Oracle® Enterprise Manager

Advanced Configuration

10*g* Release 3 (10.2.0.3.0)

B40002-01

February 2007

ORACLE®

Oracle Enterprise Manager Advanced Configuration, 10*g* Release 3 (10.2.0.3.0)

B40002-01

Copyright © 2003, 2007, 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 USA, Inc., 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 Siebel 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 xvii

Intended Audience xvii
Documentation Accessibility xvii
Related Documents xviii
Conventions xviii

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 10*g* 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 1-4
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 1-5
1.2.4 Understanding the Enterprise Manager Directories Installed with Oracle Database 10*g* 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 10*g* 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 1-11
1.2.6.4 Using an Input File for EMCA Parameters 1-15
1.2.6.5 Using EMCA with Real Application Clusters 1-16
1.2.6.6 Specifying the Ports Used By the Database Control 1-18
1.2.6.7 EMCA Troubleshooting Tips 1-19
1.2.6.7.1 Using EMCA After Changing the Database Listener Port 1-19
1.2.6.7.2 Upgrading Database or ASM Instances with 10g Release 2 Grid Control Agents 1-19
1.2.6.7.3 Using EMCA When Database Host Name or IP Address Changes 1-19

1.2.6.7.4 Using EMCA When the TNS Configuration Is Changed 1-19
1.3 Enabling Enterprise Manager Accessibility Features 1-20
1.3.1 Enabling Enterprise Manager Accessibility Mode 1-20
1.3.2 Providing Textual Descriptions of Enterprise Manager Charts 1-21

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 2-5
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 2-8
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 10*g* 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 2-18
2.7.6 Reevaluating Metric Collections 2-18

3 Grid Control Common Configurations

3.1 About Common Configurations 3-1
3.2 Deploying Grid Control Components on a Single Host 3-2
3.3 Managing Multiple Hosts and Deploying a Remote Management Repository 3-4
3.4 Using Multiple Management Service Installations 3-6
3.4.1 Understanding the Flow of Management Data When Using Multiple Management Services 3-6
3.4.2 Determining When to Use Multiple Management Service Installations 3-8

3.4.2.1 Monitoring the Load on Your Management Service Installations 3-9
3.4.2.2 Monitoring the Response Time of the Enterprise Manager Web Application Target 3-9
3.5 High Availability Configurations 3-10
3.5.1 Load Balancing Connections Between the Management Agent and the Management Service 3-11
3.5.2 Load Balancing Connections Between the Grid Control Console and the Management Service 3-12
3.5.2.1 Understanding the Flow of Data When Load Balancing the Grid Control Console 3-12
3.5.2.2 Configuring Oracle HTTP Server When Using a Load Balancer for the Grid Control Console 3-13
3.5.2.3 Configuring the Management Services 3-14
3.5.3 Configuring a Load Balancer 3-15
3.5.4 Configuring the Management Repository 3-17
3.6 Installation Best Practices for Enterprise Manager High Availability 3-18
3.6.1 Configuring the Management Agent to Automatically Start on Boot and Restart on Failure 3-18
3.6.2 Configuring Restart for the Management Agent 3-18
3.6.3 Installing the Management Agent Software on Redundant Storage 3-19
3.6.4 Install the Management Service Software on Redundant Storage 3-19
3.7 Configuration With Grid Control 3-19
3.7.1 Console Warnings, Alerts, and Notifications 3-19
3.7.2 Configure Additional Error Reporting Mechanisms 3-20
3.7.3 Component Backup 3-20
3.7.4 Troubleshooting 3-20
3.7.4.1 Upload Delay for Monitoring Data 3-20
3.7.4.2 Notification Delay of Target State Change 3-20

4 Configuring Oracle Enterprise Manager for Active and Passive Environments

4.1 Configuring Oracle Enterprise Management Agents for Use in Active and Passive Environments 4-1
4.1.1 Installation and Configuration 4-1
4.1.2 Switchover Steps 4-3
4.1.3 Performance Implications 4-3
4.1.4 Summary 4-4
4.2 Using Virtual Host Names for Active and Passive High Availability Environments in Enterprise Manager Database Control 4-4
4.2.1 Set Up the Alias for the Virtual Host Name and Virtual IP Address 4-4
4.2.2 Set Up Shared Storage 4-5
4.2.3 Set Up the Environment 4-5
4.2.4 Ensure That the Oracle USERNAME, ID, and GROUP NAME Are Synchronized on All Cluster Members 4-5
4.2.5 Ensure That Inventory Files Are on the Shared Storage 4-5
4.2.6 Start the Installer 4-6
4.2.7 Start Services 4-6

5 Enterprise Manager Security

5.1 About Oracle Enterprise Manager Security 5-1
5.1.1 Oracle Enterprise Manager Security Model 5-1
5.1.2 Classes of Users and Their Privileges 5-2
5.1.3 Resources Protected 5-2
5.1.4 Authorization and Access Enforcement 5-2
5.1.5 Leveraging Oracle Application Server Security Services 5-3
5.1.6 Leveraging Oracle Identity Management Infrastructure 5-3
5.2 Configuring Security for Grid Control 5-4
5.2.1 About Enterprise Manager Framework Security 5-4
5.2.2 Overview of the Steps Required to Enable Enterprise Manager Framework Security 5-5
5.2.3 Enabling Security for the Oracle Management Service 5-6
5.2.3.1 Checking the Security Status 5-8
5.2.4 Enabling Security for the Oracle Management Agent 5-9
5.2.5 Enabling Security with Multiple Management Service Installations 5-11
5.2.6 Restricting HTTP Access to the Management Service 5-11
5.2.7 Managing Agent Registration Passwords 5-13
5.2.7.1 Using the Grid Control Console to Manage Agent Registration Passwords 5-13
5.2.7.2 Using emctl to Change the Agent Registration Password 5-14
5.2.8 Enabling Security with a Server Load Balancer 5-14
5.2.9 Enabling Security for the Management Repository Database 5-15
5.2.9.1 About Oracle Advanced Security and the sqlnet.ora Configuration File 5-15
5.2.9.2 Configuring the Management Service to Connect to a Secure Management Repository Database 5-16
5.2.9.3 Enabling Oracle Advanced Security for the Management Repository 5-18
5.2.9.4 Enabling Security for a Management Agent Monitoring a Secure Management Repository or Database 5-18
5.3 Configuring Enterprise Manager for Use with Oracle Application Server Single Sign-On 5-19
5.3.1 Configuring Enterprise Manager to Use the Single Sign-On Logon Page 5-19
5.3.2 Registering Single Sign-On Users as Enterprise Manager Administrators 5-21
5.3.3 Grid Control as a Single Sign-On Partner Application 5-22
5.3.4 Bypassing the Single Sign-On Logon Page 5-23
5.4 Configuring Enterprise Manager for Use with Enterprise User Security 5-23
5.5 Setting Up the Auditing System for Enterprise Manager 5-24
5.5.1 Audit Data 5-24
5.5.2 Operation Codes 5-25
5.5.3 Audit APIs 5-26
5.5.4 Configuring the Enterprise Manager Audit System 5-26
5.6 Configuring the emkey 5-27
5.6.1 Generating the emkey 5-27
5.6.2 emctl Commands 5-28
5.6.2.1 emctl status emkey 5-28
5.6.2.2 emctl config emkey -repos 5-29
5.6.2.3 emctl config emkey -emkeyfile 5-29
5.6.2.4 emctl config emkey -emkey 5-30
5.6.2.5 emctl config emkey -remove_from_repos 5-30

5.6.2.6 emctl config emkey -copy_to_repos 5-30
5.6.3 Install and Upgrade Scenarios 5-31
5.6.3.1 Installing the Management Repository 5-31
5.6.3.2 Installing the First Oracle Management Service 5-31
5.6.3.3 Installing Additional Oracle Management Service 5-31
5.6.3.4 Upgrading from 10.1 to 10.2 5-31
5.6.3.5 Recreating the Management Repository 5-32
5.7 Additional Security Considerations 5-32
5.7.1 Responding to Browser-Specific Security Certificate Alerts 5-32
5.7.1.1 Responding to the Internet Explorer Security Alert Dialog Box 5-32
5.7.1.2 Responding to the Netscape Navigator New Site Certificate Dialog Box 5-34
5.7.1.3 Preventing the Display of the Internet Explorer Security Information Dialog Box 5-35
5.7.2 Configuring Beacons to Monitor Web Applications Over HTTPS 5-35
5.8 Other Security Features 5-37
5.8.1 Using ORACLE _HOME Credentials 5-37
5.8.2 Patching Oracle Homes When the User is Locked 5-39
5.8.3 Cloning Oracle Homes 5-39
5.8.4 Using the sudo Command 5-40

6 Configuring Enterprise Manager for Firewalls

6.1 Considerations Before Configuring Your Firewall 6-1
6.2 Firewall Configurations for Enterprise Management Components 6-2
6.2.1 Firewalls Between Your Browser and the Grid Control Console 6-3
6.2.2 Configuring the Management Agent on a Host Protected by a Firewall 6-4
6.2.2.1 Configuring the Management Agent to Use a Proxy Server 6-5
6.2.2.2 Configuring the Firewall to Allow Incoming Communication From the Management Service 6-6
6.2.3 Configuring the Management Service on a Host Protected by a Firewall 6-6
6.2.3.1 Configuring the Management Service to Use a Proxy Server 6-7
6.2.3.2 About the dontProxyfor Property 6-8
6.2.3.3 Configuring the Firewall to Allow Incoming Management Data From the Management Agents 6-9
6.2.4 Firewalls Between the Management Service and the Management Repository 6-9
6.2.5 Firewalls Between the Grid Control and a Managed Database Target 6-10
6.2.6 Firewalls Used with Multiple Management Services 6-10
6.2.7 Configuring Firewalls to Allow ICMP and UDP Traffic for Beacons 6-11
6.2.8 Configuring Firewalls When Managing Oracle Application Server 6-11
6.3 Viewing a Summary of the Ports Assigned During the Application Server Installation 6-12
6.4 Additional Considerations for Windows XP 6-12

7 Configuring Services

7.1 Summary of Service Management Tasks 7-1
7.2 Setting up the System 7-3
7.3 Creating a Service 7-4
7.4 Configuring a Service 7-5

7.4.1 Availability Definition ... 7-6
7.4.2 Performance Metrics ... 7-7
7.4.3 Usage Metrics ... 7-8
7.4.4 Business Metrics ... 7-9
7.4.5 Service Tests and Beacons ... 7-9
7.4.5.1 Configuring the Beacons ... 7-10
7.4.6 Root Cause Analysis Configuration ... 7-11
7.4.6.1 Getting the Most From Root Cause Analysis ... 7-12
7.5 Recording Web Transactions ... 7-12
7.6 Monitoring Settings ... 7-13
7.7 Configuring Aggregate Services ... 7-14
7.8 Configuring End-User Performance Monitoring ... 7-14
7.8.1 Configuring End-User Performance Monitoring Using Oracle HTTP Server Based on Apache 2.0 or Apache HTTP Server 2.0 ... 7-15
7.8.1.1 Setting up the Third Party Apache Server ... 7-17
7.8.2 Configuring End-User Performance Monitoring Using Oracle Application Server Web Cache ... 7-17
7.8.2.1 Configuring Oracle Application Server Web Cache 10.1.2 ... 7-18
7.8.2.2 Configuring Oracle Application Server Web Cache 9.0.4 ... 7-19
7.8.2.3 Configuring End-User Performance Monitoring Using Earlier Versions of Oracle Application Server Web Cache ... 7-20
7.8.2.3.1 Using the chronos_setup.pl Configuration Script ... 7-20
7.8.2.3.2 Configuring the Document Root For Each Web Server ... 7-21
7.8.2.3.3 Configuring Oracle Application Server Web Cache for End-User Performance Monitoring ... 7-21
7.8.2.3.4 Starting End-User Performance Monitoring ... 7-22
7.8.2.4 Configuring End-User Performance Monitoring Using Standalone Oracle Application Server Web Cache ... 7-23
7.8.2.4.1 Installing Standalone Oracle Application Server Web Cache ... 7-23
7.8.2.4.2 Configuring Standalone Oracle Application Server Web Cache ... 7-24
7.8.2.4.3 Enabling End-User Performance Monitoring for Standalone Oracle Application Server Web Cache ... 7-24
7.8.3 Starting and Stopping End-User Performance Monitoring ... 7-25
7.8.4 Verifying and Troubleshooting End-User Performance Monitoring ... 7-25
7.8.5 Enabling End-User Performance Monitoring for Third-Party Application Servers 7-27
7.9 Managing Forms Applications ... 7-28
7.9.1 Recording and Monitoring Forms Transactions ... 7-28
7.9.1.1 Setting the Permissions of the .java.policy File ... 7-28
7.9.1.2 Using a Trusted Enterprise Manager Certificate ... 7-29
7.9.1.3 Adding a Forms Certificate to the Enterprise Manager Agent ... 7-30
7.9.1.4 Configuring the Forms Server ... 7-30
7.9.1.5 Installing the Transaction Recorder to Record and Play Back Forms Transactions ... 7-31
7.9.1.6 Forms Server Version Support ... 7-31
7.9.2 Monitoring the End-User Performance of Forms Applications ... 7-32
7.9.2.1 Configuring the Forms Server for End-User Performance Monitoring ... 7-32
7.9.2.2 Configuring the OracleAS Web Cache ... 7-33
7.9.2.3 Configuring the Oracle HTTP Server / Apache HTTP Server ... 7-35

7.9.2.4 Starting and Stopping End-User Performance Monitoring 7-36
7.10 Configuring OC4J for Request Performance Diagnostics 7-36
7.10.1 Selecting OC4J Targets for Request Performance Diagnostics 7-37
7.10.2 Configuring Interactive Transaction Tracing 7-37
7.10.3 Configuring OC4J Tracing for Request Performance Data 7-38
7.10.4 Additional Configuration for Monitoring UIX Applications 7-39
7.11 Setting Up Monitoring Templates 7-39
7.11.1 Configuring Service Tests and Beacons 7-40
7.12 Configuring Service Levels 7-41
7.12.1 Defining Service Level Rules 7-42
7.12.2 Viewing Service Level Details 7-42
7.13 Configuring a Service Using the Command Line Interface 7-43

8 Locating and Configuring Enterprise Manager Log Files

8.1 Locating and Configuring Management Agent Log and Trace Files 8-1
8.1.1 About the Management Agent Log and Trace Files 8-1
8.1.2 Locating the Management Agent Log and Trace Files 8-2
8.1.3 About Management Agent Rollover Files 8-2
8.1.4 Controlling the Size and Number of Management Agent Log and Trace Files 8-3
8.1.5 Controlling the Contents of the Management Agent Trace File 8-4
8.1.6 Controlling the Size and Number of Fetchlet Log and Trace Files 8-4
8.1.7 Controlling the Contents of the Fetchlet Trace File 8-5
8.2 Locating and Configuring Management Service Log and Trace Files 8-6
8.2.1 About the Management Service Log and Trace Files 8-6
8.2.2 Locating the Management Service Log and Trace Files 8-6
8.2.3 Controlling the Size and Number of Management Service Log and Trace Files 8-6
8.2.4 Controlling the Contents of the Management Service Trace File 8-8
8.2.5 Controlling the Oracle Application Server Log Files 8-8

9 Maintaining and Troubleshooting the Management Repository

9.1 Management Repository Deployment Guidelines 9-1
9.2 Changing the SYSMAN Password 9-2
9.3 Dropping and Recreating the Management Repository 9-3
9.3.1 Dropping the Management Repository 9-3
9.3.2 Recreating the Management Repository 9-4
9.3.2.1 Using the RepManager Script to Create the Management Repository 9-4
9.3.2.2 Using a Connect Descriptor to Identify the Management Repository Database 9-5
9.4 Troubleshooting Management Repository Creation Errors 9-5
9.4.1 Package Body Does Not Exist Error While Creating the Management Repository 9-5
9.4.2 Server Connection Hung Error While Creating the Management Repository 9-6
9.4.3 General Troubleshooting Techniques for Creating the Management Repository 9-6
9.5 Improving the Login Performance of the Console Home Page 9-6

10 Using Enterprise Manager For Grid Automation With Deployment Procedures

10.1 Key Advantages of Deployment Procedures 10-1

10.1.1 Deployment Procedures Shipped In Oracle Enterprise Manager 10-2
10.2 Deployment Procedure Requirements..... 10-2
10.3 Use-Cases for Deployment Procedures 10-3
10.3.1 Using Deployment Procedures to Apply Security-Related Critical Patch Updates to Oracle Databases 10-3
10.3.2 Using Deployment Procedures for Single-Click Extend of Real Application Clusters 10-4
10.3.3 Using Deployment Procedures for Delete/Scale Down of Real Application Clusters 10-4
10.4 Customizable Deployment Procedures 10-4
10.4.1 Customization Examples 10-4
10.4.1.1 Insert Custom Step to Backup the Database Before Patching 10-5
10.4.1.2 Manual Step 10-5
10.4.1.3 Application Service Shutdown and Startup Handling 10-5
10.4.1.4 Set Notification for the Deployment Procedure Run 10-5
10.4.2 Importing or Exporting Deployment Procedures 10-5
10.5 Executing Deployment Procedures Using Pluggable Authentication Modules (PAM) and SUDO 10-7
10.5.1 Design Time Experience Example (One Time)..... 10-7
10.5.2 Run Time Experience (Every Time) 10-8
10.6 Deployment Procedure Variables..... 10-8
10.7 Using EMCLI to Execute Deployment Procedures..... 10-8
10.7.1 Creating the RuntimeData Templates 10-9
10.7.1.1 Step 1: Locate the Successful Test Run of Your Customized Procedure 10-9
10.7.1.2 Step 2: Download the Runtime Data For This Run 10-10
10.7.1.3 Step 3: Locate and Create Instantiation Variables 10-10
10.7.2 Using the RuntimeData XML Instead Of RuntimeData Templates..... 10-11
10.7.3 Creating a Properties File for Provisioning Procedures..... 10-12
10.7.3.1 Properties File Variables for Provisioning Procedures 10-12
10.7.3.2 Properties File Variables for Patching Procedures 10-14
10.7.3.3 Updating the Runtime Data Template to Fit Various Patching Use Cases..... 10-15
10.7.3.4 Getting the Host and Homes From the Instances For Patching Runtime Data 10-16
10.7.4 Creating the Properties File for Patching Procedures 10-16
10.7.5 Submitting the Deployment Procedure..... 10-16
10.8 Known Issues and Troubleshooting..... 10-17
10.8.1 Known Issues 10-17
10.8.2 Troubleshooting 10-17

