.trc`

7.1.2 Locating the Management Agent Log and Trace Files

The Management Agent log files are stored in the following directory when you install the Management Agent:

```
AGENT_HOME/sysman/log/
```

See Also: [Chapter 1, "Introduction to Enterprise Manager Advanced Configuration"](#) for information about locating the Agent home directory.

7.1.3 About Management Agent Rollover Files

Both the Management Agent log file and the Management Agent trace file are designed to increase in size over time as information is written to the files. However, they are also designed to reach a maximum size. When the files reach the predefined maximum size, the Management Agent renames (or rolls) the logging or trace information to a new file name and starts a new log or trace file. This process keeps the log files from growing too large.

To be sure you have access to important log or trace file information, the Management Agent will rollover the log and trace files four times by default. When it rolls the log or trace file over the fourth time, the Agent deletes the oldest rollover file.

As a result, you will often see a total of four log files and four trace files in the log directory. The following example shows three archived trace files and the current trace file in the `AGENT_HOME/sysman/log` directory:

```
emagent.trc
emagent.trc.1
emagent.trc.2
emagent.trc.3
```

7.1.4 Controlling the Size and Number of Management Agent Log and Trace Files

You can control how large the log file and the trace file can get before the Management Agent creates a rollover file. You can also control how many rollover files are created before the Management Agent deletes any logging or tracing data.

To control the size and number of Management Agent Log and Trace Files:

1. Stop the Management Agent.

See Also: ["Starting, Stopping, and Checking the Status of the Management Agent on UNIX"](#) on page 2-1

2. Locate the `emd.properties` file, which is located in the following directory:

`AGENT_HOME/sysman/config/` (UNIX)
`AGENT_HOME\sysman\config` (Windows)

3. Use a text editor to open the `emd.properties` file.
4. Use the information in [Table 7-1](#) to locate and modify the Agent logging and tracing properties in the `emd.properties` file.
5. Restart the Management Agent.

Table 7-1 Management Agent Log and Trace File Properties

Property	Purpose	Example
<code>LogFilewithPID</code>	When set to TRUE, this property appends the process ID of the Management Agent to the log file name. This makes it easier to identify the process ID of the Management Agent you are monitoring.	<code>LogFilewithPID=true</code>
<code>LogFileMaxSize</code>	When the Agent log file reaches this size (in kilobytes), the Management Agent copies the logging data to a new rollover file and creates a new <code>emagent.log</code> logging file.	<code>LogFileMaxSize=4096</code>
<code>LogFileMaxRolls</code>	By the default, the Agent will rollover the log file four times before it deletes any logging data. The number of rollover files is controlled by this property.	<code>LogFileMaxRolls=4</code>
<code>TrcFileMaxSize</code>	When the Agent trace file reach this size (in kilobytes), the Management Agent copies the logging data to a new rollover file and creates a new <code>emagent.trc</code> logging file.	<code>TrcFileMaxSize=4096</code>
<code>TrcFileMaxRolls</code>	By the default, the Agent will rollover the trace file four times before it deletes any tracing data. The number of rollover files is controlled by this property.	<code>TrcFileMaxRolls=4</code>

7.1.5 Controlling the Contents of the Management Agent Trace File

To modify the amount of information saved in the Management Agent trace file:

1. Stop the Management Agent.

See Also: ["Starting, Stopping, and Checking the Status of the Management Agent on UNIX"](#) on page 2-1

2. Locate the `emd.properties` file, which is located in the following directory:

```
AGENT_HOME/sysman/config
```

3. Open the `emd.properties` file using your favorite text editor and look for the following entries near the bottom of the file:

```
tracelevel.main=WARN
tracelevel.emdSDK=WARN
tracelevel.emdSDK.util=WARN
tracelevel.ResMonitor=WARN
tracelevel.Dispatcher=WARN
tracelevel.ThreadPool=WARN
tracelevel.pingManger=WARN
.
.
.
```

Each of these properties controls the level of logging detail for the various subcomponents of the Management Agent.

4. Modify the amount of information that is included in the trace file by replacing the WARN value for each property to one of the values shown in [Table 7-2](#).

Note: The values described in [Table 7-2](#) are case-sensitive.

5. Restart the Management Agent.

Table 7-2 Enterprise Manager Component Tracing Levels

Level	Purpose
ERROR	Include only critical errors in the trace file. This setting generates the least amount of tracing data. The trace file will likely grow at a relatively slow rate when you select this logging level.
WARN	Include warning information, in addition to critical errors.
INFO	Include informational messages, in addition to warning and critical error information.
DEBUG	Include debugging information, as well as informational tracing, warning, and critical errors. This setting generates the greatest amount of tracing data. Note: The trace file will likely grow at a relatively fast rate when you select this logging level.

7.1.6 Controlling the Size and Number of Fetchlet Log and Trace Files

Like the Management Agent log and trace files, the Management Agent fetchlet log and trace files are designed to reach a maximum size before the Management Agent renames (or rolls) the information to a new file name and starts a new log or trace file.

To control the maximum size of the Management Agent fetchlet log and trace files, as well as the number of rollover files:

1. Stop the Management Agent.

See Also: ["Starting, Stopping, and Checking the Status of the Management Agent on UNIX"](#) on page 2-1

2. Locate the `emagentlogging.properties` file in the following directory:

`AGENT_HOME/sysman/config`

3. Open the `emagentlogging.properties` file with a text editor and modify the entries described in [Table 7-3](#).

4. Restart the Management Agent.

Table 7-3 Management Agent Servlet Log and Trace File Properties

Property	Purpose	Example
log4j.appender. emagentlogAppender. MaxFileSize	When the fetchlet log file reaches this size, the Management Agent copies the logging data to a new rollover file and creates a new <code>emagentfetchlet.log</code> file.	log4j.appender. emagentlogAppender. MaxFileSize=2000000
log4j.appender. emagentlogAppender. MaxBackupIndex	This optional property indicates how many times the Management Agent will rollover the fetchlet log file to a new file name before deleting logging data. Note: Because the log file does not contain as much data as the trace file, it is usually not necessary to create more than one rollover file. As a result, this entry is not included in the properties file by default.	log4j.appender.emagentlogAppender. MaxBackupIndex=1
log4j.appender. emagenttrcAppender. MaxFileSize	When the fetchlet trace file reaches this size, the Management Agent copies the logging data to a new rollover file and creates a new <code>emagentfetchlet.trc</code> log file.	log4j.appender. emagenttrcAppender. MaxFileSize=5000000
log4j.appender. emagenttrcAppender. MaxBackupIndex	This property indicates how many times the Management Agent will rollover the trace file to a new file name before deleting tracing data.	log4j.appender. emagenttrcAppender. MaxBackupIndex=10

7.1.7 Controlling the Contents of the Fetchlet Trace File

By default, the Management Agent will save all critical and warning messages generated by the Management Agent fetchlets to the `emagentfetchlet.trc` file. However, you can adjust the amount of logging information that the fetchlets generate.

To change the amount of tracing information generated by the Management Agent fetchlets:

1. Stop the Management Agent.

See Also: ["Starting, Stopping, and Checking the Status of the Management Agent on UNIX"](#) on page 2-1

2. Locate the `emagentlogging.properties` file in the following directory:

`AGENT_HOME/sysman/config`

3. Open the `emagentlogging.properties` file with a text editor and locate the following entry:

```
log4j.rootCategory=WARN, emagentlogAppender, emagenttrcAppender
```

4. Change the value of the `log4j.rootCategory` parameter to one of the values shown in [Table 7-2](#).

Note: The the values described in [Table 7-2](#) are case-sensitive.

5. Restart the Management Agent.

7.2 Locating and Configuring Management Service Log and Trace Files

The following sections describe how to locate and configure the Management Service log files:

- [Locating the Management Service Log and Trace Files](#)
- [Controlling the Size and Number of Management Service Log and Trace Files](#)
- [Controlling the Contents of the Management Service Trace File](#)
- [Controlling the Oracle Application Server Log Files](#)

7.2.1 About the Management Service Log and Trace Files

Oracle Management Service log and trace files store important information that support personnel can later use to troubleshoot problems. The Management Service uses two types of log files:

- The Management Service log file (`emoms.log`)
The Oracle Management Service saves information to the log file when the Management Service performs an action (such as starting or stopping) or when the Management Service generates an error.
- The Management Service trace file (`emoms.trc`)
The Management Service trace file provides an advanced method of troubleshooting that can provide support personnel with even more information about what actions the Management Service was performing when a particular problem occurred.

7.2.2 Locating the Management Service Log and Trace Files

The Management Service log and trace files are stored in the following directory inside the Oracle Application Server Home where the Oracle Management Service is installed and deployed:

```
AS_HOME/sysman/log/
```

7.2.3 Controlling the Size and Number of Management Service Log and Trace Files

The Management Service log and trace files increases in size over time as information is written to the files. However, the files are designed to reach a maximum size. When the files reach the predefined maximum size, the Management Service renames (or

rolls) the logging information to a new file name and starts a new log or trace file. This process keeps the log and trace files from growing too large.

As a result, you will often see multiple log and trace files in the Management Service log directory. The following example shows one archived log file and the current log file in the `AS_HOME/sysman/log` directory:

```
emoms.log
emoms.log.1
```

To control the maximum size of the Management Service log and trace files, as well as the number of rollover files:

1. Stop the Management Service.

See Also: ["Controlling the Oracle Management Service"](#) on page 2-4

2. Locate the `emomslogging.properties` file in the following directory:

```
AS_HOME/sysman/config
```

3. Open the `emomslogging.properties` file with a text editor and modify the entries described in [Table 7-4](#).
4. Restart the Management Service.

Table 7-4 Management Service Log File Properties in the `emomslogging.properties` File

Property	Purpose	Example
<code>log4j.appender.emlogAppender.MaxFileSize</code>	When the Management Service log file reaches this size, the Management Service copies the logging data to a new rollover file and creates a new <code>emoms.log</code> log file.	<code>log4j.appender.emlogAppender.MaxFileSize=2000000</code>
<code>log4j.appender.emlogAppender.MaxBackupIndex</code>	This optional property indicates how many times the Management Service will rollover the log file to a new file name before deleting logging data. Note: Because the log file does not contain as much data as the trace file, it is usually not necessary to create more than one rollover file. As a result, this entry is not included in the properties file by default.	<code>log4j.appender.emlogAppender.MaxBackupIndex=1</code>
<code>log4j.appender.emtrcAppender.MaxFileSize</code>	When the Management Service trace file reaches this size, the Management Service copies the logging data to a new rollover file and creates a new <code>emoms.trc</code> log file.	<code>log4j.appender.emtrcAppender.MaxFileSize=5000000</code>
<code>log4j.appender.emtrcAppender.MaxBackupIndex</code>	This property indicates how many times the Management Services will rollover the trace file to a new file name before deleting tracing data.	<code>log4j.appender.emtrcAppender.MaxBackupIndex=10</code>

7.2.4 Controlling the Contents of the Management Service Trace File

By default, the Management Service will save all critical and warning messages to the `emoms.trc` file. However, you can adjust the amount of logging information that the Management Service generates.

To change the amount of logging information generated by the Management Service:

1. Stop the Management Service.

See Also: ["Controlling the Oracle Management Service"](#) on page 2-4

2. Locate the `emomslogging.properties` file in the following directory:

```
AS_HOME/sysman/config
```

3. Open the `emomslogging.properties` file with a text editor and locate the following entry:

```
log4j.rootCategory=WARN, emlogAppender, emtrcAppender
```

4. Modify the value of the `log4j.rootCategory` parameter to one of the values shown in [Table 7-2](#).

Note: The values described in [Table 7-2](#) are case-sensitive.

5. Restart the Management Service.

7.2.5 Controlling the Oracle Application Server Log Files

The Management Service is a J2EE application running in an Oracle Application Server Containers for J2EE (OC4J) instance within the Application Server. Different components of the Application Server generate their own log files. These files contain important information that can be used later by support personnel to troubleshoot problems.

[Table 7-5](#) lists the location of the log files for some components.

Table 7-5 Component Log File Location

Component	Location
HTTP Server	ORACLE_HOME/Apache/Apache/logs/error_log,time ORACLE_HOME/Apache/Apache/logs/access_log,time
OC4J	ORACLE_HOME/j2ee/instance_name/logORACLE_HOME/j2ee/instance_name/application-deployments/application_name/application.log
OPMN	ORACLE_HOME/opmn/logs
Web Cache	ORACLE_HOME/webcache/logs

Refer to the Oracle Application Server Administrator's Guide for instructions on controlling the size and rotation of these log files.

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](#)
- [Improving the Login Performance of the Console Home Page](#)

8.1 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 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

8.2 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.

8.2.1 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 8-1](#).

Table 8-1 *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 8-2](#).

Table 8–2 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

8.2.2 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.

8.2.3 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 8–3](#) 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 8–3 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

Table 8–3 (Cont.) Parameters for Modifying Default Data Retention Times in the Management Repository

Table Name	Parameter in MGMT_PARAMETERS Table	Default Retention Value
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');
```

8.2.4 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/oldpassword;
```

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;
```

8.3 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 2-4

2. In the Grid Control Console, click the **Targets** tab, and then click **All Targets** on the sub tab.

3. Select the **Management Services and Repository** target and click **Configure**.
Enterprise Manager displays the Monitoring Configuration page.
4. Enter the new password in the **Repository password** field and click **OK**.

See Also: ["Specifying New Target Monitoring Credentials"](#) on page 2-13

5. 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;
```

6. 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/
```

7. Locate the following entries in the `emoms.properties` file:

```
oracle.sysman.eml.mntr.emdRepPwd=ece067ffc15edc4f
oracle.sysman.eml.mntr.emdRepPwdEncrypted=TRUE
```

8. 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
```

9. Save and exit the `emoms.properties` file and restart each Management Service associated with the Management Repository.
10. After the Management Service has started, 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
```

8.4 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.

8.4.1 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=servicename)))"
-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

8.4.2 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](#)

8.4.2.1 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 8-1.
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.

8.4.2.2 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 efkl34lmm -action create
```

See Also: "Establishing a Connection and Testing the Network" in the *Oracle Database Net Services Administrator's Guide* for more information about connecting to a database using a connect descriptor

The ability to use a connect string allows you to provide an address list as part of the connection string. The following example shows how you can provide an address list consisting of two listeners as part of the `RepManager` command line. If a listener on one host becomes unavailable, the second listener can still accept incoming requests:

```
$PROMPT> ./RepManager -connect "(DESCRIPTION=
(ADDRESS_LIST=
(ADDRESS=(PROTOCOL=TCP)(HOST=host1)(PORT=1521)
(ADDRESS=(PROTOCOL=TCP)(HOST=host2)(PORT=1521)
(CONNECT_DATE=(SERVICE_NAME=servicename)))"
-sys_password efkl34lmm -action create
```

See Also: *Oracle High Availability Architecture and Best Practices*
"Configuring the Management Service to Use Oracle Net Load Balancing and Failover" on page 3-18

8.5 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.

8.5.1 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 8-9.

8.5.2 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*.

8.5.3 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 8-6

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;
```

8.6 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.

Sizing and Maximizing the Performance of Oracle Enterprise Manager

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](#)
- [Configuring Enterprise Manager for High Availability](#)

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)	2
Medium (1,000 monitored targets)	1	2 (3 GHz)	2
Large (10,000 monitored targets)	2	2 (3 GHz) 2	2

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 (GB)

Deployment Size	Total Management Repository Storage (GB)
Small	10
Medium	30
Large	100

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.

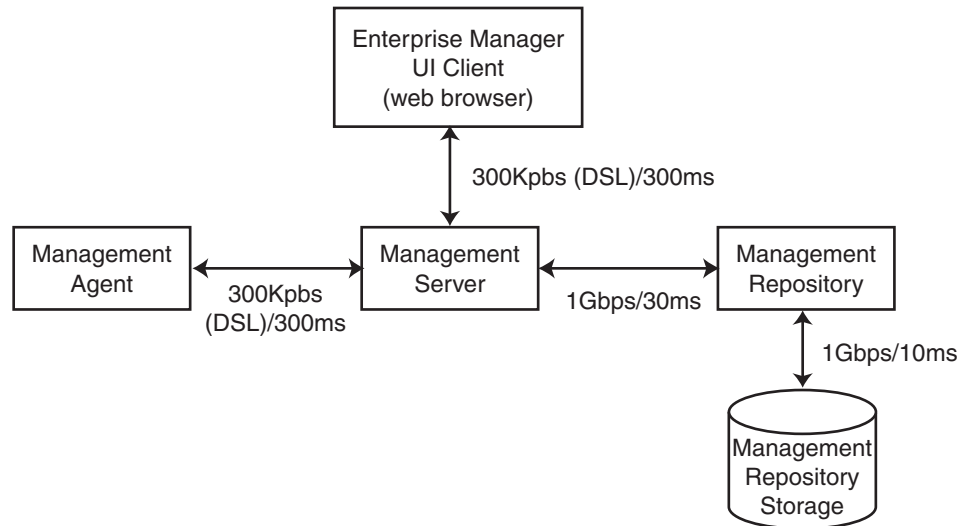
Deploying multiple Management Servers also provides basic fail-over capabilities, with the remaining servers continuing to operate in the event of the failure of one. Use of a Server Load Balancer, or SLB, provides transparent failover for Enterprise Manager UI clients in the event of a Management Server 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:

		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)
	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

		EM Site 1	EM Site 2
	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+
	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.

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

- Check the system error page and resolve the causes of all errors. Some may be related to product bugs, but resolve as many as you can. Look for applicable patches if you suspect a bug. Clear the error table from the Enterprise Manager interface when you are done or when you have resolved all that you can.
- Check the alerts and errors for any metric collection errors. Most of these will be due to configuration issues at the target being monitored. Resolve these errors by fixing the reported problem. The error should then clear automatically.
- Try to resolve any open alerts in the system. Also, if there are severities that are frequently oscillating between clear and warning or critical, try adjusting the threshold to stop frequent warning and critical alert conditions. Frequent alert oscillation can add significant load at the Management Server. Adjusting the threshold to a more reasonable level will help Enterprise Manager to work more

efficiently for you. Adjusting the threshold for an alert may be the only way to close the alert. This is perfectly acceptable in cases where the tolerances are too tight for a metric.

- Watch for monitored targets that are always listed with a down status. Try to get them up and working again, or remove them from Oracle Enterprise Manager.
- Watch the Alert Log error metric for the Management Repository database for critical (ORA-0600, for example) errors. Resolve these as soon as possible. A search on Metalink using the error details almost always will reveal some clues to its cause and provide available patches.
- 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, .000001, false, 'for all indexed columns', null, 'global', true, null, null, null);`
 - `exec dbms_stats.gather_table_stats('SYSMAN', 'MGMT_METRICS_1HOUR', null, .000001, false, 'for all indexed columns', null, 'global', true, null, null, null);`
 - `exec dbms_stats.gather_table_stats('SYSMAN', 'MGMT_METRICS_1DAY', null, .000001, false, 'for all indexed columns', null, 'global', true, null, null, null);`

9.2.3.2 Offline Monthly Tasks

- Drop old partitions. Oracle Enterprise Manager automatically truncates the data and reclaim the space used by partitions older than the default retention times for each table. Due to a database bug, however, Enterprise Manager cannot drop partitions while the Management Service is running because many SQL cursors will be invalidated incorrectly leading to some strange errors in the Enterprise Manager interface. The following command must be run with all Management Servers down:
 - `exec emd_maintenance.partition_maintenance;`
- 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 a command similar to the following example:
 - `exec dbms_redefinition.start_redef_table('SYSMAN', 'MGMT_SEVERITY');`

This command rebuilds 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_SEVERITY
- MGMT_CURRENT_SEVERITY
- MGMT_SYSTEM_ERROR_LOG
- MGMT_SYSTEM_PERFORMANCE_LOG
- MGMT_METRIC_ERRORS
- MGMT_CURRENT_METRIC_ERRORS
- MGMT_STRING_METRIC_HISTORY
- MGMT_JOB_OUTPUT

These are a sampling of tables that may require more DBA attention than others, but all non-IOT Enterprise Manager tables and indexes should be evaluated monthly for defragmentation and rebuild needs. The following query gives a rough idea of the tables that may require rebuild and reorganization or both:

```
SELECT UT.TABLE_NAME, ROUND(UT.NUM_ROWS * UT.AVG_ROW_LEN /
1024 / 1024, 2) "CALCULATED SIZE MB", ROUND(US.BYTES / 1024
/1024,2) "ALLOCATED SIZE MB", ROUND(US.BYTES / (UT.NUM_ROWS *
UT.AVG_ROW_LEN), 2) "TIMES LARGER" FROM USER_TABLES UT, USER_
SEGMENTS US WHERE (UT.NUM_ROWS > 0 AND UT.AVG_ROW_LEN > 0 AND
US.BYTES > 0) AND UT.PARTITIONED = 'NO' AND UT.IOT_TYPE IS
NULL AND UT.IOT_NAME IS NULL AND UT.TABLE_NAME = US.SEGMENT_
NAME AND ROUND(US.BYTES / 1024 /1024,2) > 5 AND
ROUND(US.BYTES / 1024 /1024,2) > (ROUND(UT.NUM_ROWS * UT.AVG_
ROW_LEN / 1024 / 1024, 2)* 2) ORDER BY 4 DESC;
```

Sample query output:

Table Name	Calculated Size MB	Allocated Size MB	Times Larger
MGMT_JOB_OUTPUT	2.25	440	195.57
MGMT_FLAT_TARGET_MEMBERSHIPS	2.33	160	68.67
MGMT_ANNOTATIONS	1.21	21	17.3
MGMT_SQL_SUMMARY	6.06	80	13.2
MGMT_JOB_EXECUTION	1.03	12	11.6
MGMT_SYSTEM_ERROR_LOG	5.6	61	10.88
MGMT_JOB_HISTORY	1.11	10	9.04
MGMT_NOTIFICATION_LOG	1.59	14	8.78

Note: This query calculates the actual size of a table based on the number of rows and average row size. It compares this actual size to the currently allocated size of the table. The final column shows how many times larger the allocated size is than the calculated actual size. Use this as a guide to determine which tables and indexes should be rebuilt. Tables that are many times larger than the actual size should be rebuilt, along with their indexes, using the commands mentioned previously.

Good housekeeping will prevent many bottlenecks from occurring on your Enterprise Manager Grid Control site, but there may be times when you should investigate performance problems on your site that are not related to housekeeping. This is where the Enterprise Manager vital signs become important.

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:

```
em.loader.threadPoolSize=n
```

Where 'n' is a positive integer [1-10]. The default is one and any value other than [1-10] will result in the thread pool size defaulting to one. 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.

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 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.

9.2.4.5 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.

9.2.4.6 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

The newest release of Oracle Enterprise Manager Grid Control incorporates a portable browser-based interface to the management console and the Oracle application server technology to serve as the middle-tier Management Service. The foundation of the tool remains rooted in database server technology to manage the Management Repository and historical data. This new architecture requires a different approach to backup, recovery and Disaster Recovery (DR) planning. This section provides practical

approaches to these availability topics and discusses different strategies when practical for each tier of Enterprise Manager.

9.3.1 Best Practices for Backup and Recovery

For the database, the best practice is to use the standard database tools for any database backup; have the database in archivelog mode, and perform regular online backup using RMAN or OS commands.

There are two cases to consider with regard to recovery:

- Full recovery of the Management Repository is possible: 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: Enterprise Manager Management Agents will be unable to communicate to the Management Repository correctly until they are reset. This is a manual process that is accomplished by shutting down the Management Agent, deleting the `agntstmp.txt` and `lastupld.xml` files in the `$AGENT_HOME/sysman/emd` directories and then going to the `/state` and `/upload` subdirectories and clearing the contents. The Management Agent can then be restarted. This would need to be done for each Management Agent.

For the case of incomplete recovery, Management Agents may not be able to upload data until the previous steps are completed. Additionally, there is no indication in the interface that the Management Agents may not 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, it is best to perform this procedure for each Management Agent.

9.3.1.1 Oracle Management Service

As 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 and restoring that in the event of failure to the same directory path. The Management Agent associated with this Management Service install should also be backed up at the same time and restored with the Management Service files if a restore is required.
- Reinstall from the original media.

For any highly available Management Service install it is a recommended practice to make sure the `/recv` directory is protected with some mirroring technology. This is the directory the Management Service uses to stage files send to it from Management Agents before writing their contents to the database Management Repository. After the Management Agent finishes transmission of its XML files to the Management Service, it will delete its copy. In the event of an Management Service disk failure, this data would be lost. Warnings and alerts sent from the Management Agents would then be lost. This may require Management Agent resynchronization steps similar to those used with an incomplete database recovery.

9.3.1.2 Management Agent

This is a similar case to the Management Service except that the Management Agent is not stateless. There are two strategies that can be used:

- A disk backup and restore is sufficient, assuming the host name has not changed. Delete the `agntstmp.txt` and the `lastupld.xml` files from the `/sysman/emd` directory. The `/state` and `/upload` sub-directories should be cleared of all entries before restarting. Starting the Management Agent will then force a rediscovery of targets on the host.
- Reinstall from the original media.

As with the Management Service, it is a recommended best practice to protect the `/state` and `/upload` directories with some form of disk mirroring.

9.3.2 Best Practice for Disaster Recovery (DR)

In the event of a node failure, the database can be restored using RMAN or OS commands. To speed this process, implement Data Guard to replicate the Management Repository to a different hardware node.

9.3.2.1 Management Repository

If 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, the `targets.xml` for each Management Service will have to be updated to reflect the new Management Repository location. If there is a data loss during recovery, see the notes above on incomplete recovery of the Management Repository.

To speed Management Repository reconnection from the Management Service in the event of a single Management Service failure, configure the Management Service with a TAF aware connect string. The Management Service can be configured 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)))
```

9.3.2.2 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 binaries from backup and modifying the Management Service and Management Agent configuration files.

Note that it is not recommended to restore the Management Service and Management Agent binaries from an existing backup to a new host in the event of a disaster as there are host name dependencies. Always do a fresh install.

9.3.2.3 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.

9.4 Configuring Enterprise Manager for High Availability

Oracle customers deploy systems that are considered critical to their business. These systems often have strict availability requirements and maintenance windows. Downtime is often measured in minutes and maintenance windows are short. Oracle has addressed this business need with the rollout of the 'Unbreakable' database and blueprints for highly available systems such as the Maximum Availability Architecture. With the release of Oracle Enterprise Manager 10g Grid Control, Oracle has increased the manageability of highly available systems. This also increases the availability requirements for the manageability infrastructure.

This section describes a highly available deployment of Grid Control. The section should help you understand the steps required to configure each component for high availability. It also discusses the strengths and limitations of the current solution and you will have an understanding of how to recover from outages of each tier.

9.4.1 Architectural Overview

The architecture for a highly available Grid Control deployment is based on two key concepts; redundancy and component monitoring. Each component of Grid Control can be configured to apply both these concepts.

The components of Grid Control discussed in this section include:

- Management Agent
- Management Service
- Management Repository

For more detail about each of these components, see [Oracle Enterprise Manager Grid Control Architecture Overview](#) on page 9-1.

The Management Agent uploads collected monitoring data to a Management Service. The Management Service in turn loads the data into the Management Repository. The Management Repository represents the persistent historic view of collected information that is presented to clients using a web user interface.

Changes in a target state either in an availability state change or detection of a notification dependent upon a metric threshold being crossed results in a notification being sent. The Management Agent detects this change and is responsible for forwarding the information to the Management Service that in turn, records the state change in the Management Repository. Any registered users requesting notification have messages posted using registered notification methods by the Management Service and the console display updated.

Details on using Enterprise Manager to configure high availability features such as RMAN and Data Guard can be found in the Oracle documentation and in the Enterprise Manager Grid Control online help.

9.4.2 Installation and Configuration for High Availability

The following sections document best practices for installation and configuration of each Grid Control component.

9.4.2.1 Management Agent

Enterprise Manager uses a software process called the Oracle Management Agent to monitor a target. The Management Agent is a system daemon that consists of two processes, a process that provides monitoring, alerting and job system capabilities, as

well as a watchdog process that is responsible for insuring the Management Agent is up and available.

The data that is collected by the Management Agent is stored temporarily on the monitored host in files. Once the Management Agent deems it necessary to upload the information to the Grid Control system, it contacts the Management Service to establish a connection and uploads the data.

The Management Service accepts the data from the Management Agent, stores the information as files local to the Management Service and acknowledges receipt of the information to the Management Agent. Depending on the volume of work the Management Service is performing, a period of time may elapse before the Management Service loads the data into the Management Repository.

Notifications of alerts, warnings and target state changes do not follow this delayed model. When the Management Agent uploads the information, the Management Service commits the data immediately to the Management Repository before acknowledgement is returned to the Management Agent.

The Management Agent and its watchdog are started through the command `'$ORACLE_HOME/bin/emctl start agent.'`

9.4.2.1.1 Configuring the Management Agent to Automatically Start on Boot and Restart on Failure

The Management Agent is started manually. It is important that the Management Agent be automatically started when the host is booted to insure monitoring of critical resources on the administered host. To that end, use any and all operating system mechanisms to automatically start the Management Agent. For example, on UNIX systems this is done by placing an entry in the UNIX `/etc/init.d` that calls the Management Agent on boot or by setting the Windows service to start automatically.

9.4.2.1.2 Configuring Restart for the Management Agent

Once the Management Agent is started, the watchdog process monitors the Management Agent and attempts to restart it in the event of a failure. The behavior of the watchdog is controlled by environment variables set before the Management Agent process starts. The 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.

9.4.2.1.3 Configuring the Connection Between Management Agents and the Management Service

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 re-attempts to send the information later.

Server Load Balancers (SLBs) such as the F5 Networks Big-IP provide logical service abstractions for network clients. Clients establish connections to the virtual service exposed by the SLB. The SLB routes the request to any one of a number of available servers that provide the requested service. The service chosen by an SLB as the destination is dependent upon the virtual service definition. One such criterion is whether a service is capable of accepting connections.

The Grid Control Management Service is a network service that can be fronted by a SLB to address the need for resiliency.

To accomplish the goal of having a highly available Management Service that the Management Agents can use for data upload, configure a virtual pool that consists of the hosts and the services that the hosts provide. In the case of the Management Services pool, the hostname and Management Agent upload port would be specified. To insure a highly available Management Service, you should have two or more Management Services defined within the virtual pool.

See Also: [Load Balancing Connections Between the Management Agent and the Management Service.](#)

9.4.2.1.4 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.

9.4.2.1.5 Configuring All Out-of-band Notifications

The Enterprise Manager Grid Control deployment is configured out of the box such that connection failures between the Management Service and the Management Agent are detected. This is through a process of heartbeats that the Management Agent performs against the Management Service. If the Management Service determines it has not heard back from the Management Agent, it pings it.

This condition does not, however, correct for a condition where the Management Agents are up and available but there are no Management Services to which to upload or process notifications. For this situation, the Management Agent has the capability of sending an emergency notification when it is still up but has lost contact with the Management Service.

This provides another mechanism to alert the administrator of a Management Service failure. For more information, see [Section 9.4.3, "Configuration Within Grid Control"](#) on page 9-23.

In the `emd.properties` file located in the `$AGENT_HOME/sysman/config` directory, modify the property values for `emd_email_address` and `emd_email_`

gateway to reflect a valid e-mail address in your system. The parameter `emd_from_email_address` should also be modified to reflect the name of the system sending the alert for faster root cause identification.

In addition, any custom notification script can be executed by the Management Agent in the event of a failure to communicate with the Management Service. This script can be set to execute by modifying the `'emdFailureScript'` entry in the Management Agent `emd.properties` file.

9.4.2.2 Management Service

The Management Service element of the Enterprise Manager Grid Control product acts both as the receiver of information from Management Agents as well as serves out the User Interface in the form of HTML pages. It does this by maintaining a connection to the configurations database Management Repository and responding to requests over HTTP.

9.4.2.2.1 Configuring the Shared Filesystem Loader

Configure the Management Services to use 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 amongst themselves the workload of uploading files into the Management Repository. Should a Management Service go down for some reason, its workload is taken up by surviving Management Services.

See Also: [Section 3.6.1.1, "Configuring the Management Services for High Availability"](#).

9.4.2.2.2 Configuring SLB to Abstract the Underlying Management Service Host Names for Easier Reconnect After Failure

A hardware server load balancer (SLB) such as F5 Networks Big-IP can be used as the front end to abstract the number and location of Management Services and appear as a single service. Under that abstraction, the SLB parcels the work to any number of Management Service processes that it has in its 'virtual pool.' For any Grid Control installation with an availability requirement there should be a minimum of two Management Service processes installed. Coupled with a SLB, this provides a method for constant communication to the Grid Control Console in the event of the failure of a Management Service.

For more details on configuring SLB for Shared Filesystem Loader, see [Section 3.6.1, "Load Balancing Connections Between the Management Agent and the Management Service"](#) on page 3-10.

9.4.2.2.3 Management Service Installation Should Be Done to Non-Clustered Servers

Management Service processes cannot be installed on any machines running under a cluster, whether it is CRS or vendor cluster software. Install Management Services to single nodes and use the method described previously for failover and availability.

9.4.2.2.4 Configuring Management Service to Use Client Side Oracle Net Load Balancing for Failover and Load Balancing

When you use a RAC cluster, a standby system, or both to provide high availability for the Management Repository, the Management Service can be configured to use an Oracle Net connect string that will take advantage of redundancy in the Management

Repository. Correctly configured, the Management service process will continue to process data from Management Agents even during a database node outage.

In the `$OMS_HOME/sysman/config` directory, modify the `emdRepConnectDescriptor` entry in the `emoms.properties` file to point to the appropriate Management Repository instances. The following example shows a connect string required to support a 2-node RAC configuration. Note the backslash (`\`) before each equal sign (`=`).

```
oracle.sysman.em1.mntr.emdRepConnectDescriptor=
(DESCRIPTION\=(ADDRESS_LIST\=(FAILOVER\=ON)
(ADDRESS\=(PROTOCOL\=TCP) (HOST\=haem1.us.oracle.com)
(PORT\=1521)) (ADDRESS\=(PROTOCOL\=TCP)
(HOST\=haem2.us.oracle.com) (PORT\=1521))) (CONNECT_
DATA\=(SERVICE_NAME\=em10))
```

9.4.2.2.5 Install the Management Service Software 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.

Similar to the Management Agent directories, availability would be further enhanced by placing the entire Management Service software tree on redundant storage. This can also be determined at the command line using the `'emctl getemhome'` or by using the Management Services and Repository tab in the Grid Control Console.

9.4.2.3 Management Repository

The Management Repository is the central location for all historical data managed by Grid Control. Redundancy at this tier is provided by standard database features and best practices.

9.4.2.3.1 Install Into an Existing RAC Management Repository

The Grid Control installation process does not directly support installation into a RAC Management Repository. The recommended installation method is to install the database software first and create a RAC database. When this is complete, install the Enterprise Manager software, selecting the **Enterprise Manager Grid Control Using an Existing Database** installation option.

The installation does not transparently support the installation of the Enterprise Manager 10g Grid Control into a RAC database. Specify the SID of one of the cluster instances when prompted for during the installation. After the installation of the Enterprise Manager 10g Grid Control Management Service, you should modify the connection string the Management Service uses to take advantage of client failover in the event of a RAC host outage (refer to [Section 9.4.2.2.4, "Configuring Management Service to Use Client Side Oracle Net Load Balancing for Failover and Load Balancing"](#) on page 9-21).

The installation process also does not allow modification of the size of the required Enterprise Manager tablespaces (although it does allow for specification of the name and location of data files that are to be used by the Enterprise Manager 10g Grid Control schema). The default sizes for the initial data file extents depend on using the AUTOEXTEND feature and as such are insufficient for a production installation. This is particularly problematic where storage for the RAC is on a raw device.

If the RAC database being used for the Management Repository is configured with raw devices there are two options for increasing the size of the Management Repository. You can create multiple raw partitions, with the first one equal to the default size of the tablespace as defined by the installation process. Alternatively, you can create the tablespace using the default size, create a dummy object that will increase the size of the tablespace to the end of the raw partition, then drop that object. Regardless, if raw devices are used, disable the default space management for these objects, which is to auto-extend.

9.4.2.3.2 Consider (Physical) Data Guard for Redundancy

Clients who require greater uptime or an off-site copy of the Management Repository can use Oracle Data Guard in conjunction with Grid Control. This alternative can be used regardless of whether or not you are using a RAC database. Currently, only the use of physical data guard is supported.

A Data Guard instance must be created manually using the steps documented in the Data Guard documentation.

9.4.3 Configuration Within 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.

9.4.3.1 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 agent unreachable and agent 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.

9.4.3.2 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.

9.4.3.3 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. Best practices for backing up these tiers are documented in the section, [Section 9.3, "Oracle Enterprise Manager Backup, Recovery, and Disaster Recovery Considerations"](#) on page 9-15.

9.4.3.4 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.

9.4.3.4.1 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.

9.4.3.4.2 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.

Reconfiguring the Management Agent and Management Service

This chapter describes how to reconfigure Enterprise Manager if you later revisit your configuration decisions after you have installed the software.

This chapter contains the following sections:

- [Reconfiguring the Oracle Management Agent](#)
- [Reconfiguring the Oracle Management Service](#)

10.1 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](#)

10.1.1 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.

To associate the Management Agent with a new Management Service after you have installed the Management Agent:

1. Stop the Management Agent.

See Also: ["Controlling the Oracle Management Agent"](#) on page 2-1

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` and `emdWalletDest` properties so they reference the new Management Service and the new Oracle home path, respectively:

For example, if the new Management Service is on a host called `mgmthost2.acme.com` and the new Oracle home is `/private/oracle/em10g`, modify the properties as follows:

```
emdWalletSrcUrl=http://mgmthost2.acme.com:4889/em/wallets/emd  
emdWalletDest=/private/oracle/em10g/sysman/config/server
```

6. Save your changes and close the `emd.properties` file.
7. Delete all the files in the following directories:

```
AGENT_HOME/sysman/emd/upload/  
AGENT_HOME/sysman/emd/state/
```

8. Restart the Management Agent.

10.1.2 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.

See Also: ["Controlling the Oracle Management Agent"](#) on page 2-1

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
```

5. Start the Management Agent.

10.1.3 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.

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.

See Also: ["Controlling the Oracle Management Agent"](#) on page 2-1

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 10-1](#).
4. Save your changes and exit the file.
5. Restart the Management Agent.

Table 10-1 Properties for Controlling the Disk Space Used by the Management Agent

Property	Explanation
UploadMaxBytesXML	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.
UploadMaxDiskUsedPct	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. 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.

10.1.4 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
```

10.1.5 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](#)
- [Troubleshooting Management Service Time Zone Problems](#)

10.1.5.1 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/suportedtzs.lst
```

10.1.5.2 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 reseTZ agent
```

This script updates `ORACLE_HOME/<hostname>_<sid>/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.

10.1.5.3 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/supportedtztz.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 10-2](#), depending upon which Enterprise Manager product you are using.

Table 10-2 Location of Time Zone Error in the Enterprise Manager Log Files

If you are using...	Look for the Time Zone Error in This File...
Grid Control Console	emagent.nohup
Application Server Control Console	em.nohup
Database Control Console	emdb.nohup

See Also: ["Locating and Configuring Management Agent Log and Trace Files"](#) on page 7-1 for more information about the Management Agent log files

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.

10.1.5.4 Troubleshooting Management Service Time Zone Problems

[Section 10.1.5.3](#) 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/nupportedtzs.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 10-2](#), 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 `timezlg.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 10-5 for instructions on changing the time zone assigned to the Management Agent

10.1.6 Adding Trust Points to the Management Agent Configuration

For Application Server components such as Oracle Portal to run on a secure sockets layer (SSL), the appropriate security certificate must be added to the Management Agent configuration files.