11 Sizing and Maximizing the Performance of Oracle Enterprise Manager

11.1 Oracle Enterprise Manager Grid Control Architecture Overview 11-1
11.2 Enterprise Manager Grid Control Sizing and Performance Methodology 11-2
11.2.1 Step 1: Choosing a Starting Platform Grid Control Deployment 11-3
11.2.1.1 Network Topology Considerations 11-4
11.2.2 Step 2: Periodically Evaluate the Vital Signs of Your Site 11-5
11.2.3 Step 3: Use DBA and Enterprise Manager Tasks To Eliminate Bottlenecks Through Housekeeping 11-7
11.2.3.1 Online Weekly Tasks..... 11-7

11.2.3.2 Offline Monthly Tasks 11-8
11.2.4 Step 4: Eliminate Bottlenecks Through Tuning 11-10
11.2.4.1 High CPU Utilization 11-10
11.2.4.2 Loader Vital Signs 11-11
11.2.4.3 Rollup Vital Signs 11-12
11.2.4.4 Job, Notification, and Alert Vital Signs 11-13
11.2.4.5 I/O Vital Signs 11-13
11.2.4.6 The Oracle Enterprise Manager Performance Page 11-14
11.2.5 Step 5: Extrapolating Linearly Into the Future for Sizing Requirements 11-15
11.3 Oracle Enterprise Manager Backup, Recovery, and Disaster Recovery Considerations 11-15
11.3.1 Best Practices for Backup and Recovery 11-16
11.3.1.1 Oracle Management Service 11-16
11.3.1.2 Management Agent 11-17
11.3.2 Best Practice for Disaster Recovery (DR) 11-17
11.3.2.1 Management Repository 11-17
11.3.2.2 Oracle Management Service 11-17
11.3.2.3 Management Agent 11-18

12 Reconfiguring the Management Agent and Management Service

12.1 Reconfiguring the Oracle Management Agent 12-1
12.1.1 Configuring the Management Agent to Use a New Management Service 12-1
12.1.2 Changing the Management Agent Port 12-2
12.1.3 Controlling the Amount of Disk Space Used by the Management Agent 12-3
12.1.4 About the Management Agent Watchdog Process 12-4
12.1.5 Setting the Management Agent Time Zone 12-4
12.1.5.1 Understanding How the Management Agent Obtains Time Zone Information 12-4
12.1.5.2 Resetting the Time Zone of the Management Agent Due to ...Inconsistency of Time Zones 12-5
12.1.5.3 Troubleshooting Management Agent Time Zone Problems 12-5
12.1.5.4 Troubleshooting Management Service Time Zone Problems 12-6
12.1.6 Adding Trust Points to the Management Agent Configuration 12-7
12.2 Reconfiguring the Oracle Management Service 12-8
12.2.1 Configuring the Management Service to Use a New Management Repository 12-8
12.2.1.1 Changing the Repository Properties in the emoms.properties File 12-8
12.2.1.2 About Changing the Repository Password 12-9
12.2.2 Configuring the Management Service to Use a New Port 12-9
12.2.3 Configuring the Management Service to Prompt You When Using Execute Commands 12-10

13 Configuring Notifications

13.1 Setting Up Notifications 13-1
13.1.1 Setting Up a Mail Server for Notifications 13-1
13.1.2 Setting Up E-mail for Yourself 13-4
13.1.2.1 Defining E-mail Addresses 13-4
13.1.2.2 Setting Up a Notification Schedule 13-5
13.1.2.3 Subscribe to Receive E-mail for Notification Rules 13-6

13.1.3 Setting Up E-mail for Other Administrators 13-9
13.2 Extending Notification Beyond E-mail 13-10
13.2.1 Custom Notification Methods Using Scripts and SNMP Traps 13-11
13.2.1.1 Adding a Notification Method based on an OS Command or Script 13-11
13.2.1.2 Adding a Notification Method Based on a PL/SQL Procedure 13-15
13.2.1.3 Adding a Notification Method Based on an SNMP Trap 13-19
13.3 Passing Corrective Action Status Change Information 13-21
13.3.1 Passing Corrective Action Execution Status to an OS Command or Script 13-21
13.3.2 Passing Corrective Action Execution Status to a PLSQL Procedure 13-22
13.4 Passing Job Execution Status Information 13-24
13.4.1 Passing Job Execution Status to a PLSQL Procedure 13-25
13.4.2 Passing Job Execution Status to an OS Command or Script 13-26
13.5 Assigning Methods to Rules 13-27
13.6 Assigning Rules to Methods 13-28
13.7 Management Information Base (MIB) 13-29
13.7.1 About MIBs 13-29
13.7.2 Reading the MIB Variable Descriptions 13-30
13.7.2.1 Variable Name 13-30
13.7.2.2 MIB Definition 13-30
13.8 Troubleshooting Notifications 13-36
13.8.1 General Setup 13-37
13.8.2 Notification System Errors 13-37
13.8.3 Notification System Trace Messages 13-37
13.8.4 E-mail Errors 13-39
13.8.5 OS Command Errors 13-39
13.8.6 SNMP Trap Errors 13-39
13.8.7 PL/SQL Errors 13-39

14 User-Defined Metrics

14.1 Extending Monitoring Capability 14-1
14.2 Creating OS-Based User-Defined Metrics 14-2
14.2.1 Create Your OS Monitoring Script 14-2
14.2.1.1 Code to check the status of monitored objects 14-2
14.2.1.2 Code to return script results to Enterprise Manager 14-2
14.2.1.3 Script Runtime Environment 14-4
14.2.2 Register the Script as a User-Defined Metric 14-5
14.2.3 OS-Based User-Defined Metric Example 14-8
14.3 Creating a SQL-Based User-Defined Metric 14-10
14.3.1 SQL-Based User-Defined Metric Examples 14-14
14.3.1.1 Example 1: Query Returning Tablespace Name and Percent Used 14-14
14.3.1.2 Example 2: Query Returning Segment Name/Type and Extent Count 14-15
14.4 Notifications, Corrective Actions, and Monitoring Templates 14-15
14.4.1 Getting Notifications for User-Defined Metrics 14-15
14.4.2 Setting Corrective Actions for User-Defined Metrics 14-16
14.4.3 Deploying User-Defined Metrics Across Many Targets Using Monitoring Templates 14-16
14.4.4 Deleting User-Defined Metrics Across Many Targets Using Monitoring Templates

..... 14-18

15 Using a Software Library

15.1 Overview of Software Library 15-1
15.2 Setting up the Software Library 15-2
15.3 Configuring the Software Library For Use With Oracle Enterprise Manager 15-2
15.4 Creating and Deleting Entities in the Software Library 15-2
15.5 Exporting and Importing Software Library Entities Across Enterprise Manager Deployments 15-3
15.6 Exporting and Importing Entities Across Oracle Enterprise Manager Deployments 15-3
15.6.1 How the Importing and Exporting Scripts Funtion 15-4
15.6.1.1 Export Script 15-4
15.6.1.2 Import Script 15-4
15.6.2 Use Cases for Import and Export 15-5
15.7 Alternate Approach for Transferring a Software Library Across Different Oracle Enterprise Manager Deployments 15-6
15.8 De-Configuring a Software Library 15-6
15.9 Deleting Entities and Purging the Contents of the Software Library 15-6
15.10 Software Library Issues 15-7

16 Configuring Remedy Connector

16.1 Introduction to Remedy Connector 16-1
16.1.1 Autoticketing 16-2
16.1.2 Manual Ticketing 16-2
16.1.3 Ticket Templates 16-2
16.1.4 Grace Period 16-2
16.2 Prerequisites 16-3
16.3 Installing Remedy Connector 16-3
16.4 Configuring the Remedy Connector 16-3
16.4.1 General Settings 16-6
16.4.1.1 Connection Settings 16-6
16.4.1.2 Web Console Settings 16-7
16.4.1.3 Grace Period 16-7
16.4.2 Working with Ticket Templates 16-7
16.4.2.1 Registering Templates 16-7
16.4.2.2 Viewing Template Code 16-7
16.4.2.3 Removing Template 16-8
16.4.2.4 Replacing Templates 16-8
16.4.2.5 Adding New Templates 16-8
16.5 Creating Remedy Trouble Tickets 16-8
16.5.1 Creating Trouble Ticket Automatically 16-8
16.5.2 Creating a Ticket Manually 16-11
16.6 Enabling SSL for HTTPS 16-13
16.6.1 Generating a Certificate Request File 16-13
16.6.2 Importing the Certificate from the Certificate Authority 16-13
16.6.3 Adding Signed Certificates to Wallet Manager 16-13

16.7 Navigating Between Remedy and Enterprise Manager 16-14
16.7.1 Navigating from Remedy to Enterprise Manager 16-14
16.7.2 Navigating from the Enterprise Manager to Remedy 16-14
16.8 Out-of-Box Templates 16-15
16.8.1 Reading Ticket Templates 16-16
16.8.1.1 Templates Explained 16-21
16.8.2 Customizing Ticket Templates 16-63
16.8.3 Defining New Templates 16-63
16.9 Remedy Connector Tips 16-67
16.9.1 Recommended Protocol 16-67
16.9.2 Supported Alerts 16-67
16.9.3 Notification Failure 16-67
16.9.4 Using Worklog 16-67
16.9.5 Web Service Details 16-68
16.9.5.1 For Default Templates (without Worklog Support) 16-68
16.9.5.2 For Worklog Templates 16-68

17 Configuring Microsoft Operations Management Connector

17.1 Introduction to MOM Connector 17-1
17.2 Installing MOM Connector 17-1
17.3 Prerequisites 17-2
17.4 Configuring the MOM Connector 17-2
17.4.1 General Settings 17-4
17.4.2 Additional Target Instances 17-4
17.4.2.1 Response Status of Targets 17-6
17.5 Enabling SSL Connection for HTTPS 17-6
17.5.1 Generating a Certificate Request File 17-6
17.5.2 Using the Certificate from the Certificate Authority 17-7
17.5.3 Adding Signed Certificates to Wallet Manager 17-7
17.6 MOM Connector Tips 17-8
17.6.1 Recommended Protocol 17-8
17.6.2 Connector Configuration Fails 17-8
17.6.3 MOM Connector Fails to Retrieve Alerts 17-8
17.6.4 Alert Logging to Additional Target Instance Fails 17-9
17.6.5 Adding Targets in the Same Transaction 17-9
17.6.6 Polling Interval 17-9
17.6.7 Alerts per Polling 17-9
17.6.8 Metric/Target Does Not Exist 17-9

18 Additional Configuration Tasks

18.1 Understanding Default and Custom Data Collections 18-1
18.1.1 How Enterprise Manager Stores Default Collection Information 18-1
18.1.2 Restoring Default Collection Settings 18-2
18.2 Enabling Multi-Inventory Support for Configuration Management 18-2
18.2.1 AGENT_HOME Versus AGENT_STATE Directories 18-3
18.3 Manually Configuring a Database Target for Complete Monitoring 18-4
18.4 Modifying the Default Login Timeout Value 18-6

18.5 Configuring Clusters and Cluster Databases in Grid Control 18-7
18.5.1 Configuring Clusters 18-7
18.5.2 Configuring Cluster Databases 18-8
18.5.3 Discovering Instances Added to the Cluster Database 18-8
18.5.3.1 Troubleshooting 18-9
18.6 Collecting Client Configurations 18-9
18.6.1 Configuring the Client System Analyzer 18-10
18.6.1.1 Client System Analyzer in Oracle Grid Control 18-10
18.6.1.2 Deploying Client System Analyzer Independently 18-10
18.6.2 Configuration Parameters 18-12
18.6.2.1 Associating the Parameters with an Application 18-15
18.6.3 Rules 18-15
18.6.4 Customization 18-17
18.6.5 CSA Deployment Examples 18-17
18.6.5.1 Using Multiple Collection Tags 18-17
18.6.5.2 Privilege Model for Viewing Client Configurations 18-18
18.6.5.3 Using the Customization API Example 18-19
18.6.5.4 Using the CSA Servlet Filter Example 18-20
18.6.5.5 Sample Deployments 18-21
18.6.5.5.1 Example 1: Helpdesk 18-21
18.6.5.5.2 Example 2: Inventory 18-22
18.6.5.5.3 Example 3: Problem Detection 18-23
18.7 Setting Up and Configuring a Software Library With Oracle Enterprise Manager 18-24
18.7.1 Setting Up a Software Library 18-24
18.7.2 Configuring a Software Library 18-24
18.7.3 Deleting or Cleaning Up a Software Library 18-24

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 10*g*. 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. Accessibility 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 more information, visit the Oracle Accessibility Program Web site at

`http://www.oracle.com/accessibility/`

Accessibility of Code Examples in Documentation

Screen readers 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, some screen readers 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.

TTY Access to Oracle Support Services

Oracle provides dedicated Text Telephone (TTY) access to Oracle Support Services within the United States of America 24 hours a day, seven days a week. For TTY support, call 800.446.2398.

Related Documents

For more information, see the following manuals in the Oracle Enterprise Manager 10*g* 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 Framework, Host, and Third-Party Metric Reference Manual*
- *Oracle Enterprise Manager Oracle Database and Database-Related Metric Reference Manual*
- *Oracle Enterprise Manager Oracle Application Server Metric Reference Manual*

  Oracle Enterprise Manager Oracle Collaboration Suite Metric Reference Manual
- *Oracle Enterprise Manager Extensibility*
- *Oracle Enterprise Manager Command Line Interface*
- *Oracle Enterprise Manager SNMP Support Reference Guide*
- *Oracle Enterprise Manager Licensing Information*

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

1

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 10*g* 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 10*g* Grid Control

When you install Oracle Enterprise Manager 10*g* 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 10*g* 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 in the form of an OC4J instance (OC4J_EM) that is installed and deployed using the Oracle Application Server J2EE and Web Cache installation type.

The installation procedure installs the Enterprise Manager components within the Oracle Application Server Home, including the Oracle Management Service.

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

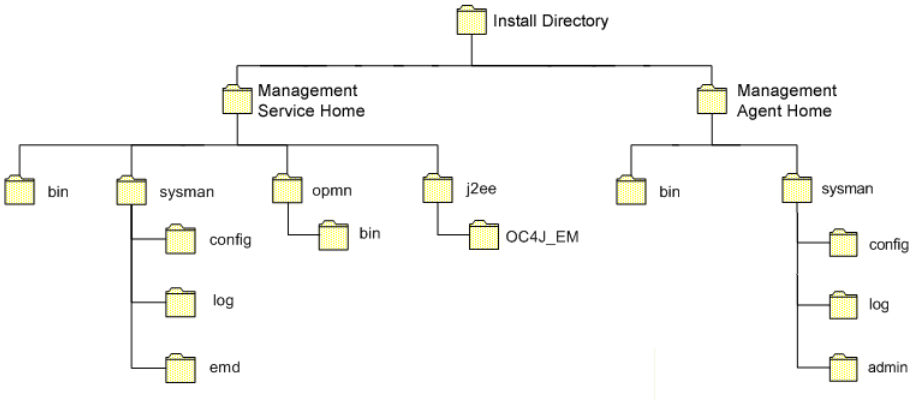

Table 1–1 describes in more detail the Management Service directories shown in Figure 1–1. In the table, ORACLE_HOME refers to the Management Service home directory in which the Oracle Management Service is installed and deployed.

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

Directory	Description
`ORACLE_HOME/bin`	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. Use the `emctl` command in this directory to start and stop the Application Server Control Console. For more information about the Application Server Control Console, see the *Oracle Application Server 10g Administrator's Guide*.

Table 1–1 (Cont.) Important Directories in the Management Service Oracle Home

Directory	Description
`ORACLE_HOME/sysman`	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. Note that the `ORACLE_HOME/sysman/log` directory contains the Oracle Management Service log files (`emoms.log`) and trace files (`emoms.trc`).
`ORACLE_HOME/opmn`	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.
`ORACLE_HOME/j2ee`	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.
ORACLE_HOME/hostname	For real application cluster agent install, this directory contains sysman files.

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`	The `agent` directory contains all the files required to configure and run the Oracle Management Agent on this host. This directory serves as the Oracle Home for the Management Agent. Later in this document, this directory is referred to as the AGENT_HOME. 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.
`AGENT_HOME/bin`	The `agent/bin` directory in the Oracle Application Server Home contains the `emctl` command that controls the Management Agent for this host. You use the `emctl` command in this directory to start and stop the Oracle Management Agent on this host.
`AGENT_HOME/sysman/admin`	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.
`AGENT_HOME/sysman/config`	This directory contains the configuration files for the Management Agent. For example, this is where Enterprise Manager stores the `emd.properties` file. The `emd.properties` file defines settings such as the Management Service upload URL for this particular agent.
`AGENT_HOME/sysman/log`	This directory contains the log files for the Management Agent.
AGENT_HOME/hostname	For real application clusters, this directory contains all configuration, log files, and system files.

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 10*g* 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 8-1

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

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

- Files that are common and shared among all instances of the database are stored in the following directory of the Oracle Database 10*g* 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 10*g* 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 8-1

If a *hostname_sid* directory does not exist in the Oracle Database 10*g* home directory, then Oracle Enterprise Manager 10*g* 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.

The installation and configuration files are stored in the `ORACLE_HOME` directory in the following sub-directories:

- `cfgtoollogs/cfgfw`
- `cfgtoollogs/dbua`
- `cfgtoollogs/netca`
- `cfgtoollogs/rconfig`
- `cfgtoollogs/dbca`
- `cfgtoollogs/emca`
- `cfgtoollogs/oui`
- `cfgtoollogs/opatch`

Figure 1–2 summarizes the location of the important Enterprise Manager directories in a typical Oracle Database 10*g* 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

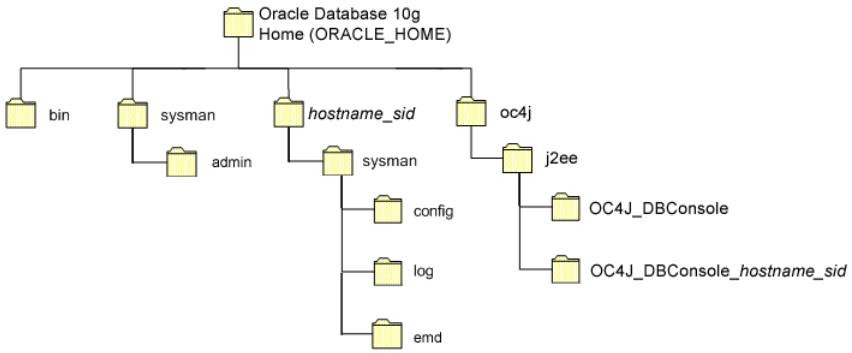

1.2.5 Tip for Identifying the Oracle Home When Using the emctl Command

When you install Grid Control, Oracle Application Server, or Oracle Database 10*g*, 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 10*g* Installation

The following sections describe how Oracle Enterprise Manager 10*g* Database Control is configured during the Oracle Database 10*g* 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 10*g*, you have the option of configuring your database so it can be managed by Oracle Enterprise Manager 10*g* Grid Control Console or by Oracle Enterprise Manager 10*g* Database Control Console.

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

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

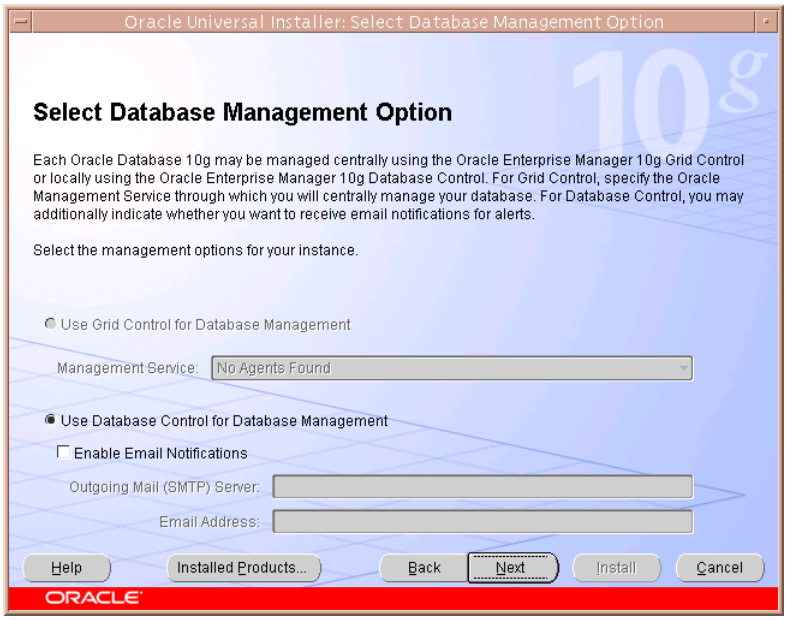

To select Grid Control Console as your management option, the Oracle Management Service must be installed on a network host. In addition, the Oracle Management Agent must be installed on the host where you are installing the database. Otherwise, the Grid Control Console option is unavailable and you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

For most of the Oracle Database 10*g* 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 10*g* 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 10*g* 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

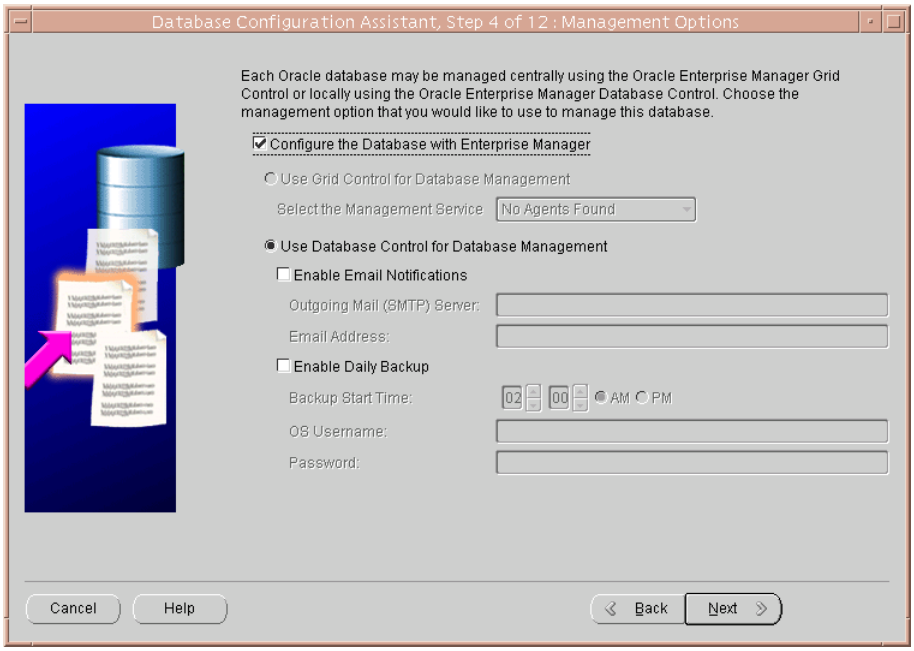

1.2.6.3 Configuring Database Control with EMCA

When you use DBCA to configure Oracle Database 10*g*, 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] [flags] [parameters]
```

> **Note:** To configure Database Console for single instance database using ASM, no extra parameters need to be passed along with the EMCA command. Run the following command to configure the Database Console which will automatically detect the ASM instance:
>
> ```
> emca -config dbcontrol db -repos create
> ```

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, andTable 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.
emca -deconfig centralAgent (db | asm) [-cluster] [-silent] [parameters]	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.

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

Command	Description
emca -deconfig all db [-repos drop] [-cluster] [-silent] [parameters]	Deconfigures both Database Control and central agent management for a database. The possible deconfiguration options are similar to those described above.
emca -addInst (db | asm) [-silent] [parameters]	Configures Enterprise Manager for a new cluster instance of a database or ASM storage. For more information, refer to Section 1.2.6.5.
emca -deleteInst (db | asm) [-silent] [parameters]	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.
emca -reconfig ports [-cluster] [parameters]	Explicitly reassigns Database Control ports. Options include performing this operation on a cluster environment. For more information, refer to Section 1.2.6.6.
emca -reconfig dbcontrol -cluster [-silent] [parameters]	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.
emca -displayConfig dbcontrol -cluster [-silent] [parameters]	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.
emca -upgrade (db | asm | db_asm) [-cluster] [-silent] [parameters]	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).
emca -restore (db | asm | db_asm) [-cluster] [-silent] [parameters]	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
db	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).
asm	Performs the operation for an ASM-only instance (including cluster ASM instances).

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

Flag	Description
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 –respFile argument. You can view a list of the available parameters by entering emca -help 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 db_recovery_file_dest 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
-MAIL_SERVER_NAME	Outgoing Mail (SMTP) server for notifications
-ASM_OH	Automatic Storage Management Oracle Home

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

Parameter	Description
-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=KDOdk432
```

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=dow3l224
SYSMAN_PWD=squN3243
HOST_USER=johnson
HOST_USER_PWD=diTf32of
SYS_PWD=qlKj4352
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. On each instance of the RAC database, the following subdirectories will be created:

```
$ORACLE_HOME/hostname1.domainname_SID1
$ORACLE_HOME/hostname2.domainname_SID2
.
.
$ORACLE_HOME/hostnamen.domainname_SIDn
```

where <SID1>...<SIDn> are database system identifiers.

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 10g Release 1 Database Control has a Database Console running on each node for the real-application cluster. 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.

For configuring Enterprise Manager for a new cluster instance of a database or ASM storage, use the following command:

```
emca -addInst db
```

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. After running the command `emca -addInst db`, enter the following information for the node and database:

```
Node name: node2
Database Unique Name: EM102
Database SID: EM1022
```

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.

Caution: If you use `emca -c` to configure the Database Control for Real Application Clusters, check TNS_ADMIN on all cluster nodes. If different TNS_ADMIN are set for each node, the listener for the target acnnot be configured correctly. If so, set the same TNS_ADMIN on all cluster nodes before executing the `emca -c` command.

1.2.6.6 Specifying the Ports Used By the Database Control

When you initially install Oracle Database 10*g* 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 10*g*.

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. The default port can be changed if the user wants to configure a specific port for Database Console. When

a port other than the default port (1521) is used, use the -RMI_PORT or -JMS_PORT options along with the emca reconfig command.

- -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. The default port can be changed if the user wants to configure a specific port for Database Console. When a port other than the default port (1521) is used, use the -RMI_PORT or -JMS_PORT options along with the emca reconfig command.

- -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.

1.2.6.7.3 Using EMCA When Database Host Name or IP Address Changes When the database host name (including the domain name) or the IP address changes, deconfigure and then reconfigure the Database Console with the repository create command. Run the following command:

```
emca -deconfig dbcontrol db -repos drop
emca -config dbcontrol db -repos create
```

or

```
emca -deconfig dbcontrol db
emca -config dbcontrol db -repos recreate
```

1.2.6.7.4 Using EMCA When the TNS Configuration Is Changed When the TNS configuration is changed, set the environment variable and then run the following command:

```
emca -config dbcontrol db
```

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 10*g* 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 10*g* 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

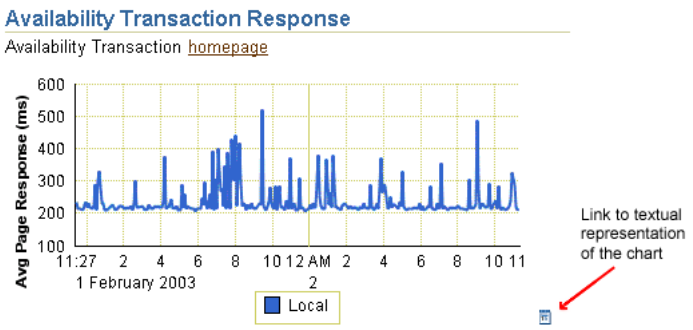

To enable the drill-down icon for the textual representation of charts:

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 10*g* 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 10*g* installation).

2

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>
```

On IBM AIX environment with a large memory configuration where the Management Agent is monitoring a large number of targets, the Agent may not start. To prevent this issue, prior to starting the Management Agent, set the following variables in the shell:

```
LDR_CNTRL="MAXDATA=0x80000000"@NOKRTL
AIX_THREADSCOPE=S
```

The `LDR_CNTRL` variable sets the data segment size and disables loading of run time libraries in kernel space. The `AIX_THREADSCOPE` parameter changes AIX Threadscope context from the default Processwide 'P' to Systemwide 'S'. This causes less mutex contention.

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 *Oracle SNMP Support Reference Guide*.

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 *Oracle SNMP Support Reference Guide*.

> **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	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. **Note:** The `emctl start oms` 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.
emctl stop oms	Stops the Management Service. Note that this command does not stop the other processes that are managed by the Oracle Process Manager and Notification Server (OPMN) utility. 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.
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 10*g*, 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 7, "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 10*g* 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.
2. 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
`opmnctl startproc ias-component=WebCache`	Starts Oracle Application Server Web Cache.
`opmnctl stopproc ias-component=WebCache`	Stops Oracle Application Server Web Cache.
`opmnctl status`	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.

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

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 10*g* 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 10*g*, 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 10*g*, 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 10*g* home directory.
2. Wait 15 seconds.
3. Use the emctl start iasconsole command to start the Application Server Control in the OracleAS Infrastructure 10*g* 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 10*g* 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 10*g* 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
`emctl upload`	Use this command to force an immediate upload of the current management data from the managed host to the Management Service. Use this command instead of waiting until the next scheduled upload of the data.
`emctl reload`	This command can be used to modify the `emd.properties` file. For example, to change the upload interval, emd.properties can be modified, and `emctl reload` 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 `targets.xml` 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 `targets.xml` 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 9-2.

When you first add an Oracle9*i* 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 10*g*, you specify the DBSNMP monitoring password during the database installation procedure.

As a result, if the password for the DBSNMP database user account is changed, you must modify the properties of the database target so that Enterprise Manager can continue to connect to the database and gather configuration and performance data.

Similarly, immediately after you add a new Oracle Database 10*g* 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 10*g* 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 10*g* 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]
[emrep10.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 10*g* 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	`emctl start blackout <Blackoutname> [<Target_name>[:<Target_Type>]]....` `[-d <Duration>]` Be sure to use a unique name for the blackout so you can refer to it later when you want to stop or check the status of the blackout. The `-d` option is used to specify the duration of the blackout. Duration is specified in `[days] hh:mm where:` ■ days indicates number of days, which is optional ■ hh indicates number of hours ■ mm indicates number of minutes If you do not specify a target or list of targets, Enterprise Manager will blackout the local host target. All monitored targets on the host are not blacked out unless a list is specified or you use the `-nodelevel` argument. If two targets of different target types share the same name, you must identify the target with its target type.
Stop an immediate blackout	`emctl stop blackout <Blackoutname>`
Set an immediate blackout for all targets on a host	`emctl start blackout <Blackoutname> [-nodeLevel] [-d <Duration>]` The `-nodeLevel` option is used to specify a blackout for all the targets on the host; in other words, all the targets that the Management Agent is monitoring, including the Management Agent host itself. The `-nodeLevel` option must follow the blackout name. If you specify any targets after the `-nodeLevel` option, the list is ignored.
Check the status of a blackout	emctl status blackout `[<Target_name>[:<Target_Type>]]....`

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

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

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

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

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

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

- To stop blackout "bk2" immediately:

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

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

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

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

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

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

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

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 12.1.5, "Setting the Management Agent Time Zone" on page 12-4

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> <colletionItemName>
```

where `<collectionItemName>` is the name of the Collection Item that collects the metric.

Performing this command causes the reevaluated value of the metric to be uploaded into the Management Repository, and possibly trigger alerts if the metric crosses its threshold.

Related metrics are typically collected together; collectively a set of metrics collected together is called a Metric Collection. Each Metric Collection has its own name. If you want to reevaluate a metric, you first need to determine the name of the Metric Collection to which it belongs, then the CollectionItem for that Metric Collection.

When you run the previous command to reevaluate the metric, all other metrics that are part of the same Metric Collection and Collection Item will also be reevaluated.

Perform the following steps to determine the Metric Collection name and Collection Item name for a metric:

1. Go to $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.

3

Grid Control Common Configurations

Oracle Enterprise Manager 10*g* 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.

The common configurations are presented 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 load balancers, Oracle Real Application Clusters, and Oracle Data Guard.

This chapter contains the following sections:

- About Common Configurations
- 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
- Installation Best Practices for Enterprise Manager High Availability
- Configuration With Grid Control

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 instance, 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 5, "Enterprise Manager Security" for information about securing the connections between Grid Control components
>
> Chapter 6, "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.

The Grid Control architecture consists of the following software components:

- Oracle Management Agent
- Oracle Management Service
- Oracle Management Repository
- Oracle Enterprise Manager 10*g* 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.

3.2 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

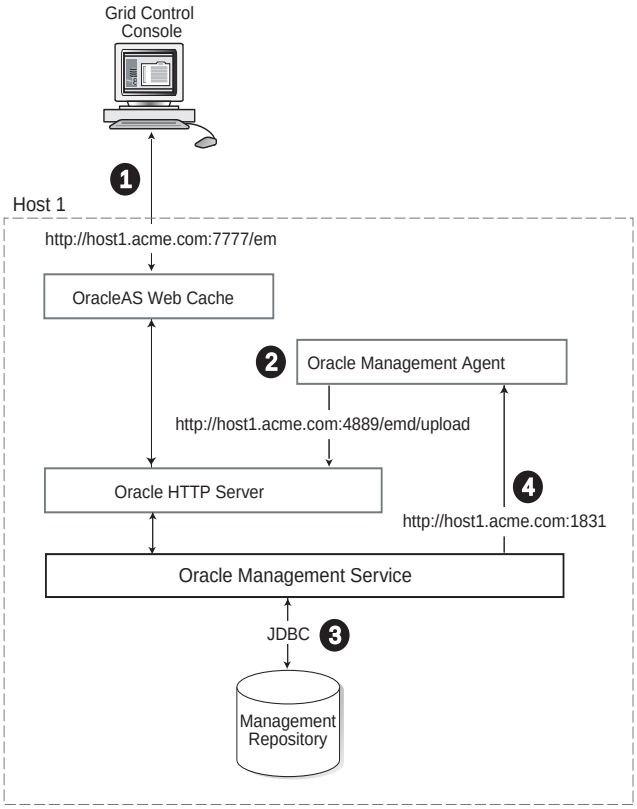

When you install all the Grid Control components on a single host, the management data travels along the following paths:

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 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 if 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 Management Repository database and to retrieve information from the Management Repository so it can be displayed in the Grid Control console. The Management Repository connection 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 12-8 for more information on modifying the Management 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.

3.3 Managing Multiple Hosts and Deploying a Remote Management Repository

Installing all the Grid Control components on a single host is an effective way to initially explore the capabilities and features available to you when you centrally manage your Oracle environment.

A logical progression from the single-host environment is to a more distributed approach, where the Management Repository database is on a separate host and does not compete for resources with the Management Service. The benefit in such a configuration is scalability; the workload for the Management Repository and Management Service can now be split. This configuration also provides the flexibility to adjust the resources allocated to each tier, depending on the system load. (Such a configuration is shown in Figure 3–2.) See Section 3.4.2.1, "Monitoring the Load on Your Management Service Installations" for additional information.