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
```

10.2 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 Use a New Port](#)
- [Configuring the Management Service to Prompt You When Using Execute Commands](#)

10.2.1 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.

This repository information is stored in the `emoms.properties` file, which can be found in the following directory where the Oracle Management Service is installed and deployed:

```
ORACLE_HOME/sysman/config/
```

The following sections describe how to modify the repository information in the `emoms.properties` file and provide details about how Enterprise Manager keeps the Management Repository password secure.

10.2.1.1 Changing the Repository Properties in the `emoms.properties` File

To associate the Management Service with a new repository, you must modify the repository properties saved in the `emoms.properties` configuration file:

1. Stop the Management Service.

See Also: ["Controlling the Oracle Management Service"](#) on page 2-4

2. Locate the `emoms.properties` file in the following directory where you installed and deployed the Management Service:

```
ORACLE_HOME/sysman/config/
```

3. Edit the `emoms.properties` file by updating the appropriate values for the properties described in [Table 10-3](#).

[Example 10-1](#) shows sample entries in the `emoms.properties` file.

4. Restart the Management Service.

Table 10–3 Repository Properties in the emoms.properties File

Property	Description
emdRepUser	The Management Repository user name. The default value is SYSMAN.
emdRepPwd	The Management Repository password. See "About Changing the Repository Password" on page 10-9 for information of how to change the password value.
emdRepConnectDescriptor	The Management Repository Oracle Net Connect String for the repository database. The values specified for properties emdRepSID, emdRepServer, and emdRepPort must be the same as that of HOST, PORT, and SERVICE_NAME in the connect string. If this property is not specified, then emRepSID, emRepServer, and emRepPort 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 "Configuring the Management Service to Use Oracle Net Load Balancing and Failover" on page 3-18
emdRepSID	The System Identifier (SID) for the database where the Management Repository schema resides.
emdRepServer	The name of the server or host computer where the repository database resides.
emdRepPort	The port number for the repository database.

Example 10–1 Sample Repository Properties in the emoms.properties File

```
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
```

10.2.1.2 About Changing the Repository Password

For security reasons, the password stored in the emoms.properties file is encrypted as soon as you start the Management Service. To change the repository password in the emoms.properties file, use the emctl setpasswd oms 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 in ORACLE_HOME/sysman/config/emoms.properties by using the following command:

```
ORACLE_HOME/bin/emctl setpasswd oms
```

3. Restart the Management Service using the following command:

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

10.2.2 Configuring the Management Service to Use a New Port

When you install the Management Service, the port number for the Management Service is automatically set to 4889. The following procedure describes how to manually change the port number after the Enterprise Manager installation. For

example, you will have to modify the port number if you attempt to install two Oracle Management Services on the same host computer.

To change the default Management Service port:

1. Stop the Management Service.

See Also: ["Controlling the Oracle Management Service"](#) on page 2-4

2. Locate the following `httpd_em.conf` file located in the following directory in the home directory where you installed and deployed the Management Service:

```
ORACLE_HOME/sysman/config/
```

3. Open the `http_em.conf` file with a text editor and change all occurrences of 4889 to the new port number you want to use.

4. Save and close the `http_em.conf` file.

5. Inform the DCM layer about the port change:

```
ORACLE_HOME/dcm/bin/dcmctl updateconfig -ct ohs
```

6. Locate the `emoms.properties` file in the same `sysman/config` directory.

7. Open the `emoms.properties` file with a text editor and change the following entry so it references the new port number of the Management Service:

```
oracle.sysman.emSDK.svlt.ConsoleServerPort=4889
```

8. Restart the Management Service.

9. Reconfigure each Management Agent on your managed hosts to use the new management port.

See Also: ["Configuring the Management Agent to Use a New Management Service"](#) on page 10-1

To change the default Management Service port to a *secure* port:

1. Stop the Management Service using:

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

2. Change the secure port using the following command:

```
ORACLE_HOME/bin/emctl secure oms -secure_port <newPortNo>
```

3. Inform the DCM layer about the port change:

```
ORACLE_HOME/dcm/bin/dcmctl updateconfig -ct ohs
```

4. Start the Management Service using:

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

10.2.3 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:

1. Stop the Management Service.
2. Locate the `emoms.properties` file where you installed the Management Service:

```
ORACLE_HOME/sysman/config/emoms.properties
```

3. Edit the `emoms.properties` file and add the appropriate lines:

- For the Execute Host Command, add the following line:

```
oracle.sysman.cmd.tgt.multiTarget.confirmExecuteHostCommand=true
```

- For Execute SQL, add the following line:

```
oracle.sysman.cmd.tgt.multiTarget.confirmExecutesQL=true
```

Note: The text in the commands is case-sensitive.

4. Save the changes and close the `emoms.properties` file.
5. Restart the Management Service.

Migrating from Previous Versions of Enterprise Manager

This chapter discusses the migration procedure used to move from a previous version of Oracle Enterprise Manager to the new Oracle Enterprise Manager 10g environment. This chapter contains the following topics:

- [Overview of the Enterprise Manager Migration Process](#)
- [Requirements for Migrating from Previous Versions of Enterprise Manager](#)
- [The Oracle Enterprise Manager 10g Migration Process](#)
- [Configuring Metric Thresholds](#)

11.1 Overview of the Enterprise Manager Migration Process

This chapter describes how to migrate from the following versions of Enterprise Manager:

- Oracle Enterprise Manager Release 2.2
- Oracle Enterprise Manager Release 9.0.1
- Oracle Enterprise Manager Release 9.2

Migrating your existing Enterprise Manager framework to the Oracle Enterprise Manager 10g environment involves two steps:

- Making targets within your managed environment monitorable using the new framework by installing Oracle Enterprise Manager 10g Management Agents on hosts that are running your managed targets
- Migrating information about users, privileges, groups, and preferred credentials from the old management repository to the new Oracle Enterprise Manager 10g Management Repository.

Once you have completed migrating to the new framework, you may wish to change the default metric thresholds for groups of managed targets within your enterprise. For more information, see [Configuring Metric Thresholds](#) on page 11-8.

11.2 Requirements for Migrating from Previous Versions of Enterprise Manager

Before beginning the migration process, ensure that the following list of requirements is satisfied:

- The previous version of complete Enterprise Manager Framework (Release 2.2, 9.0.1, or 9.2) must be up and running, including the Enterprise Manager Console, Oracle Management Server, Repository, and Intelligent Agents. The migration procedure uses the Job system in the previous version of Enterprise Manager to deploy the Oracle Enterprise Manager 10g Management Agents.
- The Oracle Enterprise Manager 10g Grid Control Console must be installed and running on one network host. Specifically, the Management Service must be up and running and available to the Oracle Enterprise Manager 10g Management Agents that you will install on your managed hosts.
- You must have the credentials for the Enterprise Manager Administrator Account for both the previous version of Enterprise Manager, as well as for Oracle Enterprise Manager 10g. Account read/write privileges are required for any machine currently running the Release 2.2, 9.0.1 or 9.2 Intelligent Agent.
- You must have the Database User and Password for the previous version of the Enterprise Manager Repository Database, as well as for the Oracle Enterprise Manager 10g Management Repository database.
- You must have 375 Megabytes of free disk space on each host where a Management Agent is to be installed.
- You must have installed the latest system and software patches for the Oracle Enterprise Manager 10g environment. Note that the system and software patch requirements for the Oracle Enterprise Manager 10g Grid Control Console are significantly different from previous versions of Enterprise Manager.

11.3 The Oracle Enterprise Manager 10g Migration Process

Migrating from a previous version of Enterprise Manager to the Oracle Enterprise Manager 10g Grid Control is a two-stage process. The following sections describe each stage in the process:

- [Deploying and Configuring Oracle Enterprise Manager 10g Management Agents](#)
- [Migrating Management Repository Data](#)

11.3.1 Deploying and Configuring Oracle Enterprise Manager 10g Management Agents

Deploying Oracle Enterprise Manager 10g Management Agents on machines running targets managed by an older version of Enterprise Manager makes these targets monitorable via Oracle Enterprise Manager 10g. To simplify and automate Management Agent deployment, a Tcl script is provided that is submitted as a job from an Enterprise Manager Release 2.2, Release 9.0.1, or Release 9.2 Job system. The deployment script (`agentInstallJob.tcl`) can be found in the Oracle Enterprise Manager 10g home directory at the following location:

```
%ORACLE_HOME/sysman/agent_download/agentInstallJob.tcl
```

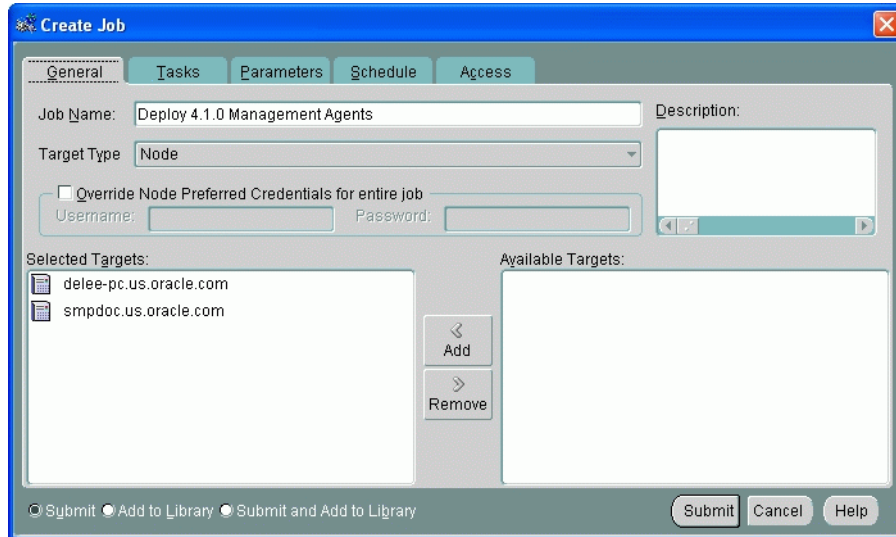
Deployment of the Oracle Enterprise Manager 10g Management Agent is carried out in two phases:

- [Deploying the Oracle Enterprise Manager 10g Management Agents Using the Release 2.2, Release 9.0.1, or Release 9.2 Job System](#)
- [Configuring the Oracle Enterprise Manager 10g Management Agents for Use with the Oracle Enterprise Manager 10g Job System \(UNIX Systems Only\)](#)

11.3.1.1 Deploying the Oracle Enterprise Manager 10g Management Agents Using the Release 2.2, Release 9.0.1, or Release 9.2 Job System

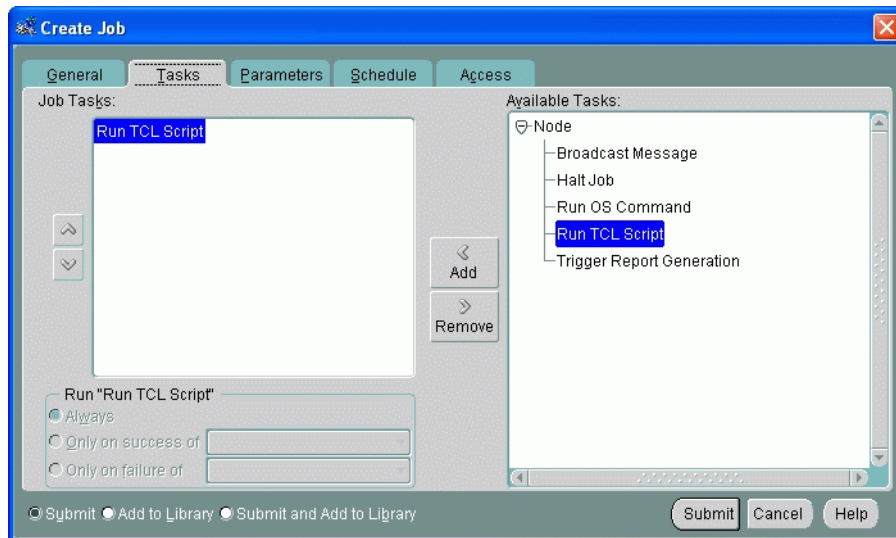
The `agentInstallJob.tcl` script must be run as a Tcl job from an Enterprise Manager Release 2.2, Release 9.0.1, or Release 9.2 Console. As shown in [Figure 11-1](#), you define the job by choosing a "Node" target type and then selecting the machines on which the Oracle Enterprise Manager 10g Management Agents are to be installed.

Figure 11-1 *Selecting Machines for Management Agent Deployment*

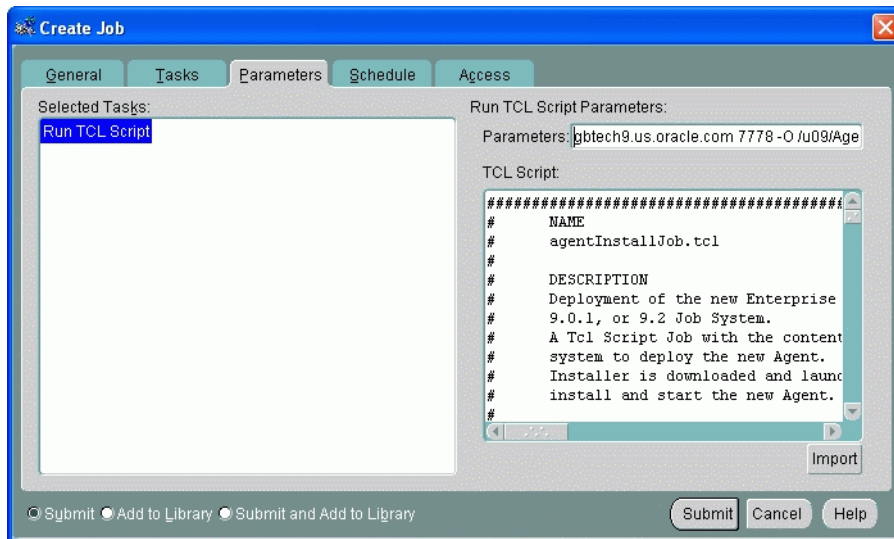


Once you have selected where the Management Agents are to be deployed, you need to define an installation task using `agentInstallJob.tcl`. As shown in [Figure 11-2](#), select the "Run TCL Script" task.

Figure 11-2 *Choosing the Run TCL Script Task*



The next step involves defining the functional core of the job. As shown in [Figure 11-3](#), you need to copy the content of the `agentInstallJob.tcl` script into the text entry area using either the Import function or manually copying and pasting the entire script into the TCL Script text entry area.

Figure 11–3 Copying agentInstallJob.tcl and Specifying Job Parameters

In addition to importing the script content, you must specify operational parameters required by the script to install the Oracle Enterprise Manager 10g Management Agent. As shown in [Figure 11–3](#), you enter these parameters in the **Run TCL Script Parameters** field. The parameters are:

- The Oracle Management Service host
Example: mgmthost1.acme.com
- HTTP Port Number
Example: 7778
- Directory Type (-o or -f)
Usage:
-o Identical installation directory structure on all machines.
-f Different installation directory structure on various machines (specified in text file)
- Directory Argument
Example: /u09/agent/agent_41

[Example 11–1](#) and [Example 11–2](#) show the format and syntax used to specify these parameters in the **Run TCL Script Parameters** field.

Example 11–1 Same Installation Directory Structure on All Machines

```
mgmthost1.acme.com 7778 -o /u09/Agent/Agent_41
```

Example 11–2 Different Installation Directory Structure on Different Machines

```
mgmthost1.acme.com 7778 -f hostname_lookup.txt
```

11.3.1.1.1 More About the Directory Type Parameter The Directory Type parameter offers two options either "-o" or "-f" plus the Directory Argument which consists of either a default directory (-o option) or host lookup file (-f option). As mentioned in the previous section, the "-o" option specifies that the same Management Agent home directory structure be created on all machines where the Management Agent is to be

installed. For example, if the `agentInstallJob.tcl` job is submitted against `MACHINE1`, `MACHINE2`, and `MACHINE3` using the following job parameters:

```
mgmthost1.acme.com 7778 -o /u09/agent/agent_41
```

The `agentInstallJob.tcl` script will create the `/u09/agent/agent_41` directory on each of the three machines. Once created, this directory is used by the Oracle Universal Installer (OUI) as an installation staging area. This directory eventually becomes the Oracle Enterprise Manager 10g Management Agent Home.

Note: The `agentInstallJob.tcl` job runs OUI in silent mode to perform the actual Management Agent installation operations.

In contrast, the `-f` option specifies that different installation/Agent Home directory structures be created for specific machines. Before creating a TCL job with this option you must first create a flat text file listing each machine name and corresponding directory for that host. The flat file **MUST** reside in the Oracle Enterprise Manager 10g Management Service Home in the following directory:

```
OMS_HOME/sysman/agent_download
```

In the following example, the lookup file parsed by the TCL job is named `hostname_lookup.txt`.

```
mgmthost1.acme.com 7778 -f hostname_lookup.txt
```

When a TCL job is submitted using the `-f` option, the job first obtains the name of the target machine by executing the `hostname` command. The result is then compared against entries in the `hostname` lookup file. If the `hostname` is found in the file, the associated directory structure is used. If the `hostname` is not found, then the directory structure specified for the `"wildcard"` character is used. The wildcard can be used as a default entry in case the TCL job cannot locate a particular `hostname` within the file. A wildcard entry is designated by a `"*"` and must be the last entry as the file is parsed from top to bottom.

[Example 11-3](#) shows the format for a sample `hostname` lookup file. In the example, you have 20 machines in your enterprise where you want Oracle Enterprise Manager 10g Management Agents installed. You want the same Agent Home directory structure created on all machines except `HOST1`, `HOST2`, and `HOST3`.

Example 11-3 Sample Hostname Lookup File

```
HOST1/oracle_home1/agent/agent_install
HOST2/ora_host2/agent_install
HOST3/orahome_host3/agent/install
*/ora_agent/agent/install
```

The TCL job (submitted to `HOST2`) runs the `hostname` command and receives `HOST2` as the output. This output is then cross-referenced with all entries within the `hostname_lookup.txt` file. Since `HOST2` is an entry in the `hostname_lookup.txt` file, the TCL job knows to create the Oracle Enterprise Manager 10g Management Agent Home in `/ora_host2/agent_install`. `HOST1`, `HOST2`, and `HOST3` will have unique directories. The Management Agent home directory for the remaining 17 machines will be `/ora_agent/agent/install`.

Because the TCL job, or more specifically OUI, creates files and directories on the target machines, full read/write privileges for the Enterprise Manager administrator account running the job are required. When the installation is complete the new Oracle

Enterprise Manager 10g Management Agent is started automatically and begins retrieving host, database, and listener information. This information is then uploaded via HTTP or HTTPS to the new Enterprise Manager 10g management repository where it becomes available to the Grid Control for viewing.

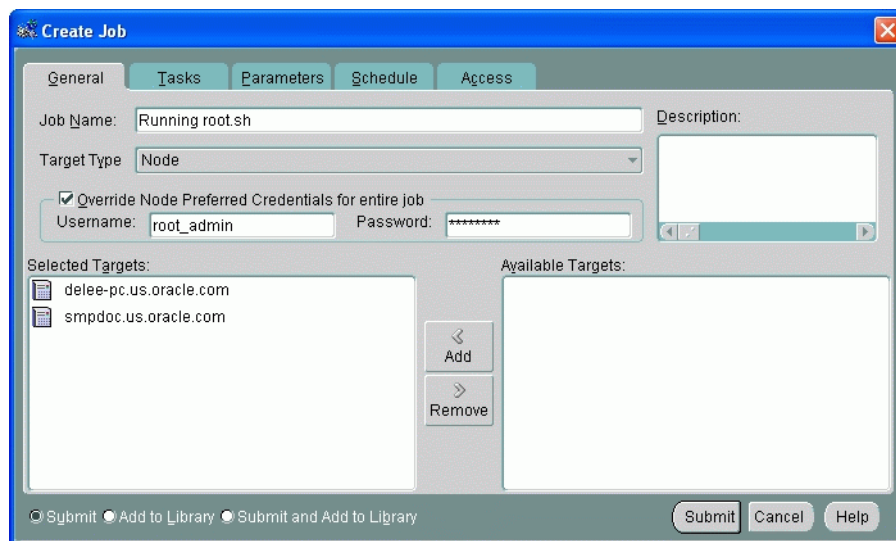
11.3.1.2 Configuring the Oracle Enterprise Manager 10g Management Agents for Use with the Oracle Enterprise Manager 10g Job System (UNIX Systems Only)

Once the Oracle Enterprise Manager 10g Management Agents are operational, you need to configure each Management Agent for use with the Oracle Enterprise Manager 10g job system. This step consists of running the `root.sh` script on each machine where the Management Agent is installed. This script is located in the Management Agent Home of the host machine. Specifically, the `root.sh` script grants root privileges to the Oracle Enterprise Manager 10g Management Agent. Therefore, the root user and password for that machine are required in order to run `root.sh`.

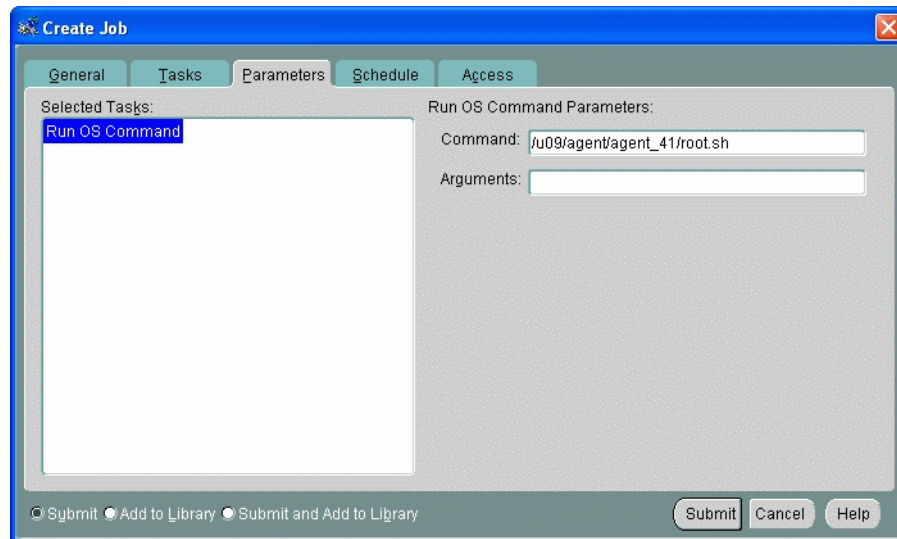
As with the `agentInstallJob.tcl` script you can automate this task by running the `root.sh` script using the Enterprise Manager Release 2.2, Release 9.0.1, or Release 9.2 Job system. To do this, you create an "OS command" job that executes "root.sh" on all machines requiring Management Agent configuration. As shown in the [Figure 11-4](#), preferred credentials should be overridden by the root user and password of the target host.

Note: The Preferred Credential Override is available with Enterprise Manager Release 9.2 systems only. For older versions of Enterprise Manager, you must run the job as a user with preferred credentials allowing root access.

Figure 11-4 Overriding Preferred Credentials



On the job Parameters page ([Figure 11-5](#)), specify "root.sh" in the **Command** field as shown in the following figure and submit the job for execution. You must specify the full path to the `root.sh` script. For example: `/u09/agent/agent_41/root.sh`

Figure 11–5 Command Parameters

The Oracle Enterprise Manager 10g Management Agent can be up and running when `root.sh` is executed and does not need to be restarted after the configuration process has been completed.

Note: In previous releases of Enterprise Manager, job system configuration was part of the Intelligent Agent install. With the Enterprise Manager 10g release, this configuration is now a separate step due to architectural differences between the old and new frameworks.

11.3.2 Migrating Management Repository Data

Once the Oracle Enterprise Manager 10g Management Agents have been deployed and configured, the next step is to migrate information about users, privileges, groups, and preferred credentials from the original Management Repository to the Oracle Enterprise Manager 10g Management Repository.

Note: Privileges, group membership, and preferred credentials are migrated for databases, listeners, and hosts only.

Both Enterprise Manager 9*i* and Oracle Enterprise Manager 10g save and encrypt all administrator accounts and preferred credentials in the repository. In order to migrate all of these accounts over to Enterprise Manager 10g, you must run the Migration Utility from the Enterprise Manager 10g home. This command line utility can be found in the following directory:

```
%EM_HOME%\bin\repo_mig
```

The Migration Utility requires the repository user and password for both the original Management Repository database and for the new Oracle Enterprise Manager 10g Management Repository database. You execute the utility and specify operational parameters using the following format:

```
repo_mig -preview|-migrate source_user/source_pwd@source_service dest_user/dest_pwd@dest_service
```

where:

- `-preview`: Generates a preliminary migration report without carrying out the migration.
- `-migrate`: Performs migration of groups, administrators, target privileges, and preferred credentials of hosts, databases, and listeners.
- `source_user`: Source OEM repository user name
- `source_pwd`: Source OEM repository password
- `source_service`: Source OEM repository service. For example, `Host:Port:SID`
- `dest_user`: Destination OEM repository user name
- `dest_pwd`: Destination OEM repository password
- `dest_service`: Destination OEM repository service. (`Host:Port:SID`)

Once the migration is complete, the account information is then saved and encrypted. The passwords on all of the accounts will remain the same.

11.4 Configuring Metric Thresholds

As mentioned previously, migration only involves transferring users, privileges, groups, and preferred credentials to the Enterprise Manager 10g framework; any older event test thresholds that existed in pre-10g versions of Enterprise Manager will not be transferred.

Enterprise Manger 10g provides out-of-the-box monitoring that simplifies a critical but potentially time-consuming task of setting up monitoring for managed targets. As you add targets to Enterprise Manager, options are automatically provided to monitor the target at a recommended or at a minimum level. Each level of monitoring consists of a set of metrics and predefined thresholds that are based on Oracle recommendations for those levels. You may choose to use these Oracle recommendations, or you can change these thresholds to suit your particular needs.

11.4.1 Copying Metric Thresholds to Multiple Targets

Your Enterprise Manager Grid Control installation may be monitoring a very large number of targets, making it inconvenient to manually change metric threshold values for each monitored target. Enterprise Manager Grid Control 10.2 provides an easy way to propagate metric thresholds to any number of targets via monitoring templates.

Monitoring templates allow you to define metric settings once and apply these settings to any number of targets of the same type. For more information about monitoring templates, see Enterprise Manager Grid Control online help.

Configuring 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](#)
- [Assigning Methods to Rules](#)
- [Assigning Rules to Methods](#)
- [Management Information Base \(MIB\)](#)
- [Troubleshooting Notifications](#)

12.1 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.

12.1.1 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 12–1). 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 12–1 shows sample notification method entries.

Example 12–1 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

Figure 12–1 Defining a Mail Server

The screenshot shows the Oracle Enterprise Manager 10g interface for configuring Notification Methods. The page title is "Notification Methods - Netscape". The navigation bar includes "Home", "Targets", "Deployments", "Alerts", "Policies", "Jobs", and "Reports". The main content area is titled "Notification Methods" and contains the following sections:

- Notification Methods:** A descriptive paragraph explaining that these methods are used for sending e-mail, SNMP traps, and running custom scripts.
- Mail Server:** A section for configuring an outgoing mail (SMTP) server. It includes fields for "Outgoing Mail (SMTP) Server", "User Name", "Password", "Confirm Password", "Identify Sender As", and "Sender's E-mail Address". There are also "Revert", "Apply", and "Test Mail Servers" buttons.
- Scripts and SNMP Traps:** A section for defining methods for sending notifications via OS commands, PL/SQL procedures, or SNMP traps. It includes an "Add" button and a dropdown menu set to "OS Command".
- Select Name / Type Table:** A table with two columns: "Select Name" and "Type". It currently shows "No notification methods found."
- TIP:** A tip reminding the user to create Notification Rules to use these methods.

The footer contains copyright information and navigation links: "Home | Targets | Deployments | Alerts | Policies | Jobs | Reports | Setup | Preferences | Help | Logout".

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 12–2 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 12-2](#), 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.

12.1.2 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.

12.1.2.1 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 12-3](#) shows a typical notification that uses the long format.

The Short Format ([Example 12-4](#)) 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 12-3 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 12-4 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 paging, SMS, etc., 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

```

12.1.2.2 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.

12.1.2.3 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 Methods 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 12–1](#) lists all the default notification rules. These are all owned by the SYSMAN user and are public rules.

Table 12–1 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 (%)
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 (%)

Table 12–1 (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
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 (%)
Listener Availability	System-generated notification rule for monitoring database Listeners' availability, and critical metric statuses.	Listeners	Down	N/A
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
Repository Operations Availability	System-generated notification rule for monitoring the availability of the DBMS jobs that are part of the Management Repository.	Oracle Management Service and Repository	Critical	DBMS Job UpDown
Violation Notification for Database Security Policies	System-generated notification rule for monitoring the secureness of the database configuration.	Databases	Critical	N/A

Table 12–1 (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
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 Methods page and select the Send Me E-Mail option.

12.1.3 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 must 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 specified the e-mail addresses 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 admin 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.

12.2 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.

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.

12.2.1 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.

12.2.1.1 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 12-13
- ["Passing Corrective Action Execution Status to an OS Command or Script"](#) on page 12-21
- ["Passing Job Execution Status to an OS Command or Script"](#) on page 12-26

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 12-5](#) 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 12-5 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 12-11.

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 12-6](#) shows information required for the notification method.

Example 12-6 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 12-28 or ["Assigning Methods to Rules"](#) on page 12-27.

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 12–2 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
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.

Table 12-2 (Cont.) Environment Variables

Environment Variable	Description
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.
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 12-7](#) 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 12-7 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 12-8](#) shows an OS script that logs alert information to the file 'alertmsg.txt'. The file is saved to the /u1/results directory.

Example 12-8 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 12–9](#) 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 `opcmsg` and forwarding the information to the HP OpenView management server.

Example 12–9 HP OpenView Script

```

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

```

12.2.1.2 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 12-17

["Passing Corrective Action Status Change Information"](#) on page 12-21

["Passing Job Execution Status Information"](#) on page 12-24

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 specification:

```
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 12-11, "PL/SQL Procedure Using a Severity Code"](#). The procedure initially retried after one minute, then two minutes, then three minutes and so on, until the notification is a day old, at which point it will be purged.

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

Log in as a Super Administrator, 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 12-15.

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 12-10](#).

Example 12-10 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 12-28 or ["Assigning Methods to Rules"](#) on page 12-27.

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 12-3 Metric Severity Attributes

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

Table 12-3 (Cont.) Metric Severity Attributes

Attribute	Datatype	Additional Information
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 12-4 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 12-4 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 12-11 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, occurred)
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;

```

12.2.1.3 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 12-29.

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 12-12](#), the SNMP host will receive your SNMP traps.

Example 12-12 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 12-13](#). Each piece of information is sent as a variable embedded in the SNMP Trap.

Example 12-13 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 Rules to Methods](#)" on page 12-28 or "[Assigning Methods to Rules](#)" on page 12-27.

12.3 Passing Corrective Action Status Change Information

Passing corrective action status change attributes (new status, job name, job type, notification rule, or rule owner, etc.) 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 (Problems, Aborted, etc.) to the script to open a trouble ticket and escalate the problem.

12.3.1 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 12-5 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'.

Table 12–5 (Cont.) Environment Variables

Environment Variable	Description
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.
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.

12.3.2 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 12–6 lists all corrective action status change attributes that can be passed:

Table 12–6 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 12–7, "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.

Table 12–6 (Cont.) Corrective Action Status Attributes

Attribute	Datatype	Additional Information
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.
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 12–4, "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 12–7 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

Table 12-7 (Cont.) Corrective Action Status Codes

Name	Datatype	Value
COMPLETED_STATUS	NUMBER(2)	5
SUSPENDED_STATUS	NUMBER(2)	6
AGENTDOWN_STATUS	NUMBER(2)	7
STOPPED_STATUS	NUMBER(2)	8
SUSPENDED_LOCK_STATUS	NUMBER(2)	9
SUSPENDED_EVENT_STATUS	NUMBER(2)	10
SUSPENDED_BLACKOUT_STATUS	NUMBER(2)	11
STOP_PENDING_STATUS	NUMBER(2)	12
SUSPEND_PENDING_STATUS	NUMBER(2)	13
INACTIVE_STATUS	NUMBER(2)	14
QUEUED_STATUS	NUMBER(2)	15
FAILED_RETRIED_STATUS	NUMBER(2)	16
WAITING_STATUS	NUMBER(2)	17
SKIPPED_STATUS	NUMBER(2)	18
REASSIGNED_STATUS	NUMBER(2)	20

Example 12-14 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;
    
```

12.4 Passing Job Execution Status Information

Passing job status change attributes (new status, job name, job type, notification rule, or rule owner, etc.) 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 (Problems, Aborted, etc.) to the script to open a trouble ticket and escalate the problem.

12.4.1 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 12–8 lists all corrective action status change attributes that can be passed:

Table 12–8 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 12–9 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

Table 12–9 (Cont.) Job Status Codes

Name	Datatype	Value
STOPPED_STATUS	NUMBER(2)	8
SUSPENDED_LOCK_STATUS	NUMBER(2)	9
SUSPENDED_EVENT_STATUS	NUMBER(2)	10
SUSPENDED_BLACKOUT_STATUS	NUMBER(2)	11
STOP_PENDING_STATUS	NUMBER(2)	12
SUSPEND_PENDING_STATUS	NUMBER(2)	13
INACTIVE_STATUS	NUMBER(2)	14
QUEUED_STATUS	NUMBER(2)	15
FAILED_RETRIED_STATUS	NUMBER(2)	16
WAITING_STATUS	NUMBER(2)	17
SKIPPED_STATUS	NUMBER(2)	18
REASSIGNED_STATUS	NUMBER(2)	20

Example 12–15 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;

```

12.4.2 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 12–10 Environment Variables

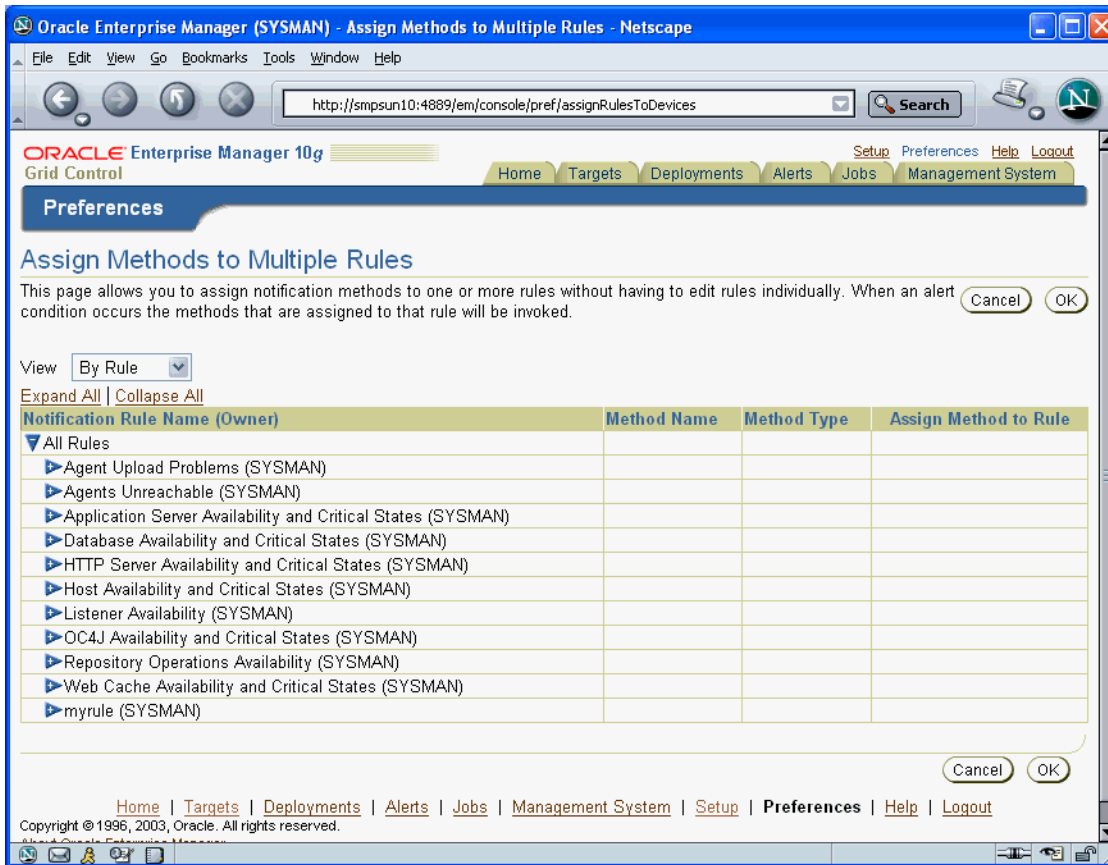
Environment Variable	Description
JOB_NAME	The name of the job.
JOB_OWNER	The owner of the job.
JOB_TYPE	The type of job.
JOB_STATUS	The job status.
TIMESTAMP	Time when the severity occurred.
NUM_TARGETS	The number of targets.
TARGET_NAME _n	The name of the <i>n</i> th target. For example, TARGET_NAME1, 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.

12.5 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 12–2 Assigning Methods to Rules



12.6 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 12–3 Assign Rules to Methods



12.7 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.

12.7.1 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.

12.7.2 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.

12.7.2.1 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.

12.7.2.2 MIB Definition

[Example 12-16](#) shows a typical MIB definition used by Enterprise Manager.

Example 12–16 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

```

12.8 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.

12.8.1 General Setup

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

OS Command, PL/SQL and SNMP Trap Notifications

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

E-mail Notifications

- Make sure an e-mail gateway is set up under the Notification Methods page of Setup. The sender's e-mail address should be valid. Clicking the Test button will send an e-mail to the sender's e-mail address. Make sure this e-mail is received. Note that the Test button ignores any Notification Schedule.
- Make sure an e-mail address is 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.

12.8.2 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.

12.8.3 Notification System Trace Messages

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

Tracing is configured by setting the following flag in `sysman/config/emomslogging.properties` file. You can set the trace level to INFO, WARN, DEBUG. For example,

```
log4j.em.notification=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
```

12.8.4 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: rgmemeasmtplib.oraclecorp.com: Sending failed;
nested exception is:
javax.mail.MessagingException: 550 5.7.1 <smtpemailtest_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

12.8.5 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 (stacbl0.us.oracle.com_Management_Service) (SYSMAN, myrule)

No execute permission on executable:

Error calling /bin/myscript: java.io.IOException: /bin/myscript: cannot execute (stacbl0.us.oracle.com_Management_Service) (SYSMAN, myrule)

Timeout because OS Command ran too long:

Timeout occurred running /bin/myscript (stacbl0.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.

12.8.6 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.

12.8.7 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](#)

13.1 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.

13.2 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](#)

13.2.1 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](#)

13.2.1.1 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 13–1 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.

13.2.1.2 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 13–2 *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 13–2 (Cont.) (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.

13.2.1.3 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 13-5. 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.

13.2.2 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.

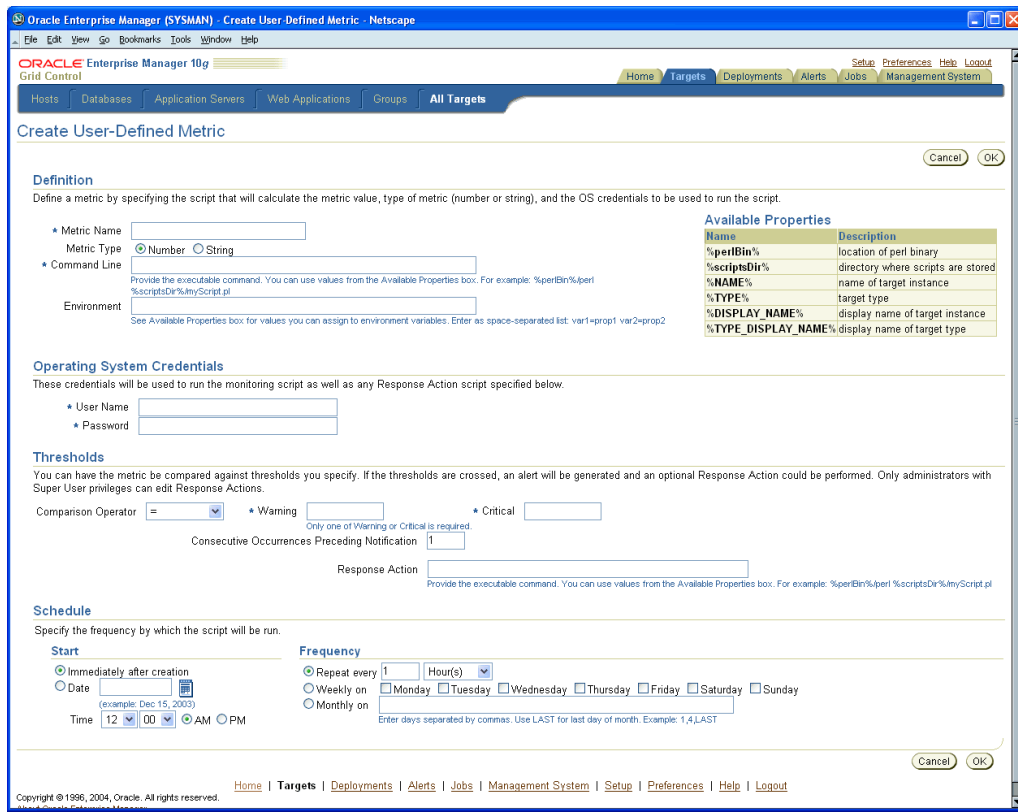
Create User-Defined Metric Page (OS-based User-Defined Metric)

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 13-1](#).

Figure 13–1 Create User-Defined Metric Page (OS-Based)



Key elements of this page are described in the following tables.

Table 13–3 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 13–3 (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 13–4 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 13–5 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 Value</th> <th>Metric Type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>=</td> <td>Number</td> <td>equal to</td> </tr> <tr> <td>></td> <td>Number</td> <td>greater than</td> </tr> <tr> <td><</td> <td>Number</td> <td>less than</td> </tr> <tr> <td>>=</td> <td>Number</td> <td>greater than or equal to</td> </tr> <tr> <td><=</td> <td>Number</td> <td>less than or equal to</td> </tr> <tr> <td>!=</td> <td>Number</td> <td>not equal to</td> </tr> <tr> <td>CONTAINS</td> <td>String</td> <td>contains at least</td> </tr> <tr> <td>MATCH</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 13-5 (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.

13.2.3 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.

13.3 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 13-2](#).

Figure 13–2 Create User-Defined Metric Page (SQL-Based)



Key elements of this page are described in the following tables.

Table 13–6 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.
SQL Query Output	Specify whether the SQL script is to return a single value (one column) or a multiple rows (two columns). <ul style="list-style-type: none"> ■ Single Value: Query is one of the following types. <ul style="list-style-type: none"> A <i>SELECT statement</i> returning a single value. Example: <code>SELECT sal FROM emp WHERE empno=7369</code> A <i>function call</i> returning a single value. Example: <code>myfunc(123, 'abc')</code> ■ Two Columns: Query is a <i>SELECT statement</i> 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.

Table 13–6 (Cont.) Create User-Defined Metric Page: Definition

User-Interface Element	Description
SQL Query	Enter a SQL query or function call that returns values of the appropriate type (STRING or NUMBER). 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.

Table 13–7 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 13–8 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																																		
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> <p><query_value> <comparison_operator> <warning_threshold></p> <p>and if the consecutive occurrences preceding notification has been reached.</p>																																				

Table 13–8 (Cont.) Create User-Defined Metric Page: Threshold

User-Interface Element	Description
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 The warning threshold is set to 250 for SMITH, 400 for JONES, and 900 for CLARK. ■ Critical threshold by key: SMITH:100;JONES:200;CLARK:500 The critical threshold is set to 100 for SMITH, 200 for JONES, and 500 for CLARK. <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>
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 AND 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>

Table 13–8 (Cont.) Create User-Defined Metric Page: Threshold

User-Interface Element	Description
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.

13.3.1 SQL-Based User-Defined Metric Examples

For a database version 9i and higher, you can run the example queries as dbsnmp, 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 dbsnmp the following privileges in order for the queries to run successfully:

For example #1:

```
grant select on sys.dba_tablespaces to dbsnmp;
grant select on sys.dba_data_files to dbsnmp;
grant select on sys.dba_free_space to dbsnmp;
```

For example #2:

```
grant select on sys.dba_extents to dbsnmp;
```

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.

13.3.1.1 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')
```