Figure 3–2 Grid Control Components Distributed on Multiple Hosts with One Management Service

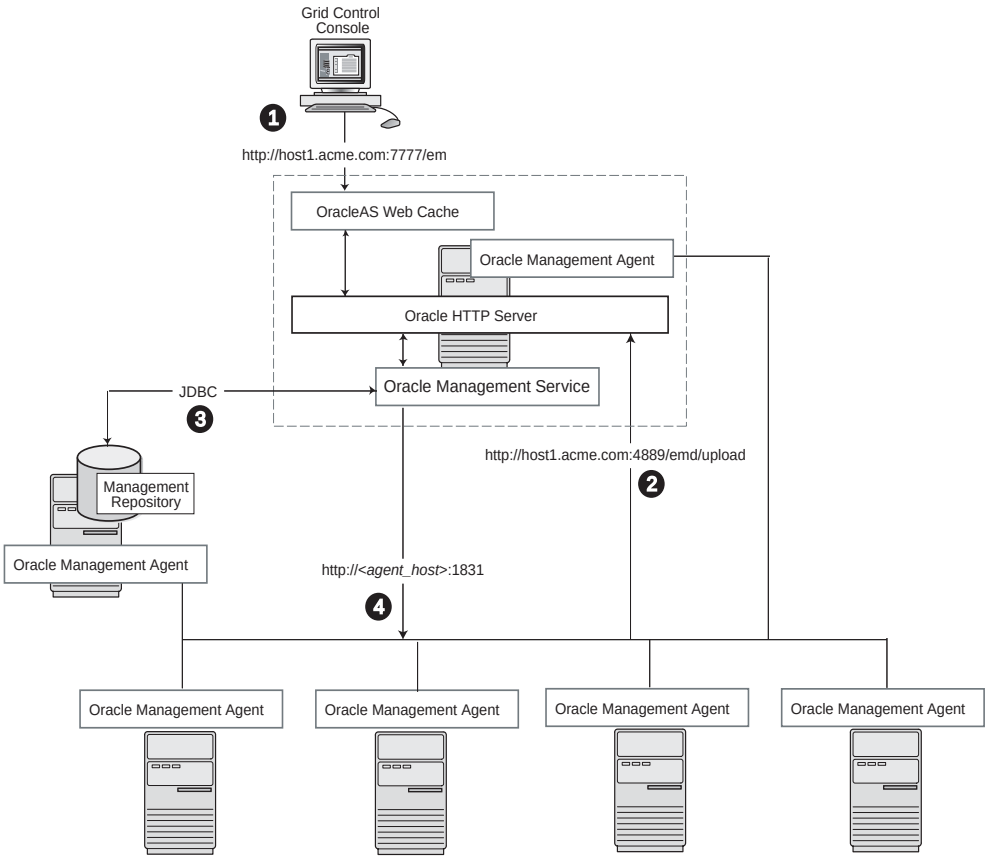

In this more distributed configuration, data about your managed targets travels along the following paths so it can be gathered, stored, and made available to administrators by way of the Grid Control console:

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 3.2.
2. Management Agents are installed on each host on the network, including the Management Repository host and the Management Service host. The Management Agents upload their data to the Management Service by way of the Management Service upload URL, which is defined in the `emd.properties` file in each Management Agent home directory. The upload URL 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 Management Repository database and to retrieve information from the Management Repository so it can be displayed in the Grid Control console. This remote connection is defined in the `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 3.2, the

Management Agent includes a built-in HTTP listener so no Oracle HTTP Server is required on the Management Agent host.

3.4 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:

- Understanding the Flow of Management Data When Using Multiple Management Services
- Determining When to Use Multiple Management Service Installations

3.4.1 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 scalability of the Grid Control environment.

Figure 3–3 Grid Control Architecture with Multiple Management Service Installations

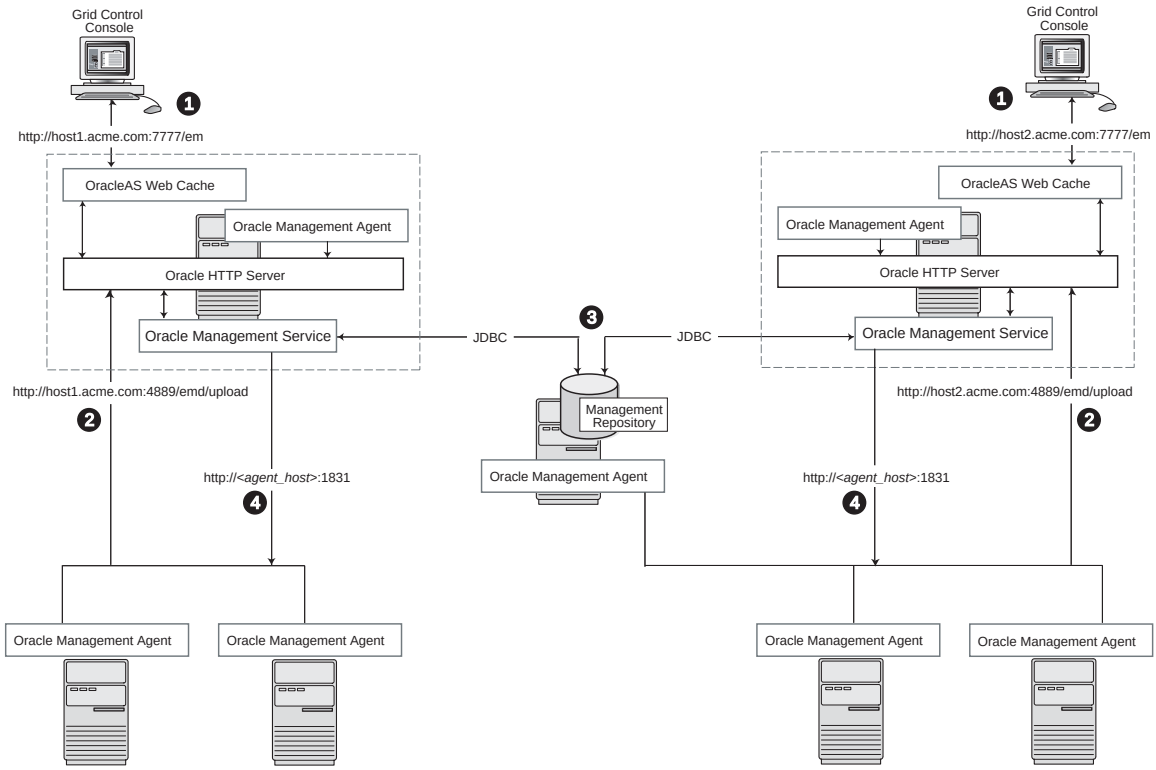

In a multiple Management Service configuration, the management data moves along the following paths:

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.

   Whenever more than one Management Service is installed, it is a best practice to have the Management Service upload to a shared directory. This allows all Management Service processes to manage files that have been downloaded from any Management Agent. Configure this functionality from the command line of each Management Service process as follows:

   ```
   emctl config oms loader -shared <yes|no> -dir <load
   directory>
   ```

Note: By adding a load balancer, you can avoid the following problems:

- Should the Management Service fail, any Management Agent connected to it cannot upload data.
- Because user accounts only know about one Management Service, users lose connectivity should the Management Service go down even if the other Management Service is up.

See Section 3.5, "High Availability Configurations" for information regarding load balancers.

---

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 3.2, 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.

3.4.2 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. 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

3.4.2.1 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 per 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 Setup on the Grid Control console.

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

- Loader backlog (files)

  The Loader is part of the Management Service that pushes metric data into the Management 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 totaled 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 consider adding another Management Service, reducing the number of notification rules, and verifying that all rules and notification methods are useful and valid.

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

The information on the Management Services and Repository page can help you determine the load being placed on your Management Service installations. More importantly, 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 7, "Configuring Services".

3.5 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 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 the Oracle white paper *Enterprise Manager Maximum Availability Architecture (MAA) Best Practices Enterprise Manager 10R2* available at:
>
> `http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_EnterpriseManagerBestPractices.pdf`

Refer to the following sections for more information about common Grid Control configurations that take advantage of high availability hardware and software solutions. These configurations are part of the Maximum Availability Architecture (MAA).

- Load Balancing Connections Between the Management Agent and the Management Service
- Load Balancing Connections Between the Grid Control Console and the Management Service
- Configuring a Load Balancer
- Configuring the Management Repository

3.5.1 Load Balancing Connections Between the Management Agent and the Management Service

Before you implement a plan to protect the flow of management data from the Management Agents to the Management Service, you should be aware of some key concepts.

Specifically, Management Agents do not maintain a persistent connection to the Management Service. When a Management Agent needs to upload collected monitoring data or an urgent target state change, the Management Agent establishes a connection to the Management Service. If the connection is not possible, such as in the case of a network failure or a host failure, the Management Agent retains the data and reattempts to send the information later.

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

However, if you decide to implement such a configuration, be sure to understand the flow of data when load balancing the upload of management data.

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

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

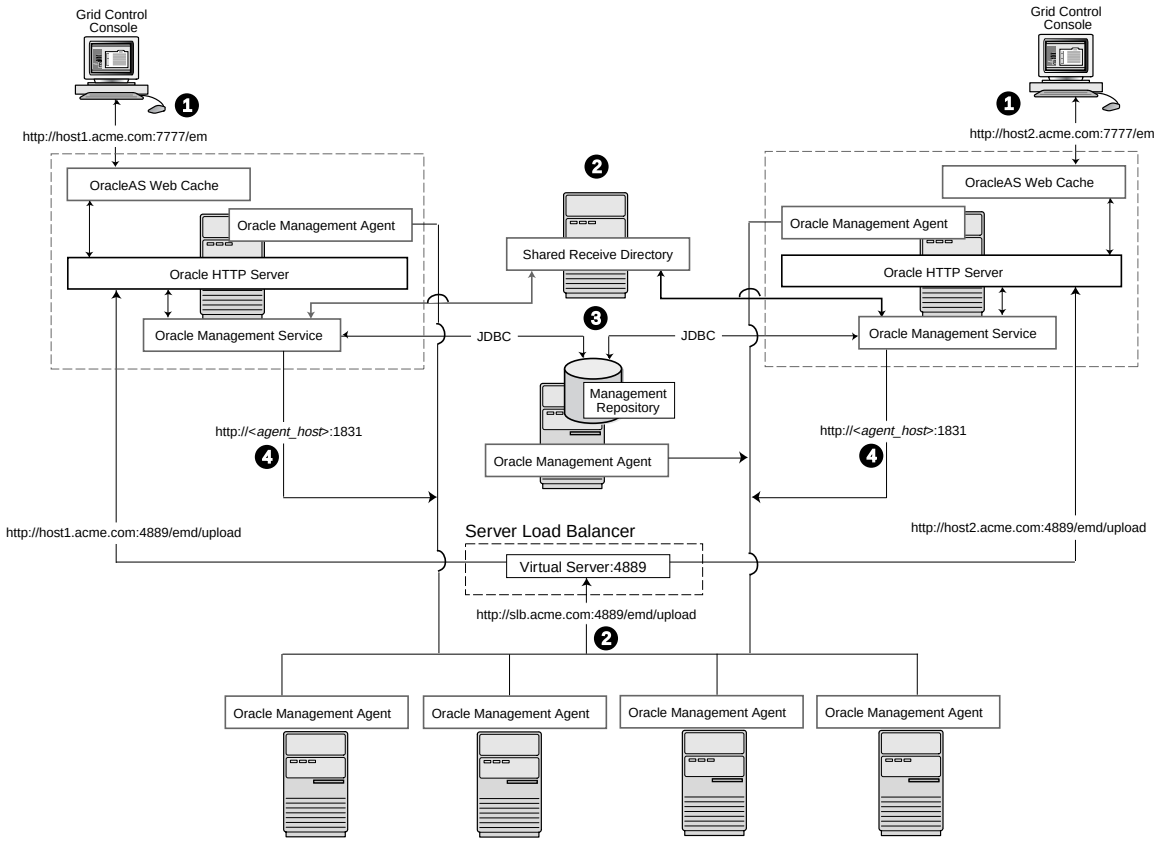

In this example, only the upload of Management Agent data is routed through the load balancer. The Grid Control console still connects directly to a single Management Service by way of the unique Management Service upload URL. This abstraction allows Grid Control to present a consistent URL to both Management

Agents and Grid Control consoles, regardless of the loss of any one Management Service component.

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

1. Administrators can use one of two URLs to access the Grid Control console just as they do in the multiple Management Service configuration.
2. Each Management Agent uploads its data to a common load balancer URL. This URL is defined in the `emd.properties` file for each Management Agent. In other words, the Management Agents connect to a virtual service exposed by the load balancer. The load balancer routes the request to any one of a number of available servers that provide the requested service.
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.

   Also, each Management Service communicates by way of JDBC with a common Management Repository, just as they do in the multiple Management Service configuration defined in Section 3.4.
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 3.4.

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

Using a 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 Oracle HTTP Server When Using a Load Balancer for the Grid Control Console
- Configuring the Management Services

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

Figure 3–5 shows a typical configuration where a 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

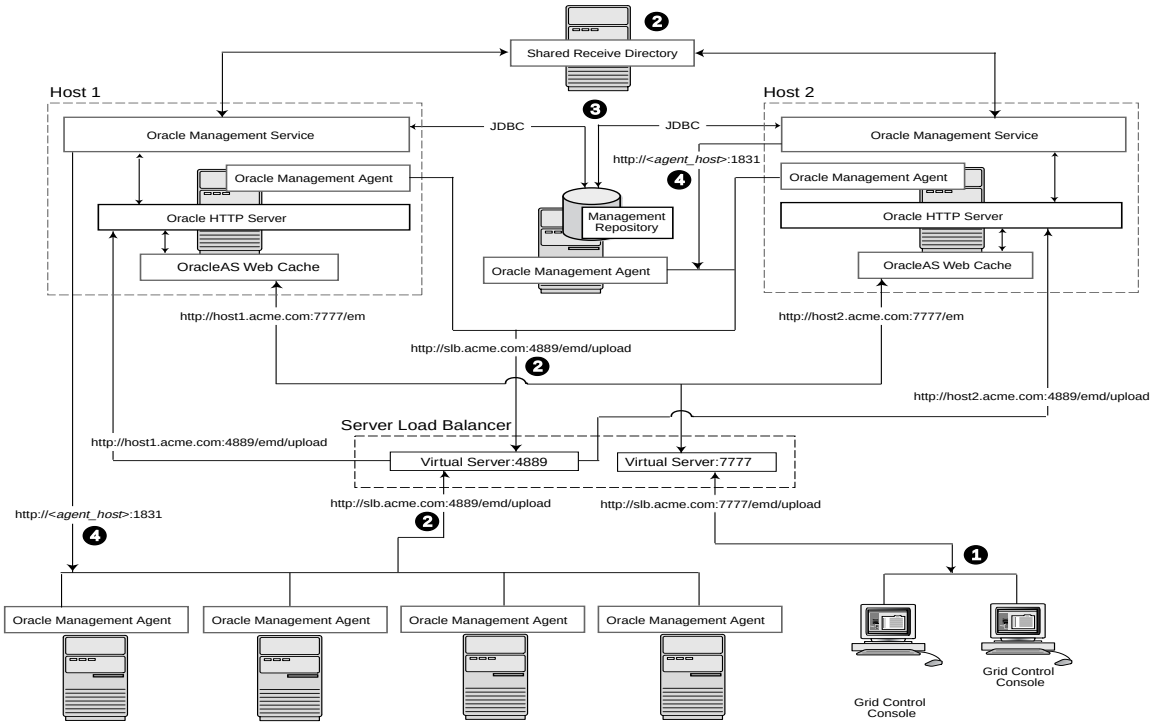

In this example, a single load balancer is used for the upload of data from the Management Agents and for the connections between the Grid Control console and the Management Service.

When you use a 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 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 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 load balancer URL (as described in Section 3.5.1) 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 3.4.
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.

3.5.2.2 Configuring Oracle HTTP Server When Using a 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 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 load balancer virtual service that you configured.

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

3.5.2.3 Configuring the Management Services

The Management Service for Grid Control 10*g* Release 2 has a high availability feature called the Shared Filesystem Loader. In the Shared Filesystem Loader, management data files received from Management Agents are stored temporarily on a common shared location called the shared receive directory. All Management Services are configured to use the same storage location for the shared receive directory. The Management Services coordinate internally and distribute 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.

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 must 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.

3.5.3 Configuring a Load Balancer

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

In the following examples, assume that you have installed two Management Service processes on Host A and Host B. For ease of installation, start with two hosts that have no application server processes installed. This ensures that the default ports are used as seen in the following table. The examples use these default values for illustration purposes.

Table 3–1 Management Service Ports

Name	Default Value	Description	Source	Defined By
Secure Upload Port	1159	Used for secure upload of management data from Management Agents.	httpd_em.conf and emoms.properties	Install. Can be modified by `emctl secure OMS - secure port <port>` command.
Agent Registration Port	4889	Used by Management Agents during the registration phase to download Management Agent wallets, for example, during `emctl secure agent`. In an unlocked Management Service, it can be used for uploading management data to the Management Service.	httpd_em.conf and emoms.properties	Install
Secure Console Port	4444	Used for secure (https) console access.	ssl.conf	Install
Unsecure Console Port	7777	Used for unsecure (http) console access.	httpd.conf	Install
Webcache Secure Port	4443	Used for secure (https) console access.	webcache.xml	Install
Webcache Unsecure Port	7779	Used for unsecure (http) console access.	webcache.xml	Install

By default, the service name on the Management Service-side certificate uses the name of the Management Service host. Management Agents do not accept this certificate when they communicate with the Management Service through a load balancer. You must run the following command to regenerate the certificate on both Management Services:

```
emctl secure -oms -sysman_pwd <sysman_pwd> -reg_pwd <agent_reg_
password> -host slb.acme.com -secure_port 1159
```

Specifically, you should use the administration tools that are packaged with your 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 Management Agent upload port. To insure a highly available Management Service, you should have two or more Management Services defined within the virtual pool. A sample configuration follows.

Sample Configuration

In this sample, both pools and virtual servers are created.

1. Create Pools

Pool abc_upload: Used for secure upload of management data from Management Agents to Management Services
Members: hostA:1159, hostB:1159
Persistence: None
Load Balancing: round robin
Pool abc_genWallet: Used for securing new Management Agents
Members: hostA:4889, host B:4889
Persistence: Active HTTP Cookie, method-> insert, expiration 60 minutes
Load balancing: round robin
Pool abc_uiAccess: Used for secure console access
Members: hostA:4444, hostB:4444
Persistence: SSL Session ID, timeout-> 3000 seconds (should be greater than the OC4J session timeout of 45 minutes)
Load balancing: round robin

2. Create Virtual Servers

   Virtual Server for secure upload
   Address: slb.acme.com
   Service: 1159
   Pool: abc_upload
   Virtual Server for Management Agent registration
   Address: slb.acme.com
   Service: 4889
   Pool: abc_genWallet
   Virtual Server for UI access
   Address: sslb.acme.com
   Service: https i.e. 443
   Pool: abc_uiAccess

3. 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 load balancer virtual service.

   See Also: "Configuring the Management Agent to Use a New Management Service" on page 12-1 for more information about modifying the REPOSITORY_URL property for a Management Agent

4. 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 distribution of 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.

Example of Configuring the Load Balancer

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 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`. See the following sample monitor configuration.

Sample Monitor Configuration

In this sample, three monitors are configured: mon_upload, mon_genWallet, and mon_uiAccess.

Monitor mon_upload
Type: https
Interval: 60
Timeout: 181
Send String: GET /em/upload HTTP 1.0\n
Receive Rule: Http Receiver Servlet active!
Associate with: hostA:1159, hostB:1159
Monitor mon_genWallet
Type: http
Interval: 60
Timeout: 181
Send String: GET /em/genwallet HTTP 1.0\n
Receive Rule: GenWallet Servlet activated
Associate with: hostA:4889, hostB:4889
Monitor mon_uiAccess
Type: https
Interval: 5
Timeout: 16
Send String: GET /em/console/home HTTP/1.0\nUser-Agent: Mozilla/4.0 (compatible; MSIE 6.0, Windows NT 5.0)\n
Receive Rule: /em/console/logon/logon;jsessionid=
Associate with: hostA:4444, hostB:4444

Note: The network bandwidth requirements on the 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 Load Balancer documentation for more information on configuring virtual pools, load balancing policies, and monitoring network traffic

3.5.4 Configuring the Management Repository

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. One way is to specify the size of the Management Repository tablespaces in a Real Application Cluster (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 three options for increasing the size of the Management Repository.

- The preferred option for high availability installs is to use Oracle Automatic Storage Management to manage database storage. The location string must 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.
- 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.

3.6 Installation Best Practices for Enterprise Manager High Availability

The following sections document best practices for installation and configuration of each Grid Control component.

3.6.1 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.

3.6.2 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.

3.6.3 Installing the Management Agent Software on Redundant Storage

The Management Agent persists its intermediate state and collected information using local files in the `$AGENT_HOME/$HOSTNAME/sysman/emd` sub tree under the Management Agent home directory.

In the event that these files are lost or corrupted before being uploaded to the Management Repository, a loss of monitoring data and any pending alerts not yet uploaded to the Management Repository occurs.

At a minimum, configure these sub-directories on striped redundant or mirrored storage. Availability would be further enhanced by placing the entire $AGENT_HOME on redundant storage. The Management Agent home directory is shown by entering the command '`emctl getemhome`' on the command line, or from the Management Services and Repository tab and Agents tab in the Grid Control console.

3.6.4 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.

3.7 Configuration With Grid Control

Grid Control comes preconfigured with a series of default rules to monitor many common targets. These rules can be extended to monitor the Grid Control infrastructure as well as the other targets on your network to meet specific monitoring needs.

3.7.1 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.

3.7.2 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.

3.7.3 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 11.3, "Oracle Enterprise Manager Backup, Recovery, and Disaster Recovery Considerations" on page 11-15.

3.7.4 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.

3.7.4.1 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.

3.7.4.2 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.

4

Configuring Oracle Enterprise Manager for Active and Passive Environments

Active and Passive environments, also known as Cold Failover Cluster (CFC) environments, refer to one type of high availability solution that allows an application to run on one node at a time. These environments generally use a combination of *cluster* software to provide a logical host name and IP address, along with interconnected host and storage systems to share information to provide a measure of high availability for applications.

This chapter contains the following sections:

- Configuring Oracle Enterprise Management Agents for Use in Active and Passive Environments
- Using Virtual Host Names for Active and Passive High Availability Environments in Enterprise Manager Database Control

4.1 Configuring Oracle Enterprise Management Agents for Use in Active and Passive Environments

In a Cold Failover Cluster environment, one host is considered the *active node* where applications are run, accessing the data contained on the shared storage. The second node is considered the *standby node*, ready to run the same applications currently hosted on the primary node in the event of a failure. The cluster software is configured to present a *Logical Host Name* and IP address. This address provides a generic location for running applications that is not tied to either the active node or the standby node.

In the event of a failure of the active node, applications can be terminated either by the hardware failure or by the cluster software. These application can then be restarted on the passive node using the same logical host name and IP address to access the new node; resuming operations with little disruption. Automating failover of the virtual host name and IP, along with starting the applications on the passive node, requires the use of the third party cluster software. Several Oracle partner vendors provide high availability solutions in this area.

4.1.1 Installation and Configuration

Enterprise Manager can be configured to support this Cold Failover Cluster configuration using the existing Management Agents communicating to the Oracle Management Service processes.

If your application is running in an Active and Passive environment, the clusterware does the job of bringing up the *passive* or *standby* database instance in case the *Active* database goes down. For Enterprise Manager to continue monitoring the application

instance in such a scenario, the existing Management Agents need additional configuration.

The additional configuration steps for this environment involve:

- Installing of an extra Management Agent using the logical host name and IP address generated through the cluster software.
- Modifying the targets monitored by each Management Agent once the third Management Agent is installed.

In summary, this configuration results in the installation of three Management Agents, one for each hardware node and one for the IP address generated by the cluster software. Theoretically, if the cluster software supports the generation of multiple virtual IP addresses to support multiple high availability environments, the solution outlined here should scale to support the environment.

The following table documents the steps required to configure agents in a CFC environment:

Table 4–1 Steps Required to Configure Management Agents in a Cold Failover Cluster Environment

Action	Method	Description/Outcome	Verification
Install the vendor specific cluster software	Installation method varies depending on the cluster vendor.	The minimal requirement is a 2-node cluster that supports Virtual or Floating IP addresses and shared storage.	Use the `ping` command to verify the existence of the floating IP address. Use `nslookup` or equivalent command to verify the IP address in your environment.
Install Management Agents to each physical node of the cluster using the physical IP address or host name as the node name.	Use the Oracle Universal Installer (OUI) to install Management Agents to each node of the cluster.	When complete, the OUI will have installed Management Agents on each node that will be visible through the Grid Control console.	Check that the Management Agent, host, and targets are visible in the Enterprise Manager environment.
Delete targets that will be configured for high availability using the cluster software.	Using the Grid Control console, delete all targets discovered during the previous installation step that are managed by the cluster software except for the Management Agent and the host.	Grid Control Console displays the Management Agent, hardware, and any target that is not configured for high availability.	Inspect the Grid Control console and verify that all targets that will be assigned to the Management Agent running on the floating IP address have been deleted from the Management Agents monitoring the fixed IP addresses.
Install a third Management Agent to the cluster using the logical IP address or logical host name as the host specified in the OUI at install time. **Note:** This installation should not detect or install to more than one node.	This Management Agent must follow all the same conventions as any application using the cluster software to move between nodes (that is, installed on the shared storage using the logical IP address). This installation requires an additional option to be used at the command line during installation time. The 'HOSTNAME' flag must be set as in the following example: (/144)- >./runInstaller HOSTNAME=<Logical IP address or hostname>	Third Management Agent installed, currently monitoring all targets discovered on the host running physical IP.	To verify the Management Agent is configured correctly, type `emctl status agent` at the command line and verify the use of the logical IP virtual host name. Also, verify that the Management Agent is set to the correct Management Service URL and that the Management Agent is uploading the files. When the Management Agent is running and uploading data, use the Grid Control console to verify that it has correctly discovered targets that will move to the standby node during a failover operation.

Table 4–1 (Cont.) Steps Required to Configure Management Agents in a Cold Failover Cluster

Action	Method	Description/Outcome	Verification
Delete any targets from the Management Agent monitoring the logical IP that will not switch to the passive node during failover.	Use the Grid Control console to delete any targets that will not move between hosts in a switchover or failover scenario. These might be targets that are not attached to this logical IP address for failover or are not configured for redundancy.	Grid Control console is now running three Management Agents. Any target that is configured for switchover using cluster software will be monitored by a Management Agent that will transition during switchover or failover operations.	The operation is also verified by inspecting the Grid Control console. All targets that will move between nodes should be monitored by the Management Agent running on the virtual host name. All remaining targets should be monitored by a Management Agent running on an individual node.
Add the new logical host to the cluster definition.	Using the All Targets tab on the Grid Control console, find the cluster target and add the newly discovered logical host to the existing cluster target definition.	It is also possible (*not required*) to use the **Add Cluster Target** option on the All Targets tab, making a new composite target using the nodes of the cluster.	The Grid Control console will now correctly display all the hosts associated with the cluster.
Place the Management Agent process running on the logical IP under the control of the cluster software.	This will vary based on the cluster software vendor.	Management Agent will transition along with applications. A suggested order of operation is covered in the next section.	Verify that the Management Agent can be stopped and restarted on the standby node using the cluster software.

4.1.2 Switchover Steps

Each cluster vendor will implement the process of building a wrapper around the steps required to do a switchover or failover in a different fashion. The steps themselves are generic and are listed here:

- Shut down the Management Agent
- Shut down all the applications running on the virtual IP and shared storage
- Switch the IP and shared storage to the new node
- Restart the applications
- Restart the Management Agent

Stopping the Management Agent first, and restarting it after the other applications have started, prevents Enterprise Manager from triggering any false *target down* alerts that would otherwise occur during a switchover or failover.

4.1.3 Performance Implications

While it is logical to assume that running two Management Agent processes on the active host may have some performance implications, this was not shown during testing. Keep in mind that if the Management Agents are configured as described in this chapter, the Management Agent monitoring the physical host IP will only have two targets to monitor. Therefore the only additional overhead is the two Management Agent processes themselves and the commands they issue to monitor a Management Agent and the operating system. During testing, it was noticed that an overhead of between 1-2% of CPU usage occurred.

4.1.4 Summary

Generically, configuring Enterprise Manager to support Cold Cluster Failover environments encompasses the following steps.

- Install a Management Agent for each virtual host name that is presented by the cluster and insure that the Management Agent is correctly communicating to the Management Service.
- Configure the Management Agent that will move between nodes to monitor the appropriate highly available targets.
- Verify that the Management Agent can be stopped on the primary node and restarted on the secondary node automatically by the cluster software in the event of a switchover or failover.

4.2 Using Virtual Host Names for Active and Passive High Availability Environments in Enterprise Manager Database Control

This section provides information to database administrators about configuring an Oracle Database release 10*g* in Cold Failover Cluster environments using Enterprise Manager Database Control.

The following conditions must be met for Database Control to service a database instance after failing over to a different host in the cluster:

- The installation of the database must be done using a Virtual IP address.
- The installation must be conducted on a shared disk or volume which holds the binaries, configuration, and runtime data (including the database).
- Configuration data and metadata must also failover to the surviving node.
- Inventory location must failover to the surviving node.
- Software owner and time zone parameters must be the same on all cluster member nodes that will host this database.

The following items are configuration and installation points you should consider before getting started.

- To override the physical host name of the cluster member with a virtual host name, software must be installed using the parameter ORACLE_HOSTNAME.
- For inventory pointer, software must be installed using the command parameter `invPtrLoc` to point to the shared inventory location file, which includes the path to the shared inventory location.
- The database software, the configuration of the database, and Database Control are done on a shared volume.

4.2.1 Set Up the Alias for the Virtual Host Name and Virtual IP Address

You can set up the alias for the virtual host name and virtual IP address by either allowing the clusterware to set it up automatically or by setting it up manually before installation and startup of Oracle services. The virtual host name must be static and resolvable consistently on the network. All nodes participating in the setup must resolve the virtual IP address to the same host name. Standard TCP tools similar to `nslookup` and `traceroute` commands can be used to verify the set up.

4.2.2 Set Up Shared Storage

Shared storage can be managed by the clusterware that is in use or you can use any shared file system volume as long as it is not unsupported. The most common shared file system is NFS.

4.2.3 Set Up the Environment

Some operating system versions require specific operating system patches to be applied prior to installing release 10*g*R2 of the Oracle database. You must also have sufficient kernel resources available when you conduct the installation.

Before you launch the installer, specific environment variables must be verified. Each of the following variables must be identically set for the account you are using to install the software on all machines participating in the cluster.

- Operating system variable TZ, time zone setting. You should unset this prior to the installation.
- PERL variables. Variables like PERL5LIB should be unset to prevent the installation and Database Control from picking up the incorrect set of PERL libraries.
- Paths used for dynamic libraries. Based on the operating system, the variables can be LD_LIBRARY_PATH, LIBPATH, SHLIB_PATH, or DYLD_LIBRARY_PATH. These variables should *only* point to directories that are visible and usable on each node of the cluster.

4.2.4 Ensure That the Oracle USERNAME, ID, and GROUP NAME Are Synchronized on All Cluster Members

The user and group of the software owner should be defined identically on all nodes of the cluster. You can verify this using the following command:

```
$ id -a
uid=1234(oracle) gid=5678(dba)groups=5678(dba)
```

4.2.5 Ensure That Inventory Files Are on the Shared Storage

To ensure that inventory files are on the shared storage, follow these steps:

- Create you new ORACLE_HOME directory.
- Create the Oracle Inventory directory under the new Oracle home

  ```
  cd <shared oracle home>
  mkdir oraInventory
  ```

- Create the oraInst.loc file. This file contains the Inventory directory path information required by the Universal Installer:

  1. `vi oraInst.loc`
  2. Enter the path information to the Oracle Inventory directory and specify the group of the software owner as the `dba` user. For example:

     ```
     inventory_loc=/app/oracle/product/10.2/oraInventory
     inst_group=dba
     ```

4.2.6 Start the Installer

To start the installer, point to the inventory location file oraInst.loc, and specify the host name of the virtual group. The debug parameter in the example below is optional:

```
run installer -invPtrloc /app/oracle/share1/oraInst.loc ORACLE_
HOSTNAME=lxdb.acme.com -debug
```

4.2.7 Start Services

You must start the services in the following order:

1. Establish IP address on the active node
2. Start the TNS listener
3. Start the database
4. Start dbconsole
5. Test functionality

In the event that services do not start, do the following:

1. Establish IP on failover box
2. Start TNS listener

   ```
   lsnrctl start
   ```

3. Start the database

   ```
   dbstart
   ```

4. Start Database Control

   ```
   emctl start dbconsole
   ```

5. Test functionality

To manually stop or shutdown a service, follow these steps:

1. Stop the application.
2. Stop Database Control

   ```
   emctl stop dbconsole
   ```

3. Stop TNS listener

   ```
   lsnrctl stop
   ```

4. Stop the database

   ```
   dbshut
   ```

5. Stop IP

5

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

5.1 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.

5.1.1 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 10*g* Grid Control Console and the Oracle Enterprise Manager 10*g* 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 5-4

5.1.2 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 10*g*. 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.

5.1.3 Resources Protected

By restricting access to privileged users and providing tools to secure communications between Oracle Enterprise Manager 10*g* 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.

5.1.4 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

5.1.5 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 10*g* Grid Control Console and Application Server Control Console.

When you configure security for the Oracle Enterprise Manager 10*g* 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 5-4 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.

5.1.6 Leveraging Oracle Identity Management Infrastructure

Oracle Enterprise Manager 10*g* 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 5-19

- 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 5-23

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.

5.2 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 with a Server Load Balancer
- Enabling Security for the Management Repository Database

5.2.1 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 5–1 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 5–1 Enterprise Manager Framework Security

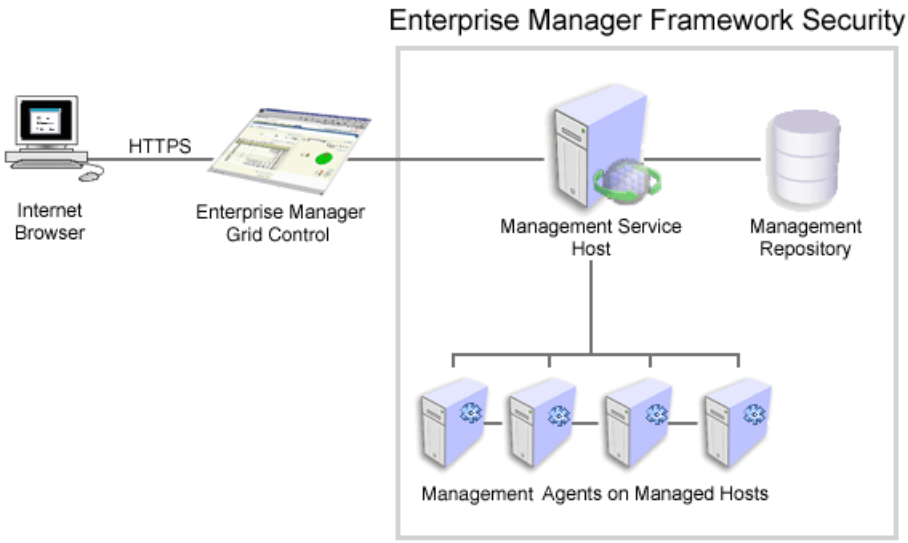

Enterprise Manager Framework Security implements the following types of secure connections between the Enterprise Manager components:

- 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*

5.2.2 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.

5.2.3 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 6-12

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 5–1 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 5–2 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>.

5.2.3.1 Checking the Security Status

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

Example 5–3 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
```

5.2.4 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 5–4 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 5–2), 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 5–4 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 5–5 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 5–2 Secure Upload Field on the Management Agent Home Page

5.2.5 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 10*g* Grid Control Console.

See Also: "Managing Agent Registration Passwords" on page 5-13

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.

5.2.6 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 Recceiver 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 10*g* 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 5–6 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 5–7 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 10*g* 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*

5.2.7 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.

5.2.7.1 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 5–3). 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 5–3 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.

5.2.7.2 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.

5.2.8 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 5–1). Enter the host name for the Server Load Balancer.

> **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 oms agent_reg_pwd  -hostname mgmt_host
> ```

5.2.9 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

5.2.9.1 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 5-18

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*.

5.2.9.2 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 5–1 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 5–1 Oracle Advanced Security Properties in the Enterprise Manager Properties File

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

5.2.9.3 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 10*g* 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.

5.2.9.4 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

5.3 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 10*g* 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

5.3.1 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 sshost1.acme.com -port 1521 -sid asdb
-pass Ch22x5xt -das http://ssohost1.acme.com:7777
```

Table 5–2 describes the arguments on the `emctl config oms sso` command line.

Example 5–8 shows the typical output generated by the `emctl config oms sso` command.