13.3.1.2 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
```

13.4 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.

13.4.1 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 for SQL-based User-Defined Metrics.

5. On the Metrics page, click Add. A list of available metrics appears. To view all metrics on 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).
 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.
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.

13.4.2 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.

13.4.3 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.

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. Under the section "Metrics with multiple thresholds", choose the option 'Duplicate threshold settings on 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.

Additional Configuration Tasks

This chapter contains the following sections:

- [Understanding Default and Custom Data Collections](#)
- [Enabling Multi-Inventory Support for Configuration Management](#)
- [Manually Configuring a Database Target for Complete Monitoring](#)
- [Modifying the Default Login Timeout Value](#)
- [Configuring Clusters and Cluster Databases in Grid Control](#)
- [Collecting Client Configurations](#)

14.1 Understanding Default and Custom Data Collections

When you install the Oracle Management Agent on a host computer, Enterprise Manager automatically begins gathering a default set of metrics that you can use to monitor the performance and availability of each targets on that host. For some of these target metrics, Enterprise Manager provides default threshold settings that determine when you will be notified that there is a problem with the metric.

See Also: "About Alerts" in the Enterprise Manager online help

For selected metrics, you can customize the default thresholds. When you make these types of customizations, Enterprise Manager saves the new settings in a file on the local disk. The following sections provide more information about how these settings are saved:

- [How Enterprise Manager Stores Default Collection Information](#)
- [Restoring Default Collection Settings](#)

14.1.1 How Enterprise Manager Stores Default Collection Information

Enterprise Manager stores the default collection criteria for each target in the following location on each Oracle Management Agent host:

`AGENT_HOME/sysman/admin/default_collection/`

For some targets, you can use the Oracle Enterprise Manager 10g Grid Control Console to modify the default metric collection settings. For example, you can modify the default thresholds for your host targets. When you make these types of modifications, Enterprise Manager creates a new default collection file in the following directory:

`AGENT_HOME/sysman/emd/collection/`

This collection file overrides the default collection information stored in the `sysman/admin/default_collection` directory.

14.1.2 Restoring Default Collection Settings

If you have made modifications to the metric thresholds for a particular target, you can restore the default metric collection settings by deleting the customized collection information in the `sysman/emd/collection` directory.

For example, if you want to restore the default collections for a particular database target, remove the customized collection file for that target from the `sysman/emd/collection` directory. Enterprise Manager will begin using the metric collection information stored in the `sysman/admin/default_collection` directory.

14.2 Enabling Multi-Inventory Support for Configuration Management

Every time you install an Oracle software product on a host computer, Oracle Universal Installer saves information about the software installation on your hard disk. The directories and files that contain this software configuration information are referred to as the Oracle Universal Installer inventory.

See Also: *Oracle Universal Installer and OPatch User's Guide*

When you use Enterprise Manager to monitor your Oracle software installations, Enterprise Manager takes advantage of information saved in the Universal Installer inventory.

As it gathers information about the configuration of your host computer, by default, Enterprise Manager assumes that you have one Oracle Universal Installer inventory on the host. Specifically, Enterprise Manager recognizes the inventory that Oracle Universal Installer uses when you run the Universal Installer on the host.

However, in some cases, you may have more than one inventory. For example, you may have worked with Oracle Support to clone your Oracle software installations. For those cases, you can use the following procedure to be sure that Enterprise Manager can track and manage the software information in multiple inventories on the same host.

Caution: Enabling support for multiple inventories is optional and available only for advanced users who are familiar with the Oracle Universal Installer inventory architecture and who have previously worked with multiple inventories on a managed host. This procedure is not required for hosts where normal installations have been performed.

To set up Enterprise Manager so it can read multiple inventories on a host:

1. Locate the `OUIinventories.add` file in the following directory:

```
$ORACLE_HOME/<nodename>_<sid>/sysman/config
```

The Management Agent state listed in this example represents an installation for Database Control. For more information about the Management Agent state to use for other installations, see [Section 14.2.1, "AGENT_HOME Versus AGENT_STATE Directories"](#) on page 14-3.

2. Open `OUIinventories.add` using a text editor.

Instructions within the file describe the format to use when identifying multiple inventories, as well other techniques for mapping Oracle Homes.

3. Carefully review the instructions within the file.
4. Add entries to the file for each additional inventory on the managed host.
5. Save your changes and close the file.

During its next collection of host configuration information, Enterprise Manager will start gathering software configuration information from the inventories that you identified in the `OUIinventories.add` file, in addition to the default inventory that Enterprise Manager normally collects.

Alternatively, to see the data gathered from the additional inventories before the next regularly-scheduled collection, navigate to the Host home page in the Grid Control Console, click the **Configuration** tab, and click the Refresh Data icon next to the page timestamp.

Note: If there any irrecoverable problems during the collection of the default inventory (for example, if the inventory file or directory protections prevent Enterprise Manager from reading the inventory), inventories listed in `OUIinventories.add` file are also not collected.

If the Enterprise Manager is able to read the default inventory, but there is a problem reading an additional inventory listed in the `OUIinventories.add` file, Enterprise Manager issues a collection warning for those inventories. However, Enterprise Manager does collect the configuration information for the other inventories.

14.2.1 AGENT_HOME Versus AGENT_STATE Directories

The Management Agent recognizes two main directory structures; its installation directory where software binaries and all unchanging metadata are stored, and its configuration/state directory where all customizations and output/log content are stored and/or generated. In a normal Management Agent installation, these two directories are the same. However, they can be different in the following cases:

- RAC Agent installation (`$ORACLE_HOME` versus `$ORACLE_HOME/<hostname>`)
- Database Control installation (`$ORACLE_HOME` versus `$ORACLE_HOME/<nodename><sid>`)
- State-only Management Agent deployment (using the `emctl deploy agent` command -- `$ORACLE_HOME` versus `$EMSTATE`)

In each of the above cases, there will be multiple instances of the Management Agent running off the same binaries installation. The different instances have different locations to maintain separate configurations but use the same set of binaries. The command `emctl agent status` provides the values of the Management Agent's binaries and state locations.

14.3 Manually Configuring a Database Target for Complete Monitoring

When you first discover an Oracle Database 10g target, you should check the monitoring credentials to be sure the password for the DBSNMP database user account is set correctly in the database target properties.

See Also: ["Specifying New Target Monitoring Credentials"](#) on page 2-13

Besides setting the monitoring credentials, no other configuration tasks are required to monitor an Oracle Database 10g target.

However, when you monitor an Oracle9i database or an Oracle8i database, there is some additional configuration required if you want to monitor certain types of database performance metrics using the Grid Control Console.

To monitor these additional performance metrics Enterprise Manager requires that Oracle Statspack and some additional Enterprise Manager packages be installed and configured in the database you are monitoring.

See Also: ["Using Statspack"](#) in *Oracle Database Performance Tuning Guide and Reference* in the Oracle9i Documentation Library

If these additional objects are not available and configured in the database, Enterprise Manager will not be able to gather the data for the complete set of performance metrics. In addition, Enterprise Manager will not be able to gather information that otherwise could be readily available from the Database home page, such as Bad SQL and the Top SQL Report.

You can use the Configure Database wizard in the Grid Control Console to install the required packages into the database, or you can use the following manual procedure.

See Also: ["Modifying Target Properties"](#) in the Enterprise Manager online help for information on configuring managed targets, including database targets

To manually install Statspack and the other required database objects into an Oracle9i database that you are managing with Enterprise Manager, you can use SQL*Plus and the following procedure:

1. Log in to the database host using an account with privileges that allow you to write to the database home directory and to the Management Agent home directory.

For each of the commands in this procedure, replace AGENT_HOME with the actual path to the Oracle Management Agent home directory and replace ORACLE_HOME with the path to the database home directory.

2. Start SQL*Plus and connect to the database using the SYS account with SYSDBA privileges.

For example:

```
$PROMPT> ./sqlplus "connect / as sysdba"
```

3. Enter the following command to run the database dbmon script:

```
SQL> @AGENT_HOME/sysman/admin/scripts/db/config/dbmon
```

The script will display the following prompt:

Enter value for dbm_password:

4. When prompted, enter the password for the DBSNMP account.

The script performs several configuration changes and returns you to the SQL*Plus prompt.

5. Connect as the DBSNMP user.

For example:

```
SQL> connect DBSNMP
```

6. Enter the following command:

```
SQL> @AGENT_HOME/sysman/admin/scripts/db/config/response.plb
SQL> grant EXECUTE on dbsnmp.mgmt_response to OEM_MONITOR;
```

7. Connect as SYS and enter the following command to create the PERFSTAT user:

```
SQL> @ORACLE_HOME/rdbms/admin/spcreate
```

Note: The spcreate script will prompt you for a default tablespace and default temporary tablespace for the PERFSTAT user. Do not specify the SYSTEM tablespace as the default tablespace for the PERFSTAT user. For more information, see "Using Statspack" in the *Oracle Database Performance Tuning Guide*.

8. Connect as the PERFSTAT user.

For example:

```
SQL> connect PERFSTAT;
```

9. Enter the following commands from the PERFSTAT user account:

```
SQL> define snap_level='6';
SQL> define cinterval='1';
SQL> define cjobno='-1';
SQL> @AGENT_HOME/sysman/admin/scripts/db/config/spset
```

10. Connect as the SYS user and enter the following command:

```
SQL> grant OEM_MONITOR to dbsnmp;
```

11. If the database you are modifying is an Oracle8i database, also enter the following commands as the SYS user:

```
grant select on sys.ts$ to OEM_MONITOR;
grant select on sys.seg$ to OEM_MONITOR;
grant select on sys.user$ to OEM_MONITOR;
grant select on sys.obj$ to OEM_MONITOR;
grant select on sys.sys_objects to OEM_MONITOR;
grant select on sys.file$ to OEM_MONITOR;
grant select on sys.attrcol$ to OEM_MONITOR;
grant select on sys.clu$ to OEM_MONITOR;
grant select on sys.col$ to OEM_MONITOR;
grant select on sys.ind$ to OEM_MONITOR;
grant select on sys.indpart$ to OEM_MONITOR;
grant select on sys.indsubpart$ to OEM_MONITOR;
grant select on sys.lob$ to OEM_MONITOR;
grant select on sys.lobfrag$ to OEM_MONITOR;
```

```
grant select on sys.partobj$ to OEM_MONITOR;
grant select on sys.tab$ to OEM_MONITOR;
grant select on sys.tabpart$ to OEM_MONITOR;
grant select on sys.tabsubpart$ to OEM_MONITOR;
grant select on sys.undo$ to OEM_MONITOR;
```

12. For any supported database version, enter the following command from the SYS account:

```
SQL> show parameter job_queue_processes
```

If the output from the `show parameter` command is zero, then perform the following steps to modify the `job_queue_processes` initialization parameter:

If you start the database using an spfile, enter the following command:

```
SQL> alter system set job_queue_processes = 2 SCOPE=BOTH;
```

Otherwise, do the following:

- a. Enter the following command:

```
SQL> alter system set job_queue_processes = 2;
```

- b. Exit SQL*PLUS and update the `init.ora` database configuration file with the following entry so the parameter will be applied whenever the database is restarted:

```
job_queue_processes=2
```

13. Exit SQL*Plus and change directory to the following directory in the home directory of the Management Agent that is monitoring the database:

```
AGENT_HOME/bin
```

14. Reload the Management Agent by entering the following command:

```
$PROMPT> ./emctl agent reload
```

15. Using the Grid Control Console, return to the Database home page and verify that the Bad SQL and Top SQL Report metrics are now being gathered.

14.4 Modifying the Default Login Timeout Value

To prevent unauthorized access to the Grid Control Console, Enterprise Manager will automatically log you out of the Grid Control Console when there is no activity for a predefined period of time. For example, if you leave your browser open and leave your office, this default behavior prevents unauthorized users from using your Enterprise Manager administrator account.

By default, if the system is inactive for 45 minutes or more, and then you attempt to perform an Enterprise Manager action, you will be asked to log in to the Grid Control Console again.

Caution: As stated in the previous paragraphs, the timeout value for logging in to the Grid Control Console is defined in order to protect your system from unauthorized logins. If you make changes to the login timeout value, be sure to consider the security implications of leaving your session open for other than the default timeout period.

To increase or decrease the default timeout period:

1. Change directory to the following location in the Oracle Application Server home directory where the Management Service was deployed:

```
IAS_HOME/sysman/config/
```

2. Using your favorite text editor, open the `emoms.properties` file and add the following entry:

```
oracle.sysman.eml.maxInactiveTime=time_in_minutes
```

3. For example, if you want to change the default timeout period to one hour, add the following entry:

```
oracle.sysman.eml.maxInactiveTime=60
```

4. Save and close the `emoms.properties` file.
5. Restart the Management Service.

Note: The default timeout value does not apply when you restart the Web server or the Oracle Management Service. In both of those cases, you will be asked to log in to the Grid Control Console, regardless of the default timeout value.

14.5 Configuring Clusters and Cluster Databases in Grid Control

This section describes how to configure clusters, cluster databases, and discovering instances.

14.5.1 Configuring Clusters

To add a cluster target that was installed but not discovered as a target automatically during installation, perform the following steps:

1. Click **All Targets** from the Targets page.
2. Select **Cluster** from the Add menu and click **Go**. The Add Target: Cluster page appears.
3. Optional: Specify the cluster name and provide the Clusterware home path if it is installed on the cluster.
4. To add hosts to the cluster, use the arrow buttons to move items from Available Hosts to Selected Hosts. The hosts you select must not already belong to a cluster.
5. Click **Add** to save the cluster target to the `targets.xml` file on every selected host.

See Also: The Enterprise Manager online help for more information about configuring clusters

14.5.2 Configuring Cluster Databases

After you have added the cluster target, you can add a cluster database target either from the Databases page or from the All Targets page.

To add a cluster database target, perform the following steps:

1. In the Enterprise Manager Grid Control Console, select one of the following entry locations:

- From the Databases page, click **Add**. The Add Database Instance Target: Specify Host page appears.
 - From the All Targets page, select **Database Instance** from the Add drop-down menu, then click **Go**. The Add Database Instance Target: Specify Host page appears.
2. Specify any host member of the cluster target where the cluster databases reside, then click **Continue**. The Add Database: Specify Source page appears.
 3. Keep the default option (on all hosts in the cluster) selected and click **Continue**. This option sends requests to all Management Agents in the cluster to perform discovery.

After target discovery completes, the newly discovered RAC databases appear in the Targets Discovered on Cluster page. If the databases do not appear, see the Troubleshooting section below.
 4. If the desired targets do not appear in the Cluster Databases table, or if the discovered targets are not configured appropriately, click **Manually Add**. The Properties page of the Configure Cluster Database wizard appears.
 5. Provide the required values for the Properties table.
 6. You must specify at least one instance in the Instances table. If no instances appear in the table, click **Add**. The Properties: Add Instance page appears. Provide the required values, then click **OK**. The Properties page of the Configure Cluster Database wizard reappears.
 7. Click **Next**. For versions 10.1 and higher, Enterprise Manager bypasses the Install Packages, Credentials, and Parameters steps, and goes directly to the Review page.
 8. Click **OK**. The Targets Discovered on Cluster page reappears, and displays the newly added cluster database and instances.

See Also: The Enterprise Manager online help for more information about configuring cluster databases

14.5.3 Discovering Instances Added to the Cluster Database

If you need to configure additional instances, follow these steps:

1. In Enterprise Manager, click **Databases** in the Targets page, and navigate to the desired **Cluster Database Home** page.
2. Click **Monitoring Configuration** in the Related Links section. The Properties page of the Configure Cluster Database wizard appears.
3. Provide the required information in the Properties table at the top of the page.
4. Examine the Instances table. One or more additional instances may exist, but may not appear in the Instances table. If this is the case, click **Add** to discover the instance in the cluster database. The Properties: Add Instance page appears.
5. Provide the required information, then click **OK**. The wizard Properties page reappears, and displays the added instance view.
6. Click **Check Connection** to ensure that the connection is working.

See Also: The Enterprise Manager online help for more information about discovering instances added to the cluster database

14.5.3.1 Troubleshooting

If you encounter configuration issues, check the following required conditions to ensure that automatic discovery is able to function correctly:

- The host user running the Management Agent is able to run the SRVCTL utility in the Oracle home and retrieve the database configuration.
- The host user running the Management Agent is able to connect to the database through SQLPLUS using OS authentication.
- The Oratab (UNIX) or Registry (Windows) contains information about the database.

If automatic discovery still does not resolve your configuration issues after you have ensured the conditions previously listed, you can manually configure cluster databases (see [Section 14.5.2, "Configuring Cluster Databases"](#)).

For more information about configurations for Oracle Enterprise Manager Grid Control, see [Chapter 3, "Grid Control Common Configurations"](#).

14.6 Collecting Client Configurations

A client is comprised of a host and operating system user. Client configuration data that is collected includes:

- Hardware for the client.
- Operating system (includes information such as operating system properties, file systems, and patches) for the client.
- Operating system-registered software.
- Network data, which includes:
 - Latency to the Web server
 - Bandwidth to the Web server
- Client-specific data items that describe the configuration of the browser used to access the client configuration collection applet, which includes:
 - Browser type (vendor)
 - Browser version
 - JVM vendor (of the JVM used to run the client configuration collection applet)
 - JVM version (of the JVM used to run the client configuration collection applet)
 - Proxy server (if specified)
 - Proxy server exceptions
 - Browser cache size (MB)
 - Browser cache update frequency
 - Supported HTTP version
- Other client-oriented data items, including:
 - Client configuration collection applet identifier (version, defined in the applet code)
 - Application URL (from which the client configuration collection applet was accessed)

- Boot drive serial number (not available from diskless systems)
- Collection timestamp (from the client configuration collection applet JSP)
- Collection durations, in milliseconds
- Client IP address
- Effective client IP address - if a network proxy server is being used between the client and the Web server providing the client configuration collection applet, the effective client IP address will be the IP address of the proxy server.

14.6.1 Configuring the Client System Analyzer

The Client System Analyzer (CSA) allows Web server administrators to collect and analyze end-user client data. The client data is collected by an applet, diagnosed and sent back to the CSA application. The Oracle Management Agent uploads this data to the Enterprise Manager Management Repository. After the client configuration data has been collected by the client configuration collection applet and written to the Web server directory specified by the CSA applet, the client configuration data is uploaded to the Oracle Management Repository.

You can either use the Client System Analyzer in the Grid Control application pre-installed with Enterprise Manager or you can deploy CSA independently to your Web server.

14.6.1.1 Client System Analyzer in Oracle Grid Control

Client System Analyzer in Grid Control - An instance of CSA is pre-installed with Enterprise Manager. If you use this option, you can collect client data without setting up a separate Web server. To activate the pre-installed CSA application in Enterprise Manager, click **Deployments**. Then click **Client System Analyzer in Grid Control** and use the button provided to activate the application. Once CSA is activated, end-users can use the URL provided to run the CSA applet. The CSA applet can collect base client configuration information from client systems and Oracle Collaboration Suite client information from Oracle Collaboration Suite client systems.

- To download the CSA applet and have it collect base client configuration information, a client should use the Client System Analyzer URL in this format:
`http[s]://management-service-host:port/em/public/ecm/csa/CSA`
- To download the CSA applet and have it collect Oracle Collaboration Suite client configuration information, a client should use the Client System Analyzer URL in this format:
`http[s]://management-service-host:port/em/public/ecm/csa/CSA?application=OCS`

14.6.1.2 Deploying Client System Analyzer Independently

The Client System Analyzer Application can be deployed independently to any J2EE-capable Web server. Click the **Deployments** tab. Then click **Getting Started with Client System Analyzer** and click **Deploy Client System Analyzer Application**. Follow these steps to deploy the CSA applet and collect the client configuration data.

1. Download the CSA Application:

The CSA application includes the CSA directory along with the necessary JSP applet files. The application is packaged as an EAR file. To download this default EAR file, click **Download Client System Analyzer Application**. You can customize the default CSA EAR file by modifying the following:

- Rules - This file contains a default set of rules against which the client data is evaluated. You can customize and add rules before deploying CSA.
- Context parameters - You can customize the context parameters in the web.xml file.
- Custom classes - You can provide customized applet classes that can be used to perform tasks like collecting additional data, changing the behavior of the applet, and performing certain operations on the client.

2. Deploy CSA to any J2EE Web server.

The CSA application is deployed on an Application Server as a regular J2EE application. Once the CSA application is deployed, context parameters can be changed similar to other web applications.

3. Direct users to the CSA.

In order for the client data to be collected, the user must access the CSA application. Users can access the CSA JSP page directly or by using a link from another application. Users can be automatically redirected to CSA using the following methods:

- HTTP Server (Apache's mod_rewrite) - This option does not require changes in the Web application.
- Servlet Filter - A servlet filter is a program that filters requests to and from the server. The CSA_filter.jar file contains the servlet filter classes. The servlet filter and the filter mapping need to be added to the Web application.
- CSA Redirection JSP - The CSA Redirection JSP (CSARedirect.jsp) page can be included into the Web application.

4. Configure Enterprise Manager.

Collected client data is recorded in the Receive File Directory on the Web server. To upload the collected client data into Enterprise Manager, you need to do the following:

- Add a CSA Collector Target to the Enterprise Manager Management Agent. To do so, click **Add Collector** and choose a target from the list.
- Specify the absolute path to the Receive File Directory. The path specified must be the same as the path specified in the outputDir parameter of the CSA application. By default, the client data is stored in the Receive File Directory "csa_results" under the context root of the Client System Analyzer Web application, but this can be configured by changing the applications's "outputDir" context parameter.

5. Test the CSA Deployment.

To verify the CSA deployment, click the URL of the CSA page and check if the client data is collected.

14.6.2 Configuration Parameters

The Client System Analyzer (CSA) can be further configured by modifying the context parameters in the CSA application's WAR file.

Table 14–1 Configuration Parameters

Parameter	Description	Default Value
alertWhenDone	If set to true, a message indicating that the applet has been executed is displayed.	false
appletJAR	The name of the JAR file.	CSA.jar
application	The name of the application associated with this CSA instance. If the application parameter value is not specified, then the Collection Tag has a value of Default.	none
autoRedir	If set to "true", this causes the CSA JSP page to automatically use the Sun JVM if JVM was set to JInitiator and the client does not have the appropriate version of JInitiator installed.	false
bwTestFile	The name of the file that is downloaded from the server during the bandwidth test.	CSA.mb (included with CSA)
bwTestMsec	The amount of time the applet should spend on the bandwidth test. The applet computes bandwidth by counting the number of bytes it can download in this interval.	200 ms
classid	The "classid" field for the OBJECT tag. Applicable only if JVM is set to "JInitiator." The classid for Sun is "clsid:8AD9C840-044E-11D1-B3E9-00805F499D93" codebase - the "codebase" field for the OBJECT tag. Applicable only if JVM is set to "JInitiator."	None – this field MUST be set if JVM is set to "JInitiator," and is ignored otherwise
codebase	The codebase field for the OBJECT tag. Applicable only if the JVM is set to "JInitiator".	The default for Sun is http://java.sun.com/products/plugin/autodl/jinstall-1_4_2-windows-i586.cab#Version=1,4,0,0
collectCookie	The list of the names of cookies to be collected. This parameter is a comma-separated list of cookie names. Only cookies for the current OS user in the current browser will be collected. The Administrator can specify asterisk (*) to collect all of the current user's cookies for the current browser.	If this field is not present, no cookies will be collected.
cookieDomain	The domain of the CSA cookie.	If either the domain or path of the cookie is not set, cookies are disabled
cookieMaxAge	The maximum duration, in seconds, of the cookie on the client machine.	1 year
cookiePath	The path of the CSA cookie	If either the domain or path is not specified, cookies are disabled.
customClass	The name of the class used to collect custom data.	none – the default behavior is for no custom code to be executed

Table 14–1 (Cont.) Configuration Parameters

Parameter	Description	Default Value
customKey1 customKey2 customKey3	The values of the three custom keys. All client collections done by a CSA JSP page that uses this deployment descriptor will have these values for the custom keys. These values can be overridden by custom code.	If no custom key values are specified, none will be collected (unless they are collected by custom code)
descriptionFile	The full path of a text file containing the description that will be displayed on the deployment page. The contents of the file should be HTML-formatted text.	None
destURL	Specifies the destination URL. This is the URL to which the "Proceed" button on the CSA JSP page is linked.	If no destURL is specified, the "Proceed" button will take the user to the referring page; if there is no referring page, the "Proceed" button will not be displayed.
destURLResultsParam	Specifies the name of the URL parameter that will be added to the "destination URL" to indicate the client's compliance level. For example, if the value was "compliance", and the client's overall compliance level was critical, then the parameter "compliance=critical" would be added to the destination URL.	Sun
JVM	This determines the type of JVM that is to be used. If the value is "Sun," the JSP page will direct the browser to use the Sun JVM. If the value is "Oracle," the page will direct the browser to use Oracle Jinitiator. If the value is "any," the JSP will write out the standard "applet" tag, which causes the client to use whichever JVM is plugged into the browser.	Sun
maxExecInterval	Parameter that is added to CSA cookie payload. When the redirection logic reads the cookie, if the timestamp of the cookie differs from the current time by more than this value, the applet is deployed again. This parameter can be overridden by the "csa execInterval" context parameter in the redirection JSP filter.	90 days
maxFileSize	Maximum amount of data, in KB, that can be posted back to the receiver in a single request. If the size of the posted data exceeds this limit, the request is rejected and any data already written to the hard drive is deleted.	100
maxOutputFiles	Maximum number of output files that can be present in XML OutputDir.	100
outputDir	Directory to which CSA configuration xml files will be written. Both the applet page and the receiver page must read this parameter, and this parameter must be identical for both pages.	By default, the output files are written into the "csa_results" subdirectory of the application root directory (if the application root directory exists, and if the subdirectory exists or can be created). Using the default value for this parameter is not recommended.

Table 14–1 (Cont.) Configuration Parameters

Parameter	Description	Default Value
outputEnabled	Enables or disables creation of output XML files. Applicable to both applet and receiver pages.	By default, the XML files are created and stored in the XMLOutputDir.
pluginspage	Used to direct the user to the JVM installer under Netscape, since Netscape does not support automatic installation. Applicable only if JVM is Jinitiator. Default for Sun is <code>http://java.sun.com/products/plugin/index.html#download</code>	none - This field must be set if JVM is set to "JInitiator" and is ignored otherwise.
receiver	The URL to which the applet should post the collected data. Note: When setting this parameter, the administrator must ensure that the version of the receiver is the same as the version of the applet.	Default is to look for "CSAr.jsp" in the same path as the CSA JSP page
ruleFile	Specifies the path on the server, relative to the web application root, of the file that contains the rules to be evaluated.	rules.xml
script	Specifies a script, provided by the administrator, which can be run on the CSA XML file before it is marked for upload by the agent.	none - If no script is specified, no script will be run.
type	The type field for the OBJECT tag rendered by the CSA JSP page to deploy the applet. This is only applicable if the JVM is set to JInitiator. If the JVM is set to Sun, the type is <code>application/x-java-applet</code> .	none - this field must be set if JVM is set to "JInitiator," and is ignored otherwise
viewData	If set to true, this parameters allows the end-user to view the collected data after it is posted to the server.	false

In addition to these parameters, the CSA redirection parameters can also be configured. Redirection can be enabled either by using a servlet filter or by including a CSA redirection JSP file in some other page. The following context parameters must be available for the redirection to work.

Table 14–2 Configuration Parameters

Parameter Name	Description	Default Value
csaURL	The URL of the CSA JSP page to which the user should be redirected.	No default: This value must be set or redirection cannot work.
execInterval	The interval, in seconds, between executions of CSA. If the difference between the cookie's age and the current server time is greater than execInterval, the user is re-directed.	None. If the execInterval is not set, then the user is only redirected if there is a CSA cookie.
redirectURL	The URL to which the user should be directed after CSA has executed	None. If this parameter is not set, the user is directed back to the originally requested page
UIMode	0 - synchronous (in the current browser window) 1 - asynchronous visible 2 - asynchronous invisible	synchronous

14.6.2.1 Associating the Parameters with an Application

In certain cases, different sets of parameters may be required for different applications. For example, two different applications may have different rule sets and custom code, and the administrator may want to associate them with different CSA Collector Targets. In this scenario, the administrator can specify the ruleFile, appletJar, script, and outputDir parameters for a particular application by using the context parameters <application name> ruleFile, <application name> appletJar, and so on. If an application is specified, either as a context parameter or through the URL, then CSA is executed using the parameter values specific to the application. If no application is specified, or if one of the parameters for an application is not overridden, the default parameters are used.

14.6.3 Rules

Custom rules can be supplied to the CSA application so that the users receive immediate feedback as to whether their systems satisfy certain constraints. A sample RULES file is shown in [Example 14–1](#) followed by a description of each tag contained in the file.

Example 14–1 Sample RULES