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

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

7. 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 5–2 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 *Oracle Application Server Single Sign-On Administrator's Guide*.
-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 5–8 Sample Output of the emctl config oms sso Command

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

Oracle Enterprise Manager 10g Release 10.2.0.0.0  Copyright (c) 1996,
2005 Oracle Corporation.  All rights reserved.
/scratch/smptest/mm9/oms10g/Apache/Apache/conf/httpd.conf has been modified.
/scratch/smptest/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/smptest/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: param25:oracle
param26:-sso_server_versi
on param27:v1.2  -DinstallType=
-DoldOracleHome=
-DoldOHSUser=root
Check /scratch/smptest/mm9/oms10g/sso/log/ssoreg.log for details of this
registration
SSO registration tool finished successfully.
Done!
```

5.3.2 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 5–4).

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 5–5).

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 5–4 Modified Enterprise Manager Logon Page When Configuring SSO

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

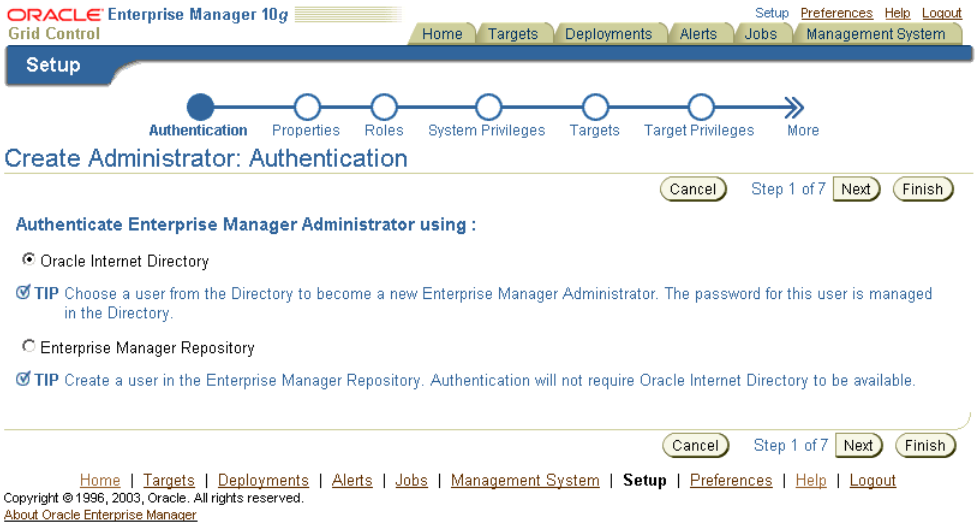

5.3.3 Grid Control as a Single Sign-On Partner Application

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

5.3.4 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
```

5.4 Configuring Enterprise Manager for Use with Enterprise User Security

Enterprise User Security enables you to create and store Oracle9*i* 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 Oracle9*i* 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 Oracle9*i* 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 10*g* 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 10*g* 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.

5.5 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.

5.5.1 Audit Data

The following data is audited for all Enterprise Manager operations:

Table 5–3 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: ▪ 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 `grant_target_priv` operation is to be audited, apart from the common data, the `user_name, target_name, target_type`, and `target_owner` 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 `create_user` operation is associated with `user_name` object, the `submit_job` operation has a job name associated with it.
Object Type	Each operation code has an object type associated with it. For example, the `create_user` operation has the `user_type` object associated with it, the `submit_job` 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.

Table 5–3 (Cont.) Common Audit Data

Field Name	Description
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.

5.5.2 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 5–4 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	

5.5.3 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_LEVL 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.

5.5.4 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 5-25.
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');`

5.6 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 5–9 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.
```

5.6.1 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
```

5.6.2 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>]
```

5.6.2.1 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 5–10 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 5–11 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 5–12 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 5–13 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>".
```

5.6.2.2 emctl config emkey -repos

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

Example 5–14 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.

5.6.2.3 emctl config emkey -emkeyfile

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

Example 5–15 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.

5.6.2.4 emctl config emkey -emkey

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

Example 5–16 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.

5.6.2.5 emctl config emkey -remove_from_repos

This command removes the emkey from the Management Repository.

Example 5–17 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.
```

5.6.2.6 emctl config emkey -copy_to_repos

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

Example 5–18 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 5.6.3). 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

5.6.3 Install and Upgrade Scenarios

This section explains the install and upgrade scenarios for emkey.

5.6.3.1 Installing the Management Repository

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

5.6.3.2 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

5.6.3.3 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.

5.6.3.4 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
```

5.6.3.5 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.

5.7 Additional Security Considerations

After you enable security for the Enterprise Manager components and framework, ther'e 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

5.7.1 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

5.7.1.1 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 5–6 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 5–6 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 5–7.

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 5–7 Certificate Path Tab on the Internet Explorer Certificate Dialog Box

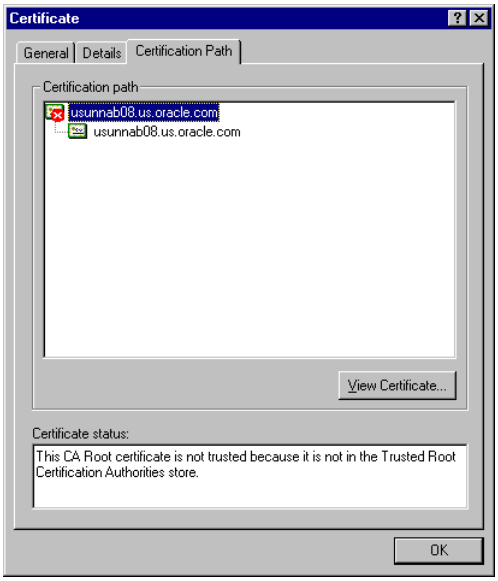

5.7.1.2 Responding to the Netscape Navigator New Site Certificate Dialog Box

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 5–8 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 5–8 Netscape Navigator New Site Certificate Dialog Box

5.7.1.3 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 5–9 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 5–9 Internet Explorer Security Information Dialog Box

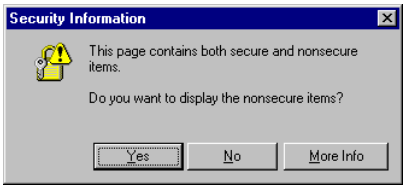

To stop this dialog box from displaying:

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.

5.7.2 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 5–7.

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 5–19.

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 5–19 Sample Content of an Exported Certificate

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

5.8 Other Security Features

This section describes Enterprise Manager security features.

5.8.1 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 "Y". "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 5–5 describes the input values to the emcli set_credential command.

Table 5–5 emcli set_credential Parameters

Parameter	Input Value	Description
-target_type	-target_type ="ttype"	Type of target. Must be "host" in case the "-oracle_homes" parameter is specified.
-target_name	[-target_name="tname"]	Name of target. Omit this argument to set enterprise preferred credentials. Must be hostname in case "-oracle_ homes" parameter is specified
-credential_set	-credential_set="cred_set"	Credential set affected.
-user	[-user="user"]	Enterprise Manager user whose credentials are affected. If omitted, the current user's credentials are affected.
-columns	-columns="col1:newval1;col2:n ewval2;..."	The name and new value of the column(s) to set. Every column of the credential set must be specified. Alternatively, a tag from the -input_ file argument may be used so that the credential values are not seen on the command line. This argument may be specified more than once.
-input_file	[-input_file="tag1:file_ path1;tag2:file_path2;..."]	Path of file that has -columns argument(s). This option is used to hide passwords. Each path must be accompanied by a tag which is referenced in the -columns argument. This argument may be specified more than once.
-oracle_homes	[-oracle_ homes="home1;home2"]	Name of Oracle Homes on the target host. Credentials will be added/updated for all specified home

5.8.2 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.

5.8.3 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.

5.8.4 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.

6

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

6.1 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 10*g* 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 12, "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.

6.2 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

Figure 6–1 provides a topology of an Enterprise Manager grid environment that is using a firewall, and also illustrates the default ports that can be used.

Figure 6–1 Firewall Port Requirements (Default)

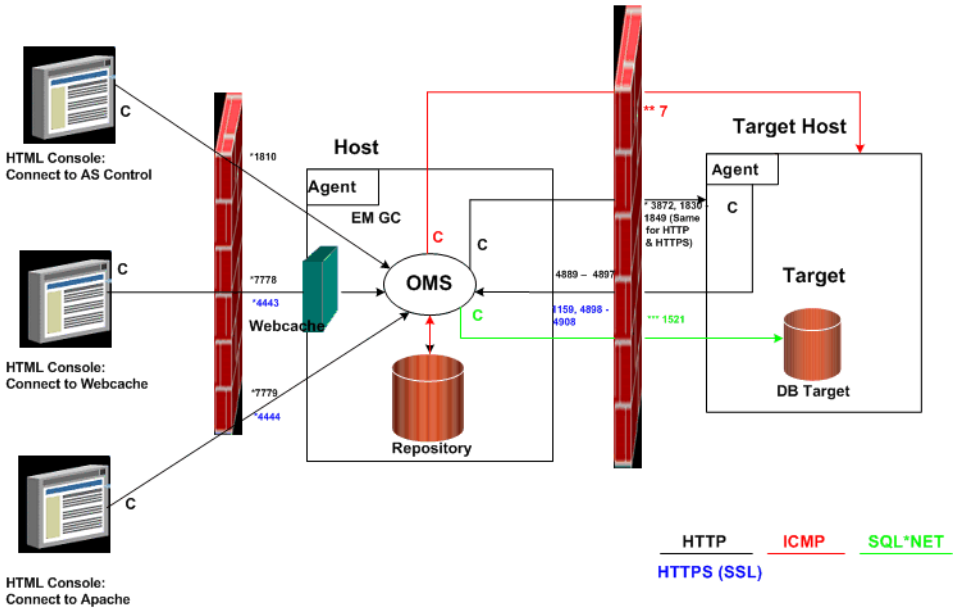

The conventions used in the preceding illustration are as follows:

Table 6–1 Conventions Used

Convention	Description
C	Is the entity that is making the call.
*	Enterprise Manager will default to the first available port within an Enterprise Manager set range.
**	Enterprise Manager will default to the first available port.
***	Are the Database listener ports.

Note:

- The direction of the arrows specify the direction of ports.
- Port 1159, 4898-4989 specify that 1159 is the default. If this port is not available, the Management Service will search in the range that is specified.
- To clone between two target hosts separated by a firewall, the agents will need to communicate to each other on the agent ports. The initiating agent will make the call.

6.2.1 Firewalls Between Your Browser and the Grid Control Console

Connections from your browser to the Oracle Enterprise Manager 10*g* 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 7778. 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 7778:

```
http://mgmthost.acme.com:7778/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 6–2 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 6–2 Firewall Between Your Browser and the Grid Control Console

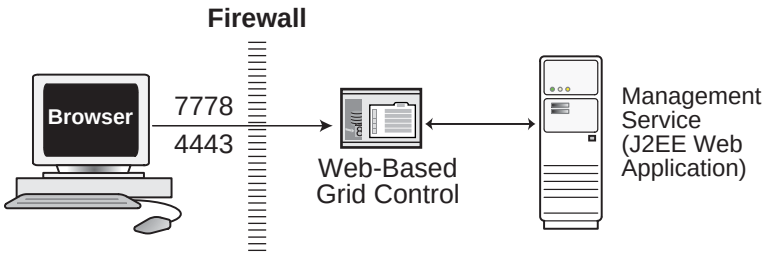

6.2.2 Configuring the Management Agent on a Host Protected by a Firewall

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 on the Management Agent port. Regardless of whether or not Enterprise Manager Framework Security has been enabled, the default port is 3872. If this default port is not available, the default port range between 1830 - 1849 is used. Incoming traffic can be received only if the port corresponding to the Management Agent is open in the firewall.

Figure 6–3 illustrates the connections the Management Agent must make when it is protected by a firewall.

Figure 6–3 Configuration Tasks When the Management Agent is Behind a Firewall

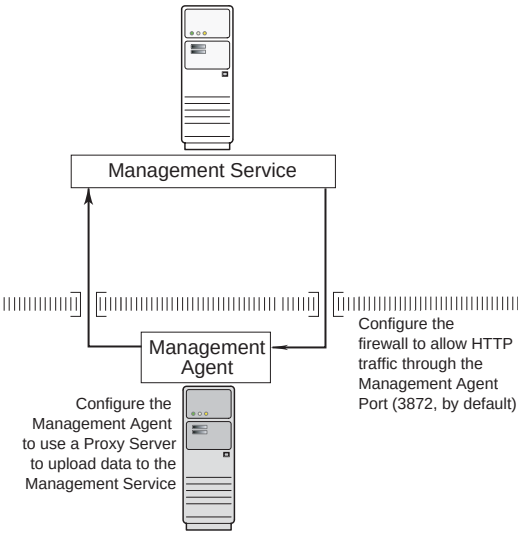

6.2.2.1 Configuring the Management Agent to Use a Proxy Server

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.

6.2.2.2 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 5-4

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

> **See Also:** "Chapter 12, "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 5-4 for information about Enterprise Manager Framework Security

6.2.3 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 1159 by default. If this port is not available, Enterprise Manager will default to first available port in the range 4898-4989. If you have *not* enabled Enterprise Manager Framework Security, the upload port is the first available port in the range 4889 - 4897.

  See also: "Enabling Security for the Oracle Management Service" on page 5-6

Figure 6–4 illustrates the connections the Management Service must make when it is protected by a firewall.

Figure 6–4 Configuration Tasks When the Management Service is Behind a Firewall

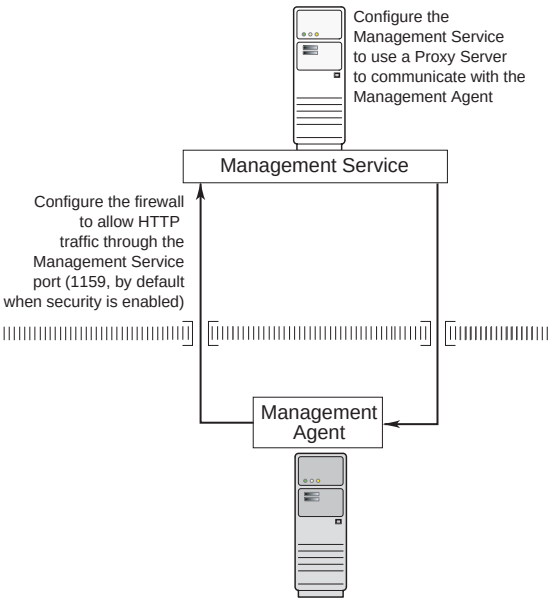

6.2.3.1 Configuring the Management Service to Use a Proxy Server

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 6-8 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.

6.2.3.2 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 Oracle*MetaLink* 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 Oracle*MetaLink* 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 Oracle*MetaLink* 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
```

6.2.3.3 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 1159 HTTPS port.

> **See Also:** "Configuring Security for Grid Control" on page 5-4 for information about Enterprise Manager Framework Security

Administrators can also change the upload port after the installation.

> **See Also:** Chapter 12, "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

6.2.4 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 6–5 shows a typical configuration of a firewall between the Management Service and the Management Repository.

Figure 6–5 Firewall Between the Management Service and the Management Repository

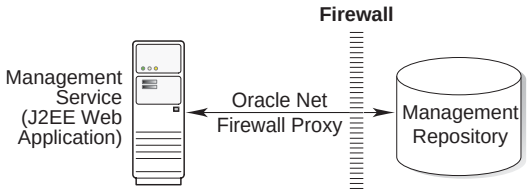

6.2.5 Firewalls Between the Grid Control and a Managed Database Target

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 6–6 shows a typical configuration of a firewall between Grid Control and the Management Repository.

Figure 6–6 Firewall Between Grid Control and Managed Database Target

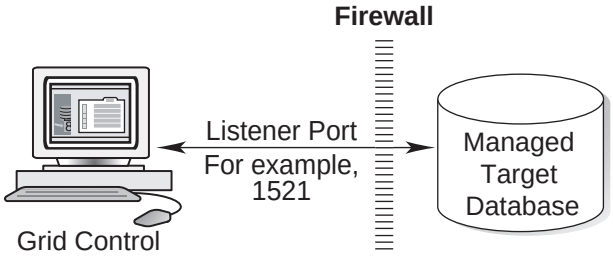

6.2.6 Firewalls Used with Multiple Management Services

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 6-4

"Configuring the Management Service on a Host Protected by a Firewall" on page 6-6

- 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

6.2.7 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 5-35

6.2.8 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

6.3 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 10*g* components before you configure your firewalls.

When you install the Oracle Application Server 10*g* or the Oracle Enterprise Manager 10*g* 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

6.4 Additional Considerations for Windows XP

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.

7

Configuring Services

This chapter describes how to configure services in Oracle Enterprise Manager 10*g* Grid Control Console. It contains the following sections:

- Summary of Service Management Tasks
- Setting up the System
- Creating a Service
- Configuring a Service
- Recording Web Transactions
- Monitoring Settings
- Configuring Aggregate Services
- Configuring End-User Performance Monitoring
- Managing Forms Applications
- Configuring OC4J for Request Performance Diagnostics
- Setting Up Monitoring Templates
- Configuring Service Levels
- Configuring a Service Using the Command Line Interface

7.1 Summary of Service Management Tasks

This table provides a summary list of all the service management features and their requirements.

Table 7–1 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.	▪ 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: ▪ Understand real end-user page response times within your application. ▪ Assess the user impact of performance problems. ▪ Analyze end user response times by by page, domain, region, visitors, and Web server.	▪ Oracle HTTP Server Based on Apache 2.0 or Apache HTTP Server 2.0 ▪ Oracle Application Server Web Cache (10.1.2, 9.0.4, 9.0.3, or 9.0.2)	Configuring End-User Performance Monitoring
Interactive Transaction Tracing	Enterprise Manager provides a mechanism for interactively tracing Web transactions. This feature allows you to: ▪ 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.	▪ Microsoft Internet Explorer 5.5 or later for creating and playing back transactions. ▪ Oracle Application Server 10*g* (9.0.4) for playing back a transaction with trace to view J2EE server activity times. **Note:** Recording a transaction is an optional feature. You can manually create a transaction by entering the required values.	Configuring Interactive Transaction Tracing

Table 7–1 (Cont.) Summary of Service Management Tasks

Feature	Description	Requirements	Refer to
Request Performance	Enterprise Manager can gather critical request performance data about your Web application. The Request Performance feature allows you to: ▪ 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 10*g* (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 ▪ Microsoft Internet Explorer 5.5 or higher ▪ Adobe SVG Viewer 3.0	Root Cause Analysis Configuration
Forms Applications	A Forms Application target in Enterprise Manager can be used to model and monitor a specific Forms application. You can: ▪ Record and monitor a Forms transaction. ▪ Measure the End-User Performance of Forms actions such as Commit, Query, Runform, Callform, Openform, and Newform.	▪ Microsoft Internet Explorer 5.5 or later for creating and playing back Forms transactions. ▪ Oracle HTTP Server or Apache HTTP Server ▪ Oracle Application Server Web Cache (10.1.2, or 9.0.4)	Recording and Monitoring Forms Transactions Monitoring the End-User Performance of Forms Applications

7.2 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.