```
<RULES>
<RULE>
<NAME>Client has sufficient memory</NAME>
<DESCRIPTION>Checks to see if the client has enough memory to run the
application</DESCRIPTION>
<VIOLATION> //ROWSET[@TABLE='MGMT_ECM_HW']/ROW/AVAIL_MEMORY_SIZE_IN_MB[number()
&lt; $arg=SIZE$] </VIOLATION>
<SEVERITY level="CRITICAL">
<PARAM id='SIZE'>100</PARAM>
<MOREINFO>
<TEXT>Application cannot run with less than 100 MB. </TEXT>
</MOREINFO>
</SEVERITY>
<SEVERITY level="WARNING">
<PARAM id='SIZE'>150</PARAM>
<MOREINFO>
<TEXT>Approaching minimum memory level</TEXT>
</MOREINFO>
</SEVERITY>
</RULE>
</RULES>
```

[Example 14–1](#) demonstrates a rule that can be used to check whether or not the client has sufficient memory to run the application. The <VIOLATION> is an XPATH expression that the applet will evaluate against an XML file that contains all of the data it has collected. Since the violation is an XPATH expression embedded in an XML file, certain characters in the XPATH, such as '<', '>', and '&', must be replaced with entities. If the XPATH expression returns a non-null node set, the rule has failed. In this case, the rule will fail if the client's available memory is less than a certain amount. The actual amount that triggers a violation can be configured by using different severity levels.

In [Table 14–3](#), the applet will first replace the substring "\$arg=SIZE\$" in the VIOLATION expression with "100" and then evaluate the expression. If the client's available memory is less than 100 MB, then the rule will fail with critical status. The applet will indicate the status along with the message "Application cannot run with less than 100 MB of memory". If the rule passes through successfully, the applet will

then replace "\$arg=SIZE\$" with 150 and try again; if the rule fails, the applet will display the message "Approaching minimum memory level." If the applet goes through all specified severity levels and does not find a violation, the rule is successful.

Table 14–3 Tags in the RULES File

Tag Name	Description
RULES	This is the top-level tag for the XML file
BUNDLE	This tag specifies the resource bundles used for translation. The value of the tag is either the name of a file or a Java class name. The rule engine reads this string and first attempts to find a file in the applet JAR that has this name. This file is expected to contain a mapping of resource IDs to strings in various languages. If such a file does not exist, then the string is treated as the name of a Java resource bundle class. Strings in a resource bundle are referenced using the syntax <resource id>@<bundle id>.
PRECONDITION	This tag is used to specify an XPATH expression that must return a non-null node set in order for a rule to be evaluated. The "id" attribute specified the ID of the precondition. A rule can specify a list of preconditions that should be evaluated by listing their IDs.
RULE	This tag represents an individual node that is to be evaluated. The rule's severity is specified using a <SEVERITY> tag. At least one severity tag must be specified for a rule. The tag has an optional "precondition" attribute, which is used to specify a list of precondition IDs separated by commas. Before the rule is evaluated, all of the preconditions must be met. If the pre-conditions are not met, the rule has a status of "Not Applicable" and is not displayed in the client UI at all. The children of a RULE tag are NAME, DESCRIPTION, VIOLATION, SEVERITY, and MOREINFO.
NAME	This tag specifies the name of the rule and identifies the tag in the repository. Note: This tag must contain a value and cannot be blank.
DESCRIPTION	This is the description of the rule.
VIOLATION	This tag lists the violations that are to be checked for a given rule. The violation is specified in the CSA Condition Language.
SEVERITY	A rule can have three severity levels: INFO, WARNING, and CRITICAL. The SEVERITY node must contain a number of ARG children equal to the number of arguments that can be accepted by the expression in the VIOLATION node. When the rule engine evaluates a rule, it evaluates the condition in VIOLATION for each of the sets of arguments specified in the severity levels, starting with CRITICAL and moving down in order of severity. As soon as the engine encounters a condition that fails, the rule is declared a failure, with a severity level equal to the severity level of the argument that caused the failure. If the conditions for all specified levels are met, the rule passes.
PARAM	This tag specifies the value of an argument that should be substituted into an expression. The 'id' attribute of the tag must match the name of one of the arguments in the expression.
MOREINFO	This tag specifies the information that is displayed if the user clicks the "more information" button that is displayed next to a failed rule. The children of MOREINFO are TEXT and ARG. Note: The MOREINFO node can be a child either of the severity node (in the case where multiple severities are specified) or of the rule itself.

Table 14-3 (Cont.) Tags in the RULES File

Tag Name	Description
TEXT	This tag specifies the text to be displayed when the "More Info" button is clicked. The "resource" attribute specifies a string in a resource bundle – if this string is not present, the value of the node is displayed instead. The text (either in the resource bundle or in the node itself) can specify a location for arguments to be inserted by using "{0}", "{1}", and so on. In this case, the expressions in the ARG nodes are evaluated and inserted into the text in the order in which they are specified. If there are more ARG nodes specified than there are slots in the string, the extra nodes are ignored.
ARG	This tag specifies an expression in the CSA Condition Language that can be evaluated and inserted into the MOREINFO text.

See Also: Enterprise Manager online help associated with the Getting Started with CSA page

14.6.4 Customization

In addition to writing custom classes to collect custom properties, the administrator can also specify custom properties in the deployment descriptor. Custom property names are specified by including a context parameter of the form `csa value_<name>`. The `<name>` field of the context parameter name is treated by the Client System Analyzer (CSA) as the custom property name, and the value of the parameter is treated as the custom property value. Similarly, administrators can specify the `type`, `type_ui`, `name_ui`, `display_ui`, and `history_tracking` fields for a custom property by using `csa_type_<name>`, `csa_type_ui_<name>`, `csa_name_ui_<name>`, `csa_display_ui_<name>`, and `csa_history_tracking_<name>` parameters, respectively. Custom properties can also be specified on the CSA Applet URL, using the same naming convention.

14.6.5 CSA Deployment Examples

The following sections outline sample use cases for client configurations.

14.6.5.1 Using Multiple Collection Tags

An administrator can check the compatibility of users with two distinct Web applications. The first is an online teaching website that delivers content using a number of various plug-ins, allowing an administrator to be sure that all users have the required installed plug-ins. The second is a software distribution portal that allows an administrator to ensure that all users downloading software from the portal have the required hardware and operating system. In this case, though both applications require their own set of rules, the administrator can use a single CSA instance for both applications through the use of collection tags displayed in the following list:

1. Choose a collection tag for each application, such as "teaching" and "distribution".
2. Create two separate rule files, one for each application.
3. Use context parameters to map each rule file to the corresponding application, as shown in [Example 14-2](#).
4. Create the appropriate links from each application to CSA. The links from the teaching and distribution applications should have "application=teaching" and "application=distribution", respectively, in the query string. This ensures that users of each application have the correct collection tags when running CSA.

Example 14–2 Using Collection Tags for Selecting a Rule File

```

<context-param>
  <param-name>csa teaching ruleFile</param-name>
  <param-value>teaching_rules.xml</param-value>
</context-param>

<context-param>
  <param-name>csa distribution ruleFile</param-name>
  <param-value>distribution_rules.xml</param-value>
</context-param>

```

[Example 14–2](#) shows only the use of collection tags for selecting a rule file. However, collection tags can be used for any of the CSA context parameters.

Collection tags also affect how client configurations are stored in the Enterprise Manager Management Repository. If the user comes to CSA using the link from the teaching application in [Example 14–2](#), then in addition to running the rules for the "teaching" collection tag, CSA also causes this tag to be stored with the client configuration data in the Management Repository. The collection tag forms part of the unique identifier for the client configuration, which makes it possible for a single client to have multiple configurations in the Management Repository, each with its own tag. Collection tags can be associated with Enterprise Manager targets in order to restrict access to client data; an Enterprise Manager user can only view a client configuration if he or she has view privileges on a target that is associated with the collection tag for that client configuration.

In [Example 14–2](#), suppose that host H1, application server A1, and database D1 are used to host the teaching application, while host H2, application server A2, and database D2 are used for the distribution application. All 6 targets are monitored by Enterprise Manager, with user X having access to A1, H1, and D1 and user Y having access to A2, H2, and D2. Since each of the two Enterprise Manager users is monitoring the resources used for one of the applications, it may also make sense to have each user also monitor the application's clients. In that case, an Enterprise Manager super user would associate the "teaching" tag with A1, D1, or H1 and associate the "distribution" tag with A2, D2, or H2. This allows user X to see all client configurations with the "teaching" tag and user Y to see all configurations with the "distribution" tag.

14.6.5.2 Privilege Model for Viewing Client Configurations

Collection Tags are used to restrict access to client data in Enterprise Manager. A client configuration is visible to the user only if the Collection Tag for that configuration is associated with a target on which the user has View privileges. For example, if collection tag C is associated with target T1, then only those users that can view target T1 will be able to see client configurations that have tag X. In [Example 14–2](#), user X will be able to see client configurations with the "teaching" tag because user X has view privileges on targets that are associated with the "teaching" tag. However, user X will not be able to see any client configurations with the "distribution" tag because that tag is not associated with any targets that user X can see. Super users can associate collection tags with targets by using the Collection Tag Associations page, which can be accessed from the Deployments tab or from the Client System Analyzer in Grid Control link on the Setup page. Super users can view all client configurations regardless of any collection tag associations.

14.6.5.3 Using the Customization API Example

If the administrator is interested in the user's settings for an e-mail client in addition to the normal CSA data, the administrator can add this information to the other data collected by CSA through the use of the customization API, as shown in [Example 14-3](#).

1. Create the Java classes required to gather the information. The administrator can create as many classes as necessary, but there must be at least one class that implements `oracle.sysman.eml.ecm.csa.CSAResultInterface` and one that implements `oracle.sysman.eml.ecm.csa.CSACustomInteface`, both of which are shown in [Example 14-3](#). Assume that the former is `acme.csa.custom` and the latter is `acme.csa.result`.
2. Set the value of the "customClass" parameter in CSA to "acme.csa.custom"

Example 14-3 Customization API

```
public interface CSACustomInterface {

    /**
     * requires: none
     * effects: returns a CSAResultInterface object that may contain custom
     * properties. Other effects are determined by the customActions method
     * in the implementing class
     * modifies: unknown - dependent on implementing class.
     * @param inputData contains client config data collected by default, plus
     * applet parameters, etc.None of the data in the inputData is guaranteed
     * to be there as there could have been collection errors.
     * @return a data structure that may contain custom properties
     */
    CSAResultInterface customActions(CSAInputInterface inputData);
}

public interface CSAResultInterface {

    /**
     * requires: none
     * effects: returns an array of custom properties
     * modifies:none
     * @return String[][7] where ...
     *
     * String[i][0] is a name
     * String[i][1] is a value of the i-th row. (Type and name must be unique.)
     * String[i][2] is a type/category of data (could be null),
     * String[i][3] is the displayed value of the name of the property
     * String[i][4] is the displayed value of the type of the property
     * String[i][5] indicates data item (ie "Y") whose history should be computed
     * String[i][6] indicates data item (ie "Y") should be displayed in default UI
     */
    String[][] getResultsData();
}

public interface CSAInputInterface {

    /**
     * Get data value for given name
     * requires: name is not null
     * effects: returns the data value associated with the name
     * modifies: none
     * @param name the name of the key whose value is to be returned
     * @return the value associated with name
     */
}
```

```

*
*/
String getDataValue(String name);

/**
 * Get table-formatted data.
 * requires: name is not null
 * effects: returns the table with this name
 * modifies: none
 * @param name the name of the table
 * @return the rows of the child tables
 *
*/
CSAInputInterface[] getDataTable(String name);
}

```

The additional data collected by the custom code will be stored in the table `MGMT_ECM_CSA_CUSTOM`. To add data to this table, the custom code returns it in an object that implements `CSAResultInterface`. The custom code can also manipulate the normal data collected by CSA by modifying the `CSAInputInterface` object passed to the `customActions` method by the applet.

Since the custom code is executed before rules are evaluated, the administrator can also write rules based on the custom data. For example, if the administrator wants to write a rule that raises a critical error if the user does not have the correct IMAP server set up his or her e-mail client, the administrator would write custom code that retrieves the IMAP server settings and stores them in the `MGMT_ECM_CSA_CUSTOM` table and then writes a rule that checks these values.

14.6.5.4 Using the CSA Servlet Filter Example

Since CSA does not involve the use of a Management Agent on the user's machine, there is no way to keep the data in the Management Repository up to date unless end users run CSA periodically. One way to ensure that they do is to check whether or not users have run CSA recently, and if they have not, to inform them to run CSA again. This check can be accomplished using the CSA servlet filter provided by Oracle.

The CSA servlet filter works by checking the cookie that CSA sets in the user's browser whenever it runs. The payload of this cookie indicates the time at which CSA was last run. To use the filter, the administrator places it in front of some frequently accessed application, such as an employee portal. The administrator then sets the interval at which he or she wants users to run CSA. Whenever a user tries to connect to the portal application, the filter intercepts the request and checks the CSA cookie. If the cookie is not present or if it is older than the execution interval specified by the administrator, the user is directed to the CSA page; if not, the user is allowed to proceed to the application.

Assume that Acme Corporation has a CSA instance deployed at `www.acme.com/csa/CSA.jsp`. Assume also that the company has a portal at `www.acme.com/portal` that can be used by employees to check e-mail, access their personal information, or display news about the company. Because the portal is accessed frequently by employees, the administrator at Acme decides that the portal can be used to keep CSA data up to date. The administrator would take the following steps:

1. Download the CSA servlet filter classes. These classes are contained in a JAR file, `CSA_filter.jar`, which can be downloaded from the "Deploy Client System Analyzer" page in the Enterprise Manager Grid Control Console.

2. Place the JAR file in the WEB-INF/lib directory of the application to which the filter will be applied.
3. Specify context parameters for the filter. In this case, the administrator wants users to run CSA every 30 days and return to the portal homepage after CSA has finished.

```
<context-param>
  <param-name>csa csaURL</param-name>
  <param-value>www.acme.com/csa/CSA.jsp</param-value>
</context-param>

<context-param>
  <param-name>csa execInterval</param-name>
  <param-value>2592000</param-value>
</context-param>
```

An alternative is to have CSA run in a separate browser window in the background. This can be set up by using the `csa_uiMode` parameter. If the parameter is set to 1, the filter will open a new browser window that is the same size as the original window and go to the CSA page. If the parameter is set to 2, CSA will run in "invisible" mode; in this case, the filter will open a new browser window and immediately minimize it, and it will close the window as soon as CSA has completed.

14.6.5.5 Sample Deployments

In the following sample deployment examples, there are three primary actors. The first is the CSA administrator, who is responsible for setting up CSA. The second is the Enterprise Manager user, who will be viewing the client data in Enterprise Manager. The third is the end user, whose data is being collected by CSA.

14.6.5.5.1 Example 1: Helpdesk

In this example, the CSA administrator is using CSA to support the operations of a helpdesk. End users who have problems running a particular application can call customer support, and the customer support technician can, if necessary, instruct the user to go to a particular URL and run CSA. The Enterprise Manager users are the support personnel who will use the data collected by CSA to assist the end user. To speed up the process of diagnosing the customer's problem, the CSA administrator creates some rules in a file called "rules.xml" so that the helpdesk personnel can quickly identify potential problems. In the simplest case, suppose that the helpdesk is being set up to provide support for a single application. The application is running on an application server on host `application.acme.com`, which has an Enterprise Manager Management Agent installed on it that sends data back to the Management Service at `oms.acme.com/em`. The helpdesk personnel who will be looking at client data can log into Enterprise Manager as the user "helpdesk," which does not have super user privileges.

1. The CSA administrator adds `rules.xml` to the `CSA.war` file contained in `CSA.ear`.
2. Deploy the EAR file to the application server using the Application Services Control Console.
3. Use the Application Services Control Console to set the necessary context parameters, such as `ruleFile` and `outputDir`.
4. Optionally, the administrator can choose a collection tag for the CSA data by specifying a value for the "application" context parameter. If no tag is chosen, the tag "Default" will be used.

5. An Enterprise Manager user with super user privileges adds a CSA Collector Target to the Management Agent on application.acme.com and sets its receive file directory to the directory specified in the "outputDir" parameter of CSA.
6. An Enterprise Manager superuser creates the collection tag associations needed to allow the helpdesk users to look at the data. For example, the superuser could associate the tag "Default" with host application.acme.com and then give the "helpdesk" Enterprise Manager user view privileges on the host.

With the setup previously described, when a user calls the helpdesk to ask for support with the application, the helpdesk technician can instruct the user to run CSA from the appropriate URL on application.acme.com. The Management Agent collects the data after a certain interval and loads it into the Management Repository. The helpdesk technician can then log into Enterprise Manager as "helpdesk" and find the customer's information by searching for an identifying field such as the customer's operating system user name or host name. By default, the Management Agent will check the output directory for new data every two minutes, but this interval can be shortened by editing the file

```
$ORACLE_HOME/sysman/admin/default_collection/oracle_csa_collector.xml.
```

14.6.5.2 Example 2: Inventory

In [Example 14-4](#), a system administrator is in charge of keeping track of the hardware and software used by employees in two different departments, Human Resources (HR) and Sales. This administrator serves as both the Enterprise Manager user and the CSA administrator. The setup for this case is similar to the one described in the example on using servlet filters, but in this case, each department has its own portal application, at hr.acme.com/portal and sales.acme.com/portal, respectively. The administrator sets up an application server on host server1.acme.com and deploys CSA with the URL `http://server1.acme.com/csa/CSA.jsp`. A Management Agent on server1.acme.com collects data and sends to a Management Server at oms.acme.com/em. The administrator would like to collect data once every 30 days and to have CSA run in invisible mode. The administrator would also like to distinguish data from the two different departments by using two separate collection tags, "hr" and "sales." The administrator can log into Enterprise Manager as sysman and will thus be able to see clients with both tags.

The administrator arranges to have users directed to CSA by deploying the CSA servlet filter on both applications. Most of the filter context parameters for the two applications will be identical. However, because each application corresponds to a different tag, the values of the "csa csaURL" parameter will be slightly different. For the HR portal, the value would be `http://server1.acme.com/csa/CSA.jsp?application=hr`, and for the sales portal, the value would be `http://server1.acme.com/csa/CSA.jsp?application=sales`.

Example 14-4 Inventory Code

```
<context-param>
  <param-name>csa csaURL</param-name>
  <param-value>www.acme.com/csa/CSA.jsp?application=sales</param-value>
</context-param>

<context-param>
  <param-name>csa execInterval</param-name>
  <param-value>2592000</param-value>
</context-param>
```

```

<context-param>
  <param-name>csa uiMode</param-name>
  <param-value>2</param-value>
</context-param>

```

Under this setup, users in the HR department who are directed to CSA from the HR portal will have the tag "hr," and users from the sales department will have the tag "sales". Thus, if the administrator wants to see information about only hardware on machines in the HR department, he or she can simply use the "Collection Tag" filter on the Client Configurations page in Enterprise Manager and set it to "hr".

14.6.5.5.3 Example 3: Problem Detection

In this example, the goal is to use CSA to inform end users of potential problems they may experience while running an application. The setup is similar to the one used in Example 2. In this example, however, the CSA administrator creates rules for each application. In addition, the administrator wants CSA to run in the original browser window to ensure that end users are aware of any potential problems.

[Example 14-5](#) displays the context parameter values for the CSA servlet filter on the sales portal.

Example 14-5 Context Parameter Values for CSA Servlet Filter

```

<context-param>
  <param-name>csa csaURL</param-name>
  <param-value>www.acme.com/csa/CSA.jsp?application=sales</param-value>
</context-param>

<context-param>
  <param-name>csa execInterval</param-name>
  <param-value>2592000</param-value>
</context-param>

<context-param>
  <param-name>csa uiMode</param-name>
  <param-value>0</param-value>
</context-param>

```

[Example 14-6](#) represents the context parameter definitions to map rules to collection tags.

Example 14-6 Context Parameter Definitions Mapping Rules to Collection Tags

```

<context-param>
  <param-name>csa sales ruleFile</param-name>
  <param-value>sales_rules.xml</param-value>
</context-param>

<context-param>
  <param-name>csa distribution ruleFile</param-name>
  <param-value>hr_rules.xml</param-value>
</context-param>

```

Index

A

accessibility

- enabling accessibility mode, 1-18
- enabling accessibility features, 1-18
- providing textual descriptions of charts, 1-19

Additional Management Agent installation type, 1-4

advanced configuration

- introduction, 1-1
- types of tasks, 1-1

Agent Registration Password, 4-6

- changing, 4-13

Agent Upload Problems

- default notification rule, 12-7

AGENT_HOME

- definition, 1-4, 1-5

AGENT_HOME/bin, 1-5

AGENT_HOME/network/admin, 4-18

AGENT_HOME/sysman, 1-5

AGENT_

- HOME/sysman/admin/scripts/db/config/resp
onse.pl, 14-5

AGENT_HOME/sysman/config, 1-5

AGENT_HOME/sysman/log, 1-5

agentInstallJob.tcl, 11-2

Agents Unreachable

- default notification rule, 12-7

aggregation and purging policies

- See* data retention policies

alerts, 9-7

Application Performance Management, 4-35, 5-10

Application Server Availability and Critical/Warning States

- default notification rule, 12-7

Application Server Control

- directory structure, 1-6
- introduction, 1-6
- Ports page, 5-10
- starting and stopping, 2-7
- starting and stopping on Windows systems, 2-6

Application Service Level Management

- using to monitor the Management Service, 3-7

archive logging

- for Management Repository database, 8-1

assistive technology, 1-18

asynchronous I/O, 9-14

B

Backup, 9-16

Bad SQL

- configuring the database to show Bad SQL, 14-4

baselines, 9-5

Beacons, 9-14

- configuring firewalls to allow ICMP traffic, 5-10
- monitoring Web Applications over HTTPS, 4-35

blackouts

- controlling with emctl, 2-15
- examples, 2-17

buffer cache, 9-12

C

capacity

- predicting, 9-2

Certificate dialog box

- Internet Explorer, 4-33, 4-37

charts

- providing textual descriptions for

 - accessibility, 1-19

collection directory, 14-1

Common Configurations

- overview, 3-1

common configurations

- deploying a remote management repository, 3-4
- deploying Grid Control on a single host, 3-2
- firewalls and other security considerations, 3-1
- high availability configurations, 3-9
- managing multiple hosts, 3-4
- using multiple Management Services, 3-6
- when deploying Grid Control, 3-1

configuring for Management Services, 3-12

Configuring Services, 6-1

- Availability, 6-4, 6-6
- Beacons, 6-5
- Key Beacons, 6-5
- Local Beacon, 6-5
- Service Test-Based, 6-5, 6-6
- System-Based, 6-5, 6-6

Command Line Interface, 6-33

Create, 6-4

End-User Performance Monitoring, 6-2, 6-13

- Access Log Format, 6-17

- chronos_setup.sh, 6-19, 6-21
- EUM, 6-15
- Manage Web Server Data Collection, 6-16, 6-18
- Oracle Forms, 6-26
- Set URLs, 6-16
- Standalone Web Cache, 6-22
- Unprocessed Samples, 6-24
- URL Pattern, 6-18
- Web Cache Log Format, 6-19
- Web Cache Manager, 6-18
- Interactive Transaction Tracing, 6-3
 - J2EE Server Activity, 6-3
- Metrics
 - Performance, 6-5
 - Usage, 6-8
 - Usage Metrics, 6-5
- Monitoring Settings, 6-12
 - Beacon Overrides, 6-12
 - Collection Settings, 6-13
 - Data Granularity, 6-12
 - Frequency, 6-12
- Monitoring Templates, 6-30
 - Beacons, 6-31
 - Service Tests, 6-31
 - Service Tests and Beacons, 6-30
 - Variables, 6-31
- Oracle Application Server 10g (9.0.4), 6-3
- Performance, 6-5
- Performance Metrics, 6-7
 - Aggregation Function, 6-7
- Recording Transactions, 6-3, 6-12
- Request Performance, 6-3, 6-27
 - Correlate Requests, 6-3
 - Database Connections, 6-27
 - Enterprise Java Beans (EJBs), 6-27
 - JDeveloper, 6-29
 - Manage OC4J Data Collection, 6-28
 - OC4J Cluster, 6-27
 - OC4J Instances, 6-27
 - OC4J Tracing, 6-28
 - Oracle Application Server 10g (9.0.4), 6-3
 - Oracle User Interface XML (UIX), 6-29
 - Service Tests and Beacons, 6-28
 - Tracing Properties, 6-28
- Root Cause Analysis, 6-3, 6-7, 6-10
 - Component Tests, 6-11
 - Topology, 6-10
 - Topology Viewer, 6-3
- Service Level Rules, 6-31
 - Actual Service Level, 6-32
 - Availability, 6-31
 - Business Hours, 6-31
 - Expected Service Level, 6-32
 - Information Publisher, 6-32
 - Performance Criteria, 6-32
 - Services Dashboard, 6-32
- Service Tests and Beacons, 6-9
 - Configuring Dedicated Beacons, 6-10
 - SSL Certificate, 6-10

- Tests, 6-9
- Web Proxy, 6-10
- Service-Test Based Availability
 - Key Service Tests, 6-6
- System, 6-3
 - Key Components, 6-4
- System-Based Availability
 - Key Components, 6-6
- Test Performance, 6-2
- Thresholds
 - Critical, 6-5
 - Warning, 6-5
- Time Zone, 6-4
- Types
 - Aggregate Service, 6-13
- Types of Service
 - Generic Service, 6-4
- Types of Services
 - Aggregate Service, 6-4
 - OCS Service, 6-4
 - Web Application, 6-4
- connect descriptor
 - using to identify the Management Repository database, 8-7, 8-8
- creating a monitoring script, 13-2

D

- data collections
 - how Enterprise Manager stores, 14-1
 - restoring default, 14-1
 - understanding default and custom, 14-1
- Data Guard
 - configuring Enterprise Manager availability, 8-1
- data retention policies
 - for Application Performance Management data, 8-3
 - for other Management data, 8-3
 - modifying default, 8-3
 - of the Management Repository, 8-2
 - when targets are deleted, 8-4
- Database Availability and Critical/Warning States
 - default notification rule, 12-7
- Database Configuration Assistant
 - See* DBCA
- Database Control
 - configuring after installation, 1-8, 1-10
 - configuring during installation, 1-8
 - configuring with DBCA, 1-9
 - configuring with EMCA, 1-10
 - directory structure, 1-6
 - introduction, 1-6
 - location of Management Agent and Management Service support files, 1-6
 - starting on UNIX, 2-8
 - stopping on UNIX, 2-8
- DBCA
 - configuring Database Control with, 1-9
 - Management Options page, 1-9
 - starting on UNIX, 1-9

- starting on Windows, 1-9
- DBSNMP database user, 2-13
 - setting the password for, 2-13
- DBSNMP user, 14-5
- default_collection directory, 14-1
- deleting targets
 - data retention policies when, 8-4
- directory structure
 - introduction to, 1-1
- Disaster Recovery, 9-17
- disk mirroring and stripping
 - Management Repository guideline, 8-1
- disk space management
 - controlling Management Agent disk space, 10-3
 - controlling the contents of trace files, 7-4
 - controlling the size and number of log and trace files, 7-3, 7-5, 7-6
 - controlling the size of log and trace files, 7-8
- dontProxyFor
 - description of property, 5-7
 - property in emoms.properties, 5-6
- dropping the Management Repository, 8-6

E

- E2E monitoring, 9-14
- EM Website
 - using to monitor the Management Service, 3-7
 - Web Application target, 3-7
- em_message, 13-4
- em_result, 13-3
- emagent.log, 7-1
- emagentlogging.properties, 7-5
 - log4j.rootCategory property, 7-6
 - MaxBackupIndex property, 7-5
 - MaxFileSize property, 7-5
- emagent.nohup, 7-2
- emagent.trc, 7-2
- EMCA
 - command-line arguments, 1-11
 - configuring Database Control for Real Application Clusters, 1-15
 - configuring Database Control with, 1-10
 - reconfiguring Database Control after changing the listener port, 1-17
 - sample EMCA input file, 1-15
 - specifying port assignments, 1-16
 - troubleshooting problems with the Database Control, 1-17
 - troubleshooting tips, 1-17
 - using an input file for EMCA parameters, 1-14
- emctl, 2-1
 - controlling blackouts, 2-15
 - listing targets on a managed host, 2-15
 - location in AGENT_HOME, 1-5
 - security commands, 4-6
 - setting monitoring credentials, 2-14
 - starting, stopping, and checking the Management Service, 2-4
- emctl config agent credentials, 2-14

- emctl config agent listtargets, 2-15
- emctl config oms
 - sample output, 4-20, 4-27
- emctl config oms sso, 4-19
- emctl getemhome, 1-8
- emctl istop, 2-3
- emctl reload, 2-13
- emctl secure agent, 4-9
 - sample output, 4-10
- emctl secure lock, 4-12
- emctl secure oms, 4-6, 4-15
 - sample output, 4-7
- emctl secure setpwd, 4-14
- emctl secure unlock, 4-12
- emctl start agent, 2-2
- emctl start blackout, 2-16
- emctl start dbconsole, 2-8
- emctl start iasconsole, 2-7
- emctl start oms, 2-5
- emctl status agent, 2-2
- emctl status blackout, 2-16
- emctl status oms, 2-5
- emctl stop agent, 2-2
- emctl stop blackout, 2-16
- emctl stop dbconsole, 2-8
- emctl stop iasconsole, 2-7
- emctl stop oms, 2-5
- emctl upload, 2-13
- EMD_URL
 - property in the emd.properties file, 10-2
- emd.properties, 7-3, 10-2, 10-3
 - EMD_URL, 10-2
 - emdWalletDest, 10-2
 - emdWalletSrcUrl, 10-2
 - location, 1-5
 - LogFileMaxRolls, 7-3
 - REPOSITORY_PROXYHOST, 5-4
 - REPOSITORY_PROXYPORT, 5-4
 - REPOSITORY_URL, 3-3, 3-5, 10-2
 - TrcFileMaxrolls, 7-3
 - TrcFileMaxSize, 7-3
 - UploadMaxBytesXML, 10-3
 - UploadMaxDiskUsedPct, 10-3
- emdRepConnectDescriptor
 - property in emoms.properties, 3-18
- emdRepPort
 - property in the emoms.properties file, 10-9
- emdRepPwd
 - property in the emoms.properties file, 10-9
- emdRepServer
 - property in the emoms.properties file, 10-9
- emdRepSID
 - property in the emoms.properties file, 10-9
- emdRepUser
 - property in the emoms.properties file, 10-9
- emdWalletDest
 - property in emd.properties, 10-2
- emdWalletSrcUrl
 - property in emd.properties, 10-2
- em.notification.emails_per_minute

- property in emoms.properties, 12-4
- em.notification.os_cmd_timeout
 - property in emoms.properties, 12-12
- emoms.log, 7-6, 7-7
- emomslogging.properties, 7-7, 7-8
 - MaxBackupIndex, 7-7
 - MaxFileSize, 7-7
- emoms.properties, 8-6, 10-8
 - configuring the JDBC connection to the Management Repository, 3-4, 3-5
 - dontProxyFor property, 5-6
 - emdRepConnectDescriptor, 3-18
 - emdRepPort, 10-9
 - emdRepPwd, 10-9
 - emdRepServer, 10-9
 - emdRepSID, 10-9
 - emdRepUser, 10-9
- em.notification.emails_per_connection, 12-3
 - property in emoms.properties, 12-3
- em.notification.emails_per_minute, 12-4
- em.notification.os_cmd_timeout, 12-12
- maxInactiveTime, 14-7
- oracle.net.crypto_checksum_client, 4-17
- oracle.net.crypto_checksum_types_client, 4-17
- oracle.net.encryption_client, 4-17
- oracle.net.encryption_types_client, 4-17
- oracle.sysman.eml.mntr.emdRepPwd, 8-6
- oracle.sysman.eml.mntr.emdRepPwdEncrypted, 8-6
- oracle.sysman.emRep.dbConn.enableEncryption, 4-16
- oracle.sysman.emSDK.sec.DirectoryAuthenticationType, 4-23
- oracle.sysman.emSDK.svlt.ConsoleServerPort, 10-10
- proxyHost property, 5-6
- proxyPort property, 5-6
- sample Management Repository properties, 10-9
- emoms.trc, 7-6
- emwd watchdog script
 - in the AGENT_HOME/bin directory, 10-4
- End-User Performance Monitoring
 - Oracle Forms
 - Forms Web Configuration, 6-26
 - Web Server
 - Apache HTTP Server 2.0, 6-2, 6-14, 6-15
 - Oracle Application Server Web Cache, 6-2, 6-16
 - Oracle HTTP Server, 6-14, 6-15
 - Oracle HTTP Server Based on Apache 2.0, 6-2, 6-14
 - OracleAS Web Cache, 6-14
- Enterprise Manager
 - See Oracle Enterprise Manager
- Enterprise Manager 10g Grid Control Using a New Database
 - installation type, 3-2
- Enterprise Manager Configuration Assistant
 - See EMCA
- Enterprise Manager Framework Security

- about, 4-4
- compared with Oracle HTTP Server security features, 4-4
- configuring, 4-4
- enabling for Management Repository, 4-15
- enabling for multiple Management Services, 4-11
- enabling for the Management Agent, 4-9
- in a firewall environment, 5-2
- overview of steps required, 4-5
- restricting HTTP access, 4-11
- types of secure connections, 4-5
- Enterprise User Security
 - configuring Enterprise Manager for, 4-23

F

- fetchlet
 - log and trace files, 7-4
- firewalls
 - between browser and the Grid Control, 5-2
 - between Grid Control and a managed database target, 5-8
 - between Management Service and Management Agents, 5-9
 - between Management Service and Management Repository, 5-8
 - configuring for ICMP traffic, 5-10
 - configuring for UDP traffic, 5-10
 - configuring the Management Agent for, 5-3
 - configuring the Management Service for, 5-5
 - configuring to allow incoming data from Management Service, 5-7
 - configuring to allow incoming traffic to Management Agent, 5-4
 - considerations before configuring, 5-1
 - considerations when using with multiple Management Services, 5-9

G

- getemhome
 - emctl command, 1-8
- Grid Control
 - common configurations, 3-1
 - components, 9-1
 - configuring notifications, 12-1
 - deploying on a single host, 3-2
 - sizing, 9-2, 9-3
 - starting, 2-10
 - starting all components of, 2-10
 - stopping, 2-11
 - stopping all components of, 2-11
 - summary of the architecture and components, 3-2
- guidelines
 - for deploying the Management Repository, 8-1

H

- High Availability
 - installation and configuration, 9-18

- Host Availability and Critical/Warning States
 - default notification rule, 12-8
- hostname_lookup.txt, 11-5
- HTTP 500 - Internal server error, 2-5
- HTTP access
 - restricting, 4-11
- HTTP Server Availability and Critical/Warning States
 - default notification rule, 12-8
- http_em.conf, 10-10
- httpd.conf
 - configuring for use with a server load balancer, 3-15
 - Oracle HTTP Server configuration file, 3-15
- HTTPS, 4-5
- Hyper-Threading, 9-11

I

- ICMP, 5-10
- initialization parameter
 - adjusting when using multiple Management Services, 3-6
- Internet Control Message Protocol, 5-10
- Internet Explorer
 - Certificate dialog box, 4-33, 4-37
 - security alert dialog box, 4-32
 - Security Information dialog box, 4-35
- introduction to advanced configuration, 1-1
- I/O Channels
 - monitoring, 9-13
- istop
 - emctl command, 2-3

J

- J2EE, 1-2, 4-4
 - directory in Oracle Management Service home, 1-4
- Java Message Service (JMS), 1-17
- javax.net.ssl.SSLException
 - SSL handshake failed, 4-36
- job_queue_processes, 14-6

L

- Listener Availability
 - default notification rule, 12-8
- Listener port
 - obtaining, 5-9
- load balancing
 - connections between the Grid Control Console and Management Service, 3-13
 - connections between the Management Agent and Management Service, 3-10
 - using Oracle Net load balancing and failover, 3-18
- Loader, 9-11
- Loader backlog (files)
 - on the Grid Control Management System tab, 3-7
- loader threads, 9-11
- log files

- controlling the content of, 7-4
- controlling the size and number of, 7-6
- controlling the size of, 7-3
- fetchlet log files, 7-4
- locating and configuring, 7-1
- locating Management Agent, 7-2
- locating Management Service, 7-6
- Management Agent, 7-1
- Oracle Management Service, 7-6
- rollover files, 7-2
- log4j.appender.emagentlogAppender.MaxBackupIndex, 7-5
- log4j.appender.emagentlogAppender.MaxFileSize, 7-5
- log4j.appender.emagenttrcAppender.MaxBackupIndex, 7-5
- log4j.appender.emagenttrcAppender.MaxFileSize, 7-5
- log4j.appender.emlogAppender.
 - MaxBackupIndex, 7-7
- log4j.appender.emlogAppender.MaxFileSize, 7-7
- log4j.appender.emtrcAppender.
 - MaxBackupIndex, 7-7
- log4j.appender.emtrcAppender.MaxFileSize, 7-7
- log4j.rootCategory property in
 - emagentlogging.properties, 7-6
- log4j.rootCategory=WARN, emlogAppender, emtrcAppender, 7-8
- LogFileMaxRolls property in emd.properties, 7-3
- Login Timeout Value
 - modifying the default, 14-6
- LVM (Logical Volume Manager), 8-1

M

- Management Agent, 9-1, 9-4, 9-5
 - additional Management Agent commands, 2-12
 - checking the status on UNIX, 2-2
 - checking the status on Windows, 2-3
 - configuring to allow incoming communication from the Management Service, 5-4
 - configuring to use a proxy server, 5-4
 - configuring trust points, 10-7
 - reinstalling, 9-17
 - starting and stopping on UNIX, 2-1
 - starting and stopping on Windows, 2-2
- Management Information Base (MIB), 12-29
 - definition, 12-29
 - MIB variable descriptions, 12-30
- Management Options page
 - in DBCA, 1-9
 - in Oracle Universal Installer, 1-8
- Management Repository, 9-22
 - See Oracle Management Repository
- Management Server, 9-10
- Management Servers
 - adding, 9-13
- Management Service, 9-2, 9-4, 9-21
 - See Oracle Management Service
 - starting and stopping on Windows systems, 2-6

- using a server load balancer, 3-10
- master agent
 - Oracle Peer SNMP Master Agent service, 2-3
- MaxBackupIndex
 - property in emomslogging.properties, 7-7
- MaxBackupIndex property in
 - emagentlogging.properties, 7-5
- MaxFileSize
 - property in emomslogging.properties, 7-7
- MaxFileSize property in
 - emagentlogging.properties, 7-5
- maxInactiveTime
 - property in emoms.properties, 14-7
- metric thresholds
 - configuring after migrating from previous versions
 - of Enterprise Manager, 11-8
- MGMT_ADMIN.DISABLE_METRIC_DELETION, 8-5
- MGMT_ADMIN.ENABLE_METRIC_DELETION, 8-5
- MGMT_METRICS_1DAY table, 8-4
- MGMT_METRICS_1HOUR table, 8-4
- MGMT_METRICS_RAW table, 8-3
- MGMT_PARAMETERS table, 8-3
- MGMT_RT_datatype_1DAY table, 8-4
- MGMT_RT_datatype_1HOUR table, 8-4
- MGMT_RT_datatype_DIST_1DAY table, 8-4
- MGMT_RT_datatype_DIST_1HOUR table, 8-4
- MGMT_RT_METRICS_RAW table, 8-4
- MIB
 - See Management Information Base (MIB)
- migrating
 - configuring metric thresholds, 11-8
 - deploying and configuring Management Agents, 11-2
 - from previous versions of Enterprise Manager, 11-1
 - migrating management repository data, 11-7
 - overview of the Enterprise Manager migration process, 11-1
 - requirements, 11-1
 - supported versions, 11-1
 - using the repo_mig script, 11-7
- monitoring credentials
 - defined, 2-13
 - example of setting, 2-14
 - setting, 2-13
 - setting in Grid Control, 2-14
 - setting with emctl, 2-14
- monitoring script creation, 13-2
- monitoring templates, 13-16

N

- Netscape Navigator
 - New Site Certificate dialog box, 4-34
- network/admin, 4-15, 4-16, 4-18
- New Site Certificate dialog box
 - Netscape Navigator, 4-34
- Notification backlog

- on the Grid Control Management System tab, 3-7
- notification methods
 - based on a PL/SQL Procedure, 12-15
 - based on an SNMP trap, 12-19
 - based on operating system commands, 12-11
- notification rules
 - definition, 12-6
 - out-of-the-box notification rules, 12-5
- notification schedules, 12-5
- notifications
 - assigning methods to rules, 12-27
 - assigning rules to methods, 12-28
 - configuring, 12-1
 - defining multiple mail servers, 12-3
 - long e-mail notifications, 12-4
 - mail server settings, 12-2
 - mail server settings in emoms.properties, 12-3
 - management information base (MIB), 12-29
 - notification schedules, 12-5
 - sample Operating System command script, 12-14
 - setting up, 12-1
 - short email notifications, 12-5
 - using custom notification methods, 12-11
- Notification Rules
 - Custom, 12-9

O

- OC4J Availability and Critical/Warning States
 - default notification rule, 12-8
- OEM_MONITOR, 14-5
- Operating System command
 - sample notification method for, 12-12
 - sample script, 12-14
- Operating System scripts, 12-11
 - while creating notification methods, 12-11
- OPMN
 - See Oracle Process Management and Notification
- opmnctl
 - using to start Web Cache, 2-6
 - using to stop Web Cache, 2-6
- opmnctl startall, 2-4, 4-7, 4-20
- opmnctl status, 2-6
- opmnctl stopall, 2-4, 4-6, 4-19
- opmnctl stopproc ias-component=WebCache, 2-6
- ORA-12645
 - Parameter does not exist, 4-15
- Oracle Advanced Security, 4-5, 4-15, 5-8
 - enabling for Management Repository, 4-17
 - enabling for the Management Agent, 4-18
- Oracle Application Server
 - Enterprise Manager directories installed with, 1-6
 - J2EE and Web Cache installation type, 1-2
- Oracle Application Server Web Cache
 - as part of a common configuration, 3-1, 3-3
 - bypassing, 3-1
 - default port number, 2-5
 - errors when not running, 2-5
 - starting and stopping, 2-5
 - starting and stopping with opmnctl, 2-6

- using with Grid Control, 2-5
 - Web Cache Manager, 6-18
- Oracle Data Guard
 - using for the Management Repository, 3-16
- Oracle Database 10g
 - Enterprise Manager directories installed with, 1-6
- Oracle Enterprise Manager
 - components, 9-4
 - Configuring for high availability, 9-18
 - directory structure, 1-1
 - log files, 7-1
 - migrating from previous versions, 11-1
 - rollup process, 9-12
 - security model, 4-1
 - starting and stopping Enterprise Manager components, 2-1
- Oracle Enterprise Manager 10g Grid Control
 - See* Grid Control
- Oracle Enterprise Manager Release 2.2, 11-1
- Oracle Enterprise Manager Release 9.0.1, 11-1
- Oracle Enterprise Manager Release 9.2, 11-1
- Oracle Enterprise Manger
 - tuning, 9-10
- Oracle HTTP Server
 - configuring for use with a server load balancer, 3-15
- Oracle Identity Management, 4-3
- Oracle Internet Directory, 4-21
- Oracle Management Agent, 9-18
 - about the log and trace files, 7-1
 - changing the port, 10-2
 - configuring when protected by a firewall, 5-3
 - controlling disk space used by, 10-3
 - controlling the content of trace files, 7-4
 - controlling the size of log and trace files, 7-3
 - directory structure, 1-4
 - directory structure on Windows, 1-5
 - enabling security for, 4-9, 4-18
 - fetchlet log and trace files, 7-4
 - installing with Grid Control, 1-4
 - location of log and trace files, 7-2
 - log and trace files, 7-1
 - log and trace rollover files, 7-2
 - reconfiguring to use a new Management Service, 10-1
 - starting and stopping, 2-1
 - Watchdog process, 10-4
- Oracle Management Repository, 9-2
 - changing the Management Repository password, 10-9
 - configuration parameters in the emoms.properties file, 10-9
 - configuring for high availability, 3-15
 - data retention policies, 8-2
 - deploying on a remote host, 3-4
 - deployment guidelines, 8-1
 - dropping, 8-6
 - enabling Oracle Advanced Security, 4-17
 - enabling security for, 4-15
 - identifying with a connect descriptor, 8-7, 8-8
 - installing in a RAC database instance, 3-17
 - load balancing with Oracle Net, 3-18
 - migrating data from previous versions, 11-7
 - protecting with Oracle Data Guard, 3-16
 - protecting with Real Application Clusters, 3-16
 - recreating, 8-6, 8-7
 - reloading data, 2-13
 - restoring, 9-17
 - specifying the size of tablespaces in a RAC database, 3-18
 - starting the Management Repository database, 2-10
 - troubleshooting, 8-9
 - uploading data, 2-13
 - using raw devices with a RAC database, 3-18
- Oracle Management Service, 9-17
 - about the log and trace files, 7-6
 - adjusting the PROCESSES initialization parameter, 3-6
 - bin directory, 1-4
 - components installed with, 1-2
 - configuring for use with a proxy server, 5-6
 - configuring to allow incoming data from Management Agent, 5-7
 - configuring to use a new Repository, 10-8
 - configuring when protected by a firewall, 5-5
 - determining when to use multiple Management Services, 3-6
 - enabling security for, 4-6
 - enabling security for multiple Management Services, 4-11
 - home directory, 1-2
 - j2ee directory, 1-4
 - location the log and trace files, 7-6
 - log and trace files, 7-6
 - modifying monitoring credentials, 2-13
 - monitoring the load, 3-6
 - monitoring the response time, 3-7
 - monitoring with Application Service Level Management, 3-7
 - opmn directory, 1-4
 - reconfiguring, 10-8
 - reconfiguring to use a new port, 10-9
 - restoring, 9-17
 - starting, stopping, and checking, 2-4
 - sysman directory, 1-4
 - tips for monitoring the load and response time, 3-6
 - using multiple management services, 3-6
- Oracle Net firewall proxy access, 5-8
- Oracle Process Management and Notification (OPMN)
 - using to start and stop the Management Service, 2-4, 2-7
- Oracle Process Manager and Notification (OPMN), 1-4
- Oracle Real Application Clusters
 - specifying the size of the Management Repository tablespaces, 3-18
 - using for the Management Repository, 3-16

- Oracle Real Applications Clusters
 - installing the Management Repository into a RAC instance, 3-17
- Oracle Technology Network (OTN), 6-22
- Oracle Universal Installer
 - Management Options page, 1-8
- ORACLE_HOME/bin, 1-4
- ORACLE_HOME/hostname_sid/, 1-6
- ORACLE_HOME/install, 5-10
- ORACLE_HOME/j2ee, 1-4
- ORACLE_HOME/network/admin, 4-15, 4-16, 4-18
- ORACLE_HOME/oc4j/j2ee, 1-7
- ORACLE_HOME/oc4j/j2ee/OC4J_DBConsole, 1-7
- ORACLE_HOME/opmn, 1-4
- ORACLE_HOME/opmn/bin, 2-5
- ORACLE_HOME/sysman, 1-4, 1-6
- ORACLE_HOME/sysman/agent_download/, 11-2
- Oracle8i database
 - configuring for monitoring, 14-4
- Oracle9i
 - configuring for monitoring, 14-4
- oracle.net.crypto_checksum_client
 - property in emoms.properties, 4-17
- oracle.net.crypto_checksum_types_client
 - property in emoms.properties, 4-17
- oracle.net.encryption_client
 - property in emoms.properties, 4-17
- oracle.net.encryption_types_client
 - property in emoms.properties, 4-17
- oracle.sysman.eml.mntr.emdRepPwd
 - property in emoms.properties, 8-6
- oracle.sysman.eml.mntr.emdRepPwdEncrypted
 - property in emoms.properties, 8-6
- oracle.sysman.emRep.dbConn.enableEncryption
 - entry in emoms.properties, 4-16
- oracle.sysman.emSDK.sec.DirectoryAuthenticationType
 - property in emoms.properties, 4-23
- oracle.sysman.emSDK.svlt.ConsoleServerPort
 - property in emoms.properties file, 10-10
- OS scripts
 - See Operating System scripts
- OTN (Oracle Technology Network), 6-22
- OUIinventories.add, 14-3

P

- password
 - changing the Management Repository password, 10-9
 - changing the SYSMAN password, 8-5
- peer encapsulator service
 - SNMP, 2-3
- Performance Metrics
 - Beacon Aggregation Function
 - Average, 6-7, 6-8
 - Maximum, 6-7
 - Minimum, 6-7, 6-8
 - Sum, 6-7, 6-8
 - System Aggregation Function

- Maximum, 6-8
- PERFSTAT, 14-5
- PL/SQL procedures, 12-11
 - sample, 12-18
 - while creating a notification method, 12-15
 - while creating notification methods, 12-11
- portlist.ini, 5-10
- ports
 - 4888, 5-5, 5-8
 - 4889, 5-5, 5-8, 10-9
 - changing the Management Agent port, 10-2
 - default port for the Management Agent upload URL, 3-3
 - default Web Cache port on Windows systems, 3-3
 - displaying in the Application Server Control, 5-11
 - portlist.ini, 5-10
 - reconfiguring Database Control after changing the listener port, 1-17
 - reconfiguring the port used by the Management Service, 10-9
 - specifying Database Control ports, 1-16
 - viewing a summary of ports assigned during installation, 5-10
- Preferred Credential Override
 - when migrating from previous versions of Enterprise Manager, 11-6
- PROCESSES, 3-6
- ProcessManager
 - service used to control the Management Service on Windows systems, 2-7
- proxy server
 - configuring Management Agent for, 5-4
 - configuring the Management Service for, 5-6
- proxyHost
 - property in emoms.properties, 5-6
- proxyPort
 - property in emoms.properties, 5-6
- Public Key Infrastructure (PKI), 4-5, 4-36
- purging policies
 - See data retention policies

R

- RAID-capable disk
 - Management Repository guideline, 8-1
- raw devices
 - when used with the Management Repository, 3-18
- Recovery, 9-16
- Remote Method Invocation (RMI), 1-17
- RepManager script, 8-7, 8-8
- repo_mig
 - script for migrating from previous versions of Enterprise Manager, 11-7
- Repository Operations Availability
 - default notification rule, 12-8
- REPOSITORY_PROXYHOST
 - property in emd.properties, 5-4

REPOSITORY_PROXYPORT
 property in emd.properties, 5-4
 REPOSITORY_URL
 property in emd.properties, 3-3, 3-5
 property in the emd.properties file, 10-2
 requirements
 for migrating from previous versions of Enterprise
 Manager, 11-1
 rollover files, 7-2
 rollup process, 9-12
 Root Cause Analysis
 Mode
 Automatic, 6-11
 Manual, 6-11
 root password
See also SYSMAN
 when enabling security for the Management
 Service, 4-6
 root.sh
 when migrating from previous versions of
 Enterprise Manager, 11-6

S

scalability
 determining when to use multiple Management
 Services, 3-6
 screen readers, 1-18
 script registration, UDM, 13-5
 script results, returning, 13-2
 security
 about Enterprise Manager security, 4-1
 authorization and access enforcement, 4-2
 classes of users and their privileges, 4-2
 Enterprise Manager security model, 4-1
 leveraging Oracle Application Server security
 services, 4-3
 leveraging Oracle Identity Management
 Infrastructure, 4-3
 overview of steps required to enable Enterprise
 Manager Framework Security, 4-5
See also Enterprise Manager Framework Security
 security alert dialog box
 Internet Explorer, 4-32
 security certificate alerts
 responding to, 4-32
 security features
See Enterprise Manager Framework Security
 Security Information dialog box
 Internet Explorer, 4-35
 self-monitoring
 feature of the Management Agent, 10-4
 Server Connection Hung
 error while creating the repository, 8-9
 Server Load Balancer, 4-15
 server load balancer, 3-11
 configuring a virtual pool, 3-14
 configuring a virtual service, 3-14
 configuring for Management Agent data
 upload, 3-12

modifying the httpd.conf file, 3-15
 modifying the ssl.conf file, 3-15
 using with Management Services, 3-10
 using with the Grid Control Console, 3-13
 Server Load Balancers, 9-20
 ServerName directive
 in the httpd.conf file, 3-15
 in the ssl.conf file, 3-15
 Service Tests and Beacons
 Tests
 DNS, 6-9
 FTP, 6-9
 SOAP, 6-9
 Web Transaction, 6-9
 Services control panel
 using to start and stop the Management
 Agent, 2-7, 2-9
 using to start the Management Service, 2-6
 session timeout
 modifying, 14-6
 setupinfo.txt, 2-5
 Single Sign-On
 bypassing the Single Sign-On logon page, 4-22
 configuring Enterprise Manager with, 4-19
 registering Enterprise Manager as a partner
 application, 4-22
 registering Single Sign-On users as Enterprise
 Manager administrators, 4-21
 using Single Sign-On to authorize Enterprise
 Manager users, 4-19
 SNMP
 Oracle Peer SNMP Master Agent service, 2-3
 Oracle SNMP Peer Encapsulator service, 2-3
 SNMP traps, 12-11, 12-19
 sample, 12-20
 SQLNET.CRYPTO_SEED
 entry in sqlnet.ora, 4-18
 SQLNET.ENCRYPTION_SERVER
 entry in sqlnet.ora, 4-18
 sqlnet.ora, 4-15, 4-16
 SQLNET.CRYPTO_SEED, 4-18
 SQLNET.ENCRYPTION_SERVER, 4-18
 ssl.conf
 configuring for use with a server load
 balancer, 3-15
 Oracle HTTP Server configuration file, 3-15
 starting and stopping
 Enterprise Manager components, 2-1
 state directory
 in the Management Agent home, 10-2
 Statspack, 14-4
 Status Codes, Corrective Actions, 12-24, 12-26
 SYSMAN
 changing the SYSMAN password, 8-5
 checking for existence of, 8-9
 entering SYSMAN password when enabling
 security, 4-6
 sysman/admin/default_collection, 14-2
 sysman/emd/collection, 14-2

T

- target monitoring credentials
 - defined, 2-13
 - example of setting, 2-14
 - setting, 2-13
 - setting in Grid Control, 2-14
- targets
 - listing targets on a managed host, 2-15
- tasks
 - advanced configuration tasks, 1-1
- thresholds, 9-5, 9-7
- Top SQL Report
 - configuring the database to show the Top SQL Report, 14-4
- trace files
 - component tracing levels, 7-4
 - controlling the content of, 7-4
 - controlling the contents of Management Service, 7-8
 - controlling the size and number of, 7-6
 - controlling the size of, 7-3
 - fetchlet trace files, 7-4
 - locating Management Agent, 7-2
 - locating Management Service, 7-6
 - Management Agent, 7-1
 - Oracle Management Service, 7-6
 - rollover files, 7-2
- TrcFileMaxRolls property in emd.properties, 7-3
- TrcFileMaxSize property in emd.properties, 7-3
- Troubleshooting
 - when using EMCA, 1-17
- troubleshooting
 - general techniques while creating the Management Repository, 8-9
 - problems starting or configuring the Database Control, 1-17
 - while creating the Management Repository, 8-9
 - with EMCA, 1-17
- troubleshooting, notifications, 12-36
- trust points
 - Management Agent Configuration, 10-7

U

- UDP, 5-10
- uix-config.xml, 1-18
- upload directory
 - in the Management Agent home, 10-2, 10-3
- UploadMaxBytesXML
 - property in the emd.properties file, 10-3
- UploadMaxDiskUsedPct
 - property in the emd.properties file, 10-3
- Usage Metrics
 - Aggregation Function
 - Average, 6-8
 - Maximum, 6-8
 - Minimum, 6-8
 - Sum, 6-8
- User Datagram Protocol, 5-10
- User-defined metric page

- Command Line, 13-6
- Comparison Operator, 13-7, 13-12
- Consecutive Occurrences Preceding Notification, 13-8, 13-13
- Critical, 13-8, 13-13
- Environment, 13-7
- Metric Name, 13-6, 13-11
- Metric Type, 13-6, 13-11
- Operating System User Name and Password, 13-7
- Response Action, 13-8, 13-13
- Warning, 13-7, 13-12
- User-defined metric page, Central Console, 13-5, 13-10
- user-defined metric, example, 13-8
- User-defined metrics, 13-1
- user-defined metrics, 13-16

V

- virtual pool
 - when configuring a server load balancer, 3-14
- virtual service
 - when configuring a server load balancer, 3-14

W

- watchdog process
 - for the Management Agent, 10-4
- Web Application
 - Source
 - Step, 6-7
 - Step Group, 6-7
 - Transaction, 6-7
- Web Application target
 - using to monitor the Management Service response time, 3-7
- Web Applications
 - monitoring over HTTPS, 4-35
- Web Cache
 - See Oracle Application Server Web Cache
- Web Cache Availability and Critical/Warning States
 - default notification rule, 12-9
- web.xml, 1-19