7.3 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 7–1 Create Service - General Page

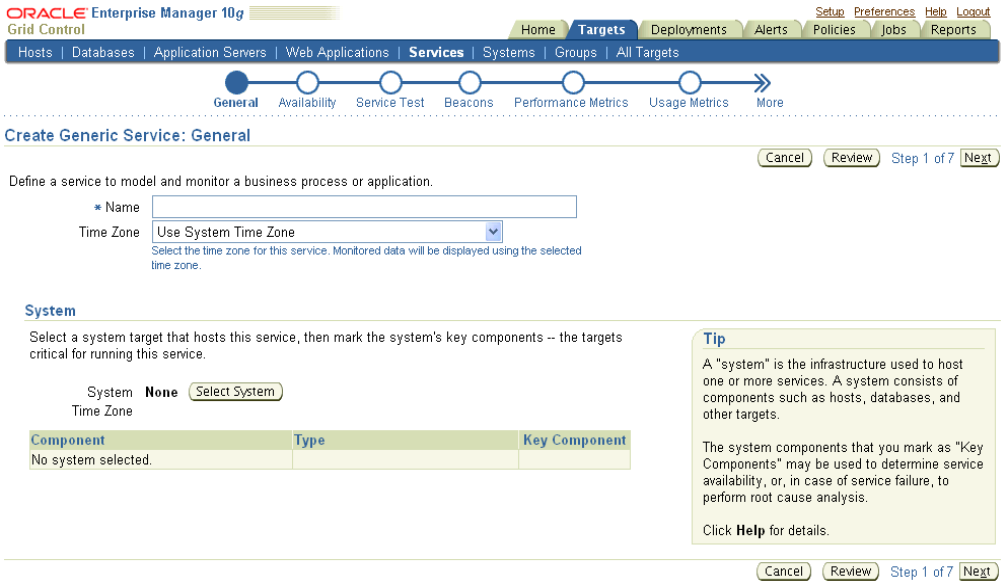

Follow the steps in the wizard to create your service. This involves the following:

- Identifying the type of service to be created. You can define different types of services based on your requirement. Some of the services that you can define are Generic Service, Web Application, Aggregate Service, and Forms Application. A Generic Service is used to monitor a variety of different protocol based services. A Web Application is used to monitor Web transactions. Enterprise Manager provides additional monitoring and diagnostics features for Web applications. A Forms Application is used to monitor Forms transactions. Each Forms transaction can consist of one or more actions that can be monitored. You can also define other services that are specific to an application such as the OCS Service. You can combine one or more services to form an Aggregate Service.
- Specifying the name and time zone for the service.
- Selecting a system target that hosts this service and then marking the system's key components that are critical for running the service. These key components are used to determine the availability of the service and identify possible causes of

service failure. For more information on defining systems and monitoring them, refer to the Service Management chapter in *Oracle Enterprise Manager Concepts*.

- Setting up the availability definition for the service. This can be service test-based or system-based. If you select service test, the service's availability is based on the execution of the service test by the one or more key beacons. If availability is based on system, availability is based on the status of one or more key components of the system.
- Adding one or more beacons to monitor service tests. Click **Add** to add one or more beacons for monitoring the service. It is recommended that you use beacons that are strategically located in your key user communities in order for them to proactively test the availability of the service from those locations. If no beacons exist, click **Create** to create a new beacon.

  Note: Beacons marked as key beacons will be used to determine the availability of the service. The service is available if one or more service tests can be successfully executed from at least one key beacon.

  For Web applications, you can compare the performance of the service test execution from each remote beacon against the local beacon.

- Defining the metrics that will be used to measure the performance of the service. Performance metrics can be based on service tests or system components. After defining the metrics, you can specify the critical and warning thresholds. You can also specify the metric that is to be displayed in a graphical format on the Service Home page.
- Defining the metrics that will be used to measure the user demand for the service. Usage metrics can be based on one or more system components. After defining the metrics, you can specify the critical and warning thresholds. You can also specify the metric that is to be displayed in a graphical format on the Service Home page.

  Note: You can define usage metrics for system-based services only.

- After you have completed all the steps in the wizard, click **Finish** to create your service. Refer to the Enterprise Manager Online Help for more details on these pages.

7.4 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 7–2 Monitoring Configuration Page

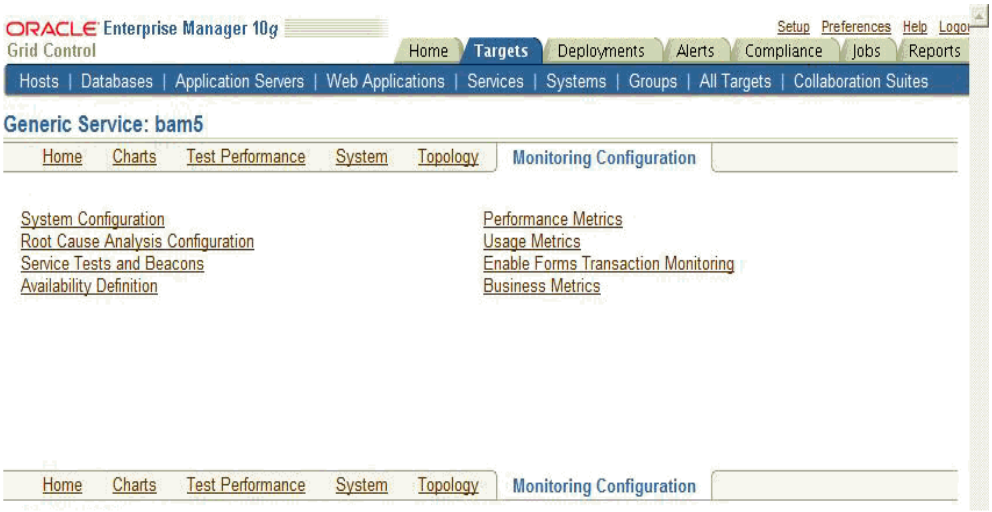

The following options are available:

- Availability Definition
- Performance Metrics
- Usage Metrics
- Business Metrics
- Service Tests and Beacons
- Root Cause Analysis Configuration

Apart from these options, for Web applications, the end-user and request performance monitoring features can also be configured. For more information, refer to the following sections:

- Configuring End-User Performance Monitoring
- Configuring OC4J for Request Performance Diagnostics

7.4.1 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 7-11.

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.

7.4.2 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 7–2 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 7–3 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.

7.4.3 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 7–4 Aggregation Functions - Usage Metrics

Function	Description
Maximum	The maximum value of the metric across all components will be used as the value of this usage metric for the service.
Minimum	The minimum value of the metric across all components will be used as the value of this usage metric for the service.
Average	The average value of the metric across all components will be used.
Sum	The sum of the values of metrics across all components will be calculated.

- Edit a usage metric that has been defined.
- Delete a usage metric that has been defined.

7.4.4 Business Metrics

Business metrics are used to measure the performance of business in an organization. These metrics are based on business indicators that can assess the business performance. You can define one or more system based metrics and specify critical and warning thresholds for these metrics. You can define business metrics for Generic Services and Aggregate Services.

> **Note:** This option is available only if one of the system components is a service and has business metrics associated with it.

You can do the following:

- Add a business metric. After selecting a metric for a target, you can choose to:
  - Use the metric from a specific component. Use this option if you want the business metric to be based on the performance of one specific system component
  - Aggregate the metric across multiple components. Use this option if you want to measure the business performance from multiple components. Select the appropriate aggregation function from the drop down list. If you choose this option, you need select the appropriate aggregation function.

Table 7–5 Aggregation Functions - Usage Metrics

Function	Description
Maximum	The maximum value of the metric across all components will be used as the value of this business metric for the service.
Minimum	The minimum value of the metric across all components will be used as the value of this business metric for the service.
Average	The average value of the metric across all components will be used.
Sum	The sum of the values of metrics across all components will be calculated.

- Edit a business metric that has been defined.
- Delete a business metric that has been defined.

You can define system based metrics only. You can configure non-system based metrics by using the Data Exchange feature which facilitates data transfer between Enterprise Manager Grid Control and other external monitoring systems. For details, refer to the *Oracle Enterprise Manager Integration Guide*.

7.4.5 Service Tests and Beacons

You can add additional service tests and specify one or more beacons that will execute these service tests. To add a service test or modify an existing service test, click the **Service Test and Beacons** link on the Monitoring Configuration page. The Service Tests and Beacons page appears. You can do the following:

- Add one or more service tests for your service. Select the Test Type and click **Add**. Some of the test types that can be defined are FTP, Web Transaction, DNS, SOAP

and others. The Create Service Test page is displayed. Refer to the Enterprise Manager Online Help for details on the various types of service tests.

Note: While defining a SOAP (Simple Object Access Protocol) service test, if the WSDL URL to be accessed is outside the company's intranet, proxy settings need to be added to the `$OMS_HOME/sysman/config/emoms.properties` file.

For example, to set up `www-proxy.us.oracle.com` as proxy, specify the values as follows:

```
proxyHost=www-proxy.us.oracle.com
proxyPort=80
dontProxyFor=us.oracle.com,oraclecorp.com
```

The `proxyUser`, `proxyPwd`, `proxyRealm`, and `proxyPropsEncrypted` properties are used to configure an authenticated proxy. After you have modified the proxy settings, you must restart the Oracle Management Service for the changes to be effective.

- After you have created the service test, you must enable it. If your service test is not enabled, it will not be executed by any of the beacons.
- Create, add, or remove a beacon. When you identify the beacon locations, select locations on your internal network or on the Internet that are important to your e-business. These are typical locations where your end users are located. For example, if your business is hosted in Canada and you have customers in the United States, use a beacon installed on a host computer in the United States to measure the availability and performance of your applications.
- After you have created the service test, you can verify it by clicking **Verify Service Test**.

Note: You can define one or more service tests as key tests. These key tests are used to monitor the availability and performance of your service. Only service tests that are enabled can be designated as key tests. To set up a service test as a key test, click the **Availability Definition** link at the bottom of the page.

For more details on creating different types of service tests, refer to the Enterprise Manager Online Help.

7.4.5.1 Configuring the Beacons

This section lists additional beacon related configuration tasks.

- **Configuring SSL Certificates for the Beacon**: To configure SSL certificates for Web transaction and Port Checker service tests, follow the steps given below:
  - For Web transactions, refer to instructions in the "Configuring Beacons to Monitor Web Applications Over HTTPS" on page 5-35.
  - 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 12-7.

- **Configuring Dedicated Beacons**: Beacon functionality on an agent requires the 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
```

7.4.6 Root Cause Analysis Configuration

You can use Root Cause Analysis (RCA) to filter a set of events to determine the cause of a higher level system, service, or application problem. RCA can help you to eliminate apparent performance problems that may otherwise appear to be root causes but which are only side effects or symptoms of the actual root cause of the problem, allowing you to more quickly identify problem areas. You can view the RCA results on the Home page or Topology page of any service that is currently down. The Topology page allows you to see a graphical representation of the service, system and component dependencies with the targets highlighted that RCA has implicated as causing the service failure.

Before running RCA, you can choose to:

- Configure the tool to run automatically whenever a service fails.
- Disable RCA by changing the default Analysis Mode to Manual.
- Define component tests for the service and thresholds for individual tests.

To configure Root Cause Analysis, follow these steps:

1. From the Service Home page, click **Monitoring Configuration**.
2. From the Monitoring Configuration page, click **Root Cause Analysis Configuration**.
3. If the current mode is set to Automatic, click **Set Mode Manual** to disable RCA. If you choose to perform the analysis manually, you can perform the analysis from the Service home page at anytime by choosing **Perform Analysis** if the service is down. If the current mode is set for Manual, click **Set Mode Automatic** to enable RCA when the state of the service and its components change

4. Click the link in the **Component Tests** column of the table for the key component you want to manage. You can then manage component tests for the service on the Component Tests page by adding, removing, or editing tests. Refer to the Enterprise Manager Online Help for details on defining component tests.

Note: When you disable RCA and set it back to automatic mode, RCA does not store the previous history results for you, thus providing no history for later reference.

7.4.6.1 Getting the Most From Root Cause Analysis

Root Cause Analysis (RCA) can provide you with great value by filtering through large amounts of data related to your services and identifying the most significant events that have occurred that are affecting your service's availability. If you are constructing your own services to manage in Enterprise Manager it is important that the services are defined with some thought and planning in order to get the most out of RCA.

The first item to consider in getting the most from RCA is the set of dependencies that your service has on other services or system components. Be sure to identify all of the system components that your service utilizes in order to accomplish its task. If you omit a key component and the service fails, RCA will not be able to identify that component as a possible cause. Conversely, if you include components in the service definition that the service does not actually depend on, RCA may erroneously identify the component as a cause of service failures.

When building service dependencies, keep in mind that you can take advantage of the aggregate service concept that is supported by Enterprise Manager. This allows you to break your service into smaller sub-services, each with its own set of dependencies.

Your services may be easier to manage in the modular fashion, and RCA will consider not only the status of a sub-service (a service that you depend on) but also on any of the system components or service that the sub-service depends on in turn and provides you with the power to encapsulate the services a key component exposes to you in the form of a managed service that your service may then depend on.

The second item to consider in getting the most from RCA is the use of component tests. As you define the system components that your service depends on, consider that there may be aspects of these components that may result in your service failure without the component itself failing. Component tests allow RCA to test the status not only of the target itself but also the status of its key aspects.

The RCA system allows you to create component tests based on any metric that is available for the key component. Remember, this includes any user-defined metrics that you have created for the component, allowing you great flexibility in how RCA tests the aspects of that component should your service fail.

7.5 Recording Web Transactions

You can record a transaction using an intuitive playback recorder that automatically records a series of user actions and navigation paths. You can play back transactions interactively, know whether it is internal or external to the data center, and understand the in-depth break-out of response times across all tiers of the Web application for quick diagnosis.

You must install the transaction recorder in your computer to record transactions. The transaction recorder is also used for playing back and tracing transactions. The

transaction recorder is downloaded from the Enterprise Manager Grid Control server the first time any of these actions is performed. The transaction recorder requires some Microsoft libraries to be installed in your computer. If these libraries are not present during installation, they are automatically downloaded and installed from the Microsoft site. Make sure that your computer has access to the Internet to download these files. After the installation has been completed, you may need to restart your computer to make the changes effective.

7.6 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.

7.7 Configuring Aggregate Services

Aggregate services consist of one or more services, called subservices. A subservice is any service created in Enterprise Manager. The availability, performance, business criteria, and usage for the aggregate service depend on the availability, performance, business criteria, and usage for the individual subservices comprising the service. To create an aggregate service, navigate to the Services main page, select Aggregate Service from the Add drop-down list and click **Go**. The Add Aggregate Service page appears. Creating an Aggregate Service involves the following:

- Specifying the name and time zone for the service.
- Adding the services that make up this aggregate service.
- Establishing the availability definition for the aggregate service. Availability of an aggregate service depends on the availability of its constituent subservices. The availability for a subservice may depend on the successful execution of a service test or on the availability of the system components on which the subservice runs, depending how the subservice was defined.
- Defining the metrics used to measure the performance of your aggregate service. You can add performance metrics from single subservices, or based on statistical aggregations of more than one metric. Once you have selected the performance metrics, you can set the thresholds used to trigger critical and warning alerts, or remove metrics you no longer want.
- Defining the metrics used to measure the usage of your aggregate service. Usage metrics can be based on the metrics of one or more system components. You can add usage metrics from single subservices, or based on statistical aggregations of more than one metric. Once you have selected the usage metrics, you can set the thresholds used to trigger critical and warning alerts, or remove metrics you no longer want.
- Defining the metrics that are used to measure of the performance of business in the organization. These metrics are based on business indicators that can assess the business performance.You can add business metrics from single subservices, or based on statistical aggregations of more than one metric. Once you have selected the business metrics, you can set the thresholds to trigger critical and warning alerts, or remove metrics you no longer want.

After you have created an aggregate service, you can add or remove its constituent subservices, modify the availability definition and add or delete performance or usage metrics. Refer to the Enterprise Manager Online Help for details on these operations.

WARNING: If you delete or remove a subservice from an aggregate service, the aggregate service performance, usage, and business metrics may be affected if they are based on a deleted subservice's metrics.

7.8 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

7.8.1 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 7-14.

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 7–3 Manage Web Server Data Collection

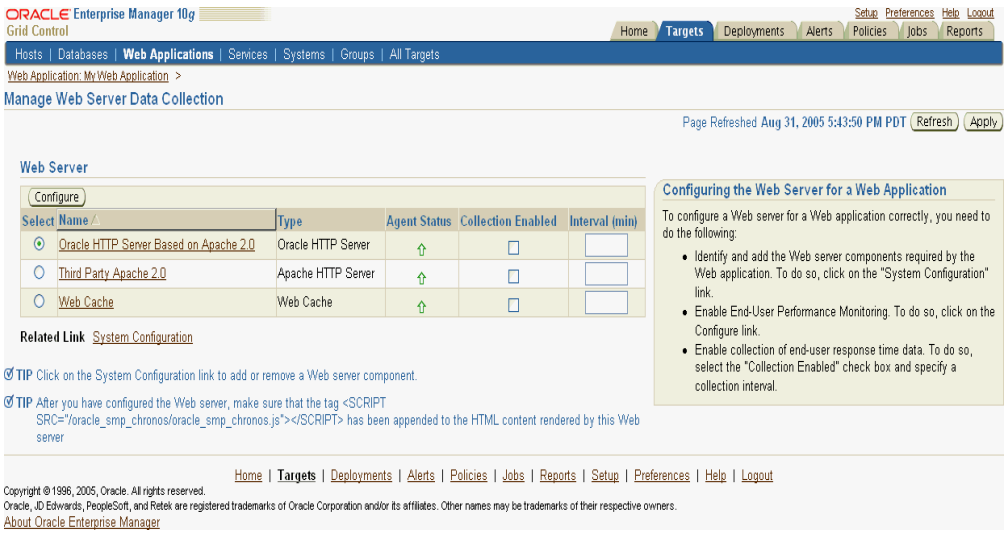

3. Select the Oracle HTTP Server or Apache HTTP Server from the table and click **Configure**. Enter the host credentials required for modifying the Apache configuration file.
4. After logging in, you will see a table containing the list of sites that are being hosted by the Apache server. These include 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 7-25. 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 7-25.

7.8.1.1 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.

7.8.2 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

7.8.2.1 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 7-25. If you do not see data after end-user performance has been enabled, refer to "Verifying and Troubleshooting End-User Performance Monitoring" on page 7-25.

7.8.2.2 Configuring Oracle Application Server Web Cache 9.0.4

To configure the Oracle Application Server Web Cache Manager 9.0.4, follow the instructions given in these sections:

1. Navigate to the Web Application home page in the Grid Control Console and click **Monitoring Configuration**.

2. Click **Manage Web Server Data Collection**. Enterprise Manager displays the Manage Web Server Data Collection page.

3. Select the Web Cache target and click **Configure**. Enterprise Manager displays the login dialog box for the Web Cache Manager.

   Tip: If the login dialog box does not appear or if you receive an error message in your browser window, you may have to start the Oracle Application Server Web Cache Manager. For more information about starting and using Oracle Application Server Web Cache Manager, refer to the *Oracle Application Server Web Cache Administrator's Guide*.

4. Enter the username and password for the Web Cache administrator account. The first time you log in to the Oracle Application Server Web Cache administrator account, the password is `administrator`. The password for the `ias_admin` account is defined during the installation of Oracle Application Server.

5. Enable OracleAS Web Cache logging for End-User Performance Monitoring:

   a. Select **Logging and Diagnostics** and then select End-User Performance Monitoring in the OracleAS Web Cache Manager navigator frame.

      You can enable monitoring for a particular Web cache or for an entire site.

   b. To enable monitoring for a particular Web cache, select the Web cache from the **Cache-Specific End-User Performance Monitoring** section and click **Enable**.

      Be sure to enable the Web cache that you are using as a front-end to your Web application.

   c. To enable monitoring for the entire site, select the site from the **Site-Specific End-User Performance Monitoring** section and click **Enable**.

6. Configure Oracle Application Server Web Cache to use the Web Cache Log Format (WCLF):
   a. Select **Logging and Diagnostics** and then select Access Logs in the OracleAS Web Cache Manager navigator frame.
   b. In the Cache-Specific Access Log Configuration table, click **Edit Selected** and enable the access log for your selected cache.
   c. In the Site-Specific Access Log Configuration table, make sure that the Format style of the selected Site Name is `WCLF` and that it is enabled.
7. Click **Apply Changes** at the top of the Web Cache Manager window and restart OracleAS Web Cache by clicking **Restart** on the Web Cache Manager Cache Operations page.
8. Close the Web Cache Manager browser window and return to the Manage Web Server Data Collection page in the Grid Control Console. You can now enable the collection of end-user performance data. For more details, refer to "Starting and Stopping End-User Performance Monitoring" on page 7-25. If you do not see data after end-user performance has been enabled, refer to "Verifying and Troubleshooting End-User Performance Monitoring" on page 7-25.

7.8.2.3 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 10*g* 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.

7.8.2.3.1 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.

7.8.2.3.2 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
     ```

7.8.2.3.3 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>`

7.8.2.3.4 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 10*g* 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

7.8.2.4 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 10*g* 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

7.8.2.4.1 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.

7.8.2.4.2 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.

7.8.2.4.3 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 7-25.

7.8.3 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.

7.8.4 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, proceed to the next step.
- If it is not present, check the configuration of your OracleAS Web Cache, Oracle HTTP Server, or Apache HTTP Server. Make sure that all configurations are correct, the site has been enabled, and the Web server has been successfully restarted after saving any configuration changes.

g. Go to the OracleAS Web Cache or Apache server target home page, click **Monitoring Configuration**, and check if the log file in the defined Log file directory contains any recent data.

- If it does not have data, go to the next step.
- If the log file does contain data and the Web server is OracleAS Web Cache, login to Oracle Application Server Control or Web Cache Manager and make sure that the access log is in WCLF or End-User Performance Monitoring format.

h. Verify that the OracleAS Web Cache / Apache server Monitoring Configuration properties that specify the location and name of the log file are accurate.

i. Check the Web Server target Home page for any collection errors. Often, the collection error will provide information describing why performance data cannot be collected.

j. Navigate to the All Metrics page for the Web server target and check to be sure the APM Mining Performance Details metrics are being collected successfully.

7.8.5 Enabling End-User Performance Monitoring for Third-Party Application Servers

For enabling End-User Performance Monitoring for third-party application servers like IBM WebSphere Application Server, BEA WebLogic Managed Server, and JBoss Application Server, after you configure one of the Web servers as explained in this chapter, you have to enable the Application Server Diagnostics Pack for the Web applications hosted on these servers.

To do so, perform the following steps:

1. Click **Setup** on the top-right corner of the Grid Control console and navigate to the Overview of Setup page.
2. Click **Management Pack Access** from the panel to the left.
3. On the Management Pack Access page, select the **All Targets** option in the View Options section of this page.
4. Select **Web Application** from the Search menu, and click **Go**. The table lists all the Web applications monitored.
5. For the Web application for which you want to enable End-User Performance Monitoring, check **Application Server Diagnostics Pack** and **Pack Access Agreed**, and then click **Apply**.
6. Now return to the Web Application Home Page and click **Page Performance** to see the end-user performance monitoring data that has been collected.

Note: End-User Performance Monitoring for a Web application is not supported if the J2EE container hosting that application is SSL enabled. This applies to Oracle J2EE containers, that is OC4J, and any non-Oracle J2EE containers for third-party application servers like BEA WebLogic Managed Server, IBM WebSphere Application Server, or JBoss Application Server. To activate End-User Performance Monitoring for such a Web application, disable SSL for that J2EE container.

For information about configuring SSL for Oracle Application Servers, refer to the Security Guide for your Oracle Application Server release. Documentation for all the Oracle Application releases is available from the Oracle Technology Network: `http://www.oracle.com/technology/documentation/index.html`.

For information about configuring SSL for third-party servers, refer to your third-party documentation.

7.9 Managing Forms Applications

A Forms Application target in Enterprise Manager can be used to model and monitor a specific Forms application. You can:

- Record and monitor a Forms transaction. See Recording and Monitoring Forms Transactions for details.
- Measure the End-User Performance of Forms actions such as Commit, Query, Runform, Callform, Openform, and Newform. See Monitoring the End-User Performance of Forms Applications for details.

7.9.1 Recording and Monitoring Forms Transactions

A Forms transaction consists of a set of user actions within a single application when using Forms. For example, an Update Employee Salary transaction may consist of several user actions like open salary form, update salary form, and save salary form. You can record multiple Forms transactions by using the intuitive playback recorder that automatically records a series of Forms actions.

Before recording a Forms transaction, you must do the following:

- Set the permissions of the `.java.policy` file on each Windows client. See Setting the Permissions of the .java.policy File on page 7-28.
- Ensure that a trusted Enterprise Manager certificate is used. See Using a Trusted Enterprise Manager Certificate on page 7-29.
- Add a certificate to the Enterprise Manager Agent to play back secure Forms transactions. See Adding a Forms Certificate to the Enterprise Manager Agent on page 7-30.
- Configure the Forms server so that Forms transactions can be recorded. See Configuring the Forms Server on page 7-30.

After you have performed these steps, you can install the transaction recorder to record and play back the Forms transaction. See Installing the Transaction Recorder to Record and Play Back Forms Transactions on page 7-31.

7.9.1.1 Setting the Permissions of the .java.policy File

You must the set the permissions of the `.java.policy` file on each Windows client on which the Forms transaction is being recorded. To set the permissions, follow these steps:

- Ensure that the `.java.policy` file is present under the user home directory. If the `.java.policy` file does not exist, you must create one as follows:
  - Create a `java.policy` (without the ".") file
  - Click **Start** and **Run** from your Windows desktop.
  - Type **cmd** and click **OK**.
  - At the DOS prompt, rename the file as follows:

    ```
    move java.policy .java.policy
    ```

- After you have created the `.java.policy` file, set the permissions for each Forms server or Oracle Applications server as follows:

  ```
  grant codeBase URL {
       permission java.security.SecurityPermission "putProviderProperty.SunJSSE";
  };
  ```

where *URL* needs to be replaced with the code source location of the Forms applet. By specifying the `codeBase`, you grant permissions to the code present in that location. For example, for an out-of-box Forms installation, you must specify the `codeBase` as follows:

```
http://formsServerHost:port/forms/java/*
```

where `formsServerHost` and `port` must be replaced with the host name and port number of the Forms server.

For Oracle Applications, you must specify the `codeBase` as follows:

```
http://appsHost:appsPort/OA_JAVA/oracle/apps/fnd/jar/*
```

where `appsHost` and `appsPort` must be replaced with the host name and port number of the Oracle Applications.

7.9.1.2 Using a Trusted Enterprise Manager Certificate

If you are using secure Enterprise Manager to record a Forms transaction running on Oracle Jinitiator or a Java plug-in, you must ensure that the Enterprise Manager certificate is trusted by Oracle Jinitiator and JPI. For Oracle Jinitiator, you must append the Enterprise Manager certificate to Jinitiator's certdb.txt file. For the Java Plug-in, you must set the certificate as trusted by JPI.

To ensure that the Enterprise Manager certificate is trusted by Jinitiator and JPI, follow these steps:

1. Export the Enterprise Manager certificate to a file.
   - When you launch secure Enterprise Manager, if Enterprise Manager is using a self generated certificate, you may see a "Certificate Error". Double click on the error and click **View Certificates**. The Certificate window is displayed.
   - Click the **Details** tab and then click **Copy to file...** to export the certificate to a file. The Certificate Export Wizard is displayed.
   - Click **Next** in the Welcome page.
   - In the Export File Format page, select `Base-64 encoded X.509(.CER)` and click **Next**.
   - Click **Browse** to select the name and the location of the file to which the certificate is to be saved.
   - Click **Finish**. The certificate has now been exported to a file.
2. After the certificate has been exported, you must set the certificate as trusted by Jinitiator or JPI.

   For Forms applications running on Oracle Jinitiator:

   - Open `certdb.txt` under `[Jinitiator InstallRoot]\lib\security\` directory. Usually Jinitiator is installed under `C:\ProgramFiles\Oracle\Jinitiator [version]`).
   - Use a text editor to open the file to which the certificate has been exported. Copy the contents and append it to `certdb.txt`.

   For Forms applications running on Java plug-in:

   - In the Control Panel, double click the Java program that is used to run the Forms application.

- Click the **Security** tab and then click **Certificates**.
- From the **Certificate Type** drop down list, select **Secure Site**.
- Click **Import** to import the file to the location in which the Enterprise Manager certificate has been saved.
- Close the certificate windows and the Java Control Panel.

3. Close the browser window. When the Forms application is accessed again, Jinitiator or JPI is restarted. This ensures that the changes to the security settings have been saved.

7.9.1.3 Adding a Forms Certificate to the Enterprise Manager Agent

To play back a secure Forms transaction, you must add a Forms certificate to the Enterprise Manager Agent by following these steps:

1. Stop the agent by entering the `emctl stop agent` command.
2. Create an importable certificate file from the forms server certificate (Base64 encoded X.509 format) and name this file as `forms.cer`.
3. Copy the `forms.cer` to `%AGENT_HOME%/jdk/jre/lib/security/` directory.
4. Run `keytool` with the following parameters (the `keytool` executable can be found under `the jdk/jre/bin` directory)

   ```
   keytool -import -alias forms -file %AGENT_HOME%/jdk/jre/lib/security/forms.cer
   -keystore AGENT_HOME%/jdk/jre/lib/security/cacerts
   ```

5. You will be prompted for the `cacerts` password. Enter `changeit` as the password.
6. Start the agent by entering the `emctl start agent` command.

7.9.1.4 Configuring the Forms Server

Before recording a Forms transaction, you must configure the Forms server by following these steps:

1. Create a system based Forms Application target that contains Forms, OracleAS Web Cache or Oracle HTTP Server / Apache HTTP Server targets. These targets must be a part of the system of the Forms Application. They must also be key components of your Forms Application or part of a key Redundancy Group. If you are using the Oracle HTTP Server, the Redundancy Group is referred to as the HTTP Server HA Group.
2. Set up the Forms server for recording transactions:
   1. Navigate to the Forms Application Home page in the Grid Control Console and click **Monitoring Configuration**.
   2. Click **Enable Forms Transaction Monitoring**.

      The Enable Forms Transaction Monitoring page is displayed.
   3. Select a Forms server from the list and click **Configure**.

      The Configure Forms Server: Login page is displayed.
   4. Enter the login credentials of the host on which Forms server is installed and click **Continue**.

      The jar files required for Forms Transaction Monitoring (`formsRecorder.jar`, `jsse.jar`, `jnet.jar`, and `jcert.jar`) are copied

into the Forms applet's archive directory (`ORACLE_HOME/forms/java`) and a confirmation message is displayed.

For Oracle Applications, the archive directory is located at `$JAVA_TOP/oracle/apps/fnd/jar.`

5. Click **Yes** to configure the Forms server and return to the Enable Forms Transaction Monitoring page.

After you have configured the system-based Forms Application target, you can record and play back Forms transactions to monitor the availability of the Forms application. To do so, navigate to the Monitoring Configuration page and click Availability Definition. In this page, change the Availability Definition to Service Test.

7.9.1.5 Installing the Transaction Recorder to Record and Play Back Forms Transactions

After you have configured the Forms server, you can install the transaction recorder on your computer. The transaction recorder is downloaded from the Enterprise Manager Grid Control server the first time you access the Record Forms Transaction page. The transaction recorder requires some Microsoft libraries to be installed in your computer. Make sure that your computer has access to the Internet to download these files. If these libraries are not present during installation, they are automatically downloaded and installed from the Microsoft site. After you have recorded a Forms transaction, if you need to record another one in the same browser, you must use the same JVM version for the new transaction.

You can record multiple Forms transactions on the Forms Application target and monitor these transactions periodically. Before recording a Forms transaction, ensure that all other Forms applications are closed. When you a record a Forms transaction, make sure that the following parameters are specified correctly:

- **Login URL**: If you selected the Login Type as **Single Sign-On (SSO)** or **Oracle Applications Login**, the Login URL must be explicitly specified.
- **Connection Type**: This can be:
  - **Socket**: Ensure that the Forms server host name and port number are specified correctly.
  - **HTTP / HTTPS**: If the Connection Type is HTTPS and a non-standard certificate is being used, you must import the certificate into the Agent Home directory.

For more details on recording a Forms transaction and metrics collected, refer to the Enterprise Manager Online Help.

7.9.1.6 Forms Server Version Support

When playing back a Forms transaction, you may encounter a "`Version Not Supported`" error. To resolve this error, follow these steps:

1. Login to the system on which the Forms server has been installed. Locate the `frmall.jar` (if you are using Forms 10.1 or later) or `f90all.jar` (if are using Forms 9.0.4 or later) under the `$FORMS_HOME/forms/java` directory.
2. Login to the system on which the beacon has been deployed and copy the jar file to the `$ORACLE_HOME/jlib/forms/<version>/` directory. The version you specify here should be the same as the version string in the error message. Make sure that the directory is empty before you copy over the jar file.

If you are using Oracle Applications R12 and you encounter this error, follow these steps to resolve the error:

1. Login to the system in which the Oracle Application server has been deployed. Locate the following files:

   ```
   $JAVA_TOP/oracle/apps/fnd/jar/fndforms.jar
   $JAVA_TOP/oracle/apps/fnd/jar/fndewt.jar
   ```

2. Login to the system on which the beacon has been deployed and copy these files to the `$ORACLE_HOME/jlib/forms/apps/` directory. Make sure that the directory is empty before you copy over the jar files.

> **Note:** You cannot monitor two deployments of Oracle Applications from the same beacon if different versions of Oracle Applications have been used.

7.9.2 Monitoring the End-User Performance of Forms Applications

The End-User Performance Monitoring utility allows you to measure the response time of your applications by viewing information about how quickly the responses are delivered to the end users. When you access a Forms application, the End-User Performance Monitoring utility measures the response time of Forms actions such as Commit, Query, Runform, Callform, Newform, and Openform.

You can monitor the Forms actions and view reports based on the response times experienced by the user. You can also define a Watch List of the most important Forms actions to monitor and view the response metrics of these critical operations at a glance.

> **Note:** End-User Performance Monitoring is supported with Forms server version 6i Patch 16, 10g R2. For version 6i Patch 16, only the Commit operation can be monitored.

Before you can begin monitoring the End-User Performance of a Forms Application, you must configure the Forms and Web server to enable data collection for End-User Performance Monitoring. To configure the Forms Application for End-User Performance Monitoring, follow these steps:

- Configure the Forms server to enable End-User Performance Monitoring.
- Configure the Web server (OracleAS Web Cache or Oracle HTTP Server / Apache HTTP Server) so that it can be used for End-User Performance Monitoring.
- Enable the collection of end-user performance data.

7.9.2.1 Configuring the Forms Server for End-User Performance Monitoring

Before you can enable the collection of end-user performance data, you must first configure the Forms server. To configure the Forms server, follow these steps:

1. Navigate to the Forms Application Home page in Enterprise Manager Grid Control.
2. Click **Monitoring Configuration**.
3. Click **Manage Web Server Data Collection**.

4. On the Manage Web Server Data Collection page, select the Forms server and click **Configure**. The Configure Forms Server for End-User Performance Monitoring: Login page is displayed.
5. Enter the host login credentials and click **Continue**. The Configure Forms for End-User Performance Monitoring: Configuration Sections page is displayed.
6. Select a section and check the Enable Monitoring checkbox to enable End-User Performance Monitoring on that section. Click **Enable All** or **Disable All** to enable or disable all the sections. You can also click **Add New Section** to add a section without affecting existing sections. After adding the section, you can enable End-User Performance Monitoring by selecting the checkbox. You can also delete a section that you have added.

   Tip: A **section** is a parameter defined in the `formsweb.cfg`. It specifies which section of Forms configuration the user wants to run. The section usually includes the application name and other relevant parameters which are required for successful execution of the application.

7. Set the value of the End-User Performance Monitoring URL column to `http://<hostname:portnumber>/oracle_smp_chronos/oracle_smp_chronos_sdk.gif`. The hostname and port number are for the Web Server that is serving the Forms application.
8. After you have configured the Forms server, click **OK** to save the changes and return to the Manage Web Server Data Collection page.

7.9.2.2 Configuring the OracleAS Web Cache

You can use the 10.1.2 or 9.0.4 versions of OracleAS Web Cache to collect end-user performance data.

- **OracleAS Web Cache 10.1.2:** To configure OracleAS Web Cache 10.1.2, follow these instructions:
  1. You can configure OracleAS Web Cache by using the Oracle Application Server Control. Navigate to the Forms Application home page in the Enterprise Manager Grid Control.
  2. Click **Monitoring Configuration**.
  3. Click **Manage Web Server Data Collection**.
  4. On the Manage Web Server Data Collection page, select the Web Cache target and click **Configure**. The Application Server Control login dialog box is displayed.

     Tip: If the login dialog box does not appear or if you receive an error message in your browser window, navigate to the Web Cache Home page and click Administer under the Related Links. You will be prompted for the user name and password for Application Server Control. Click **Administration**, scroll down and click End-User Performance Monitoring.

     If Application Server Control is not available, you can also use the Oracle Application Server Web Cache Manager to configure the OracleAS Web Cache for End-User Performance Monitoring. For more information about starting

and using Oracle Application Server Web Cache Manager, refer to the Oracle Application Server Web Cache Administrator's Guide.

5. Enter the username and password for the Web Cache administrator account or the `ias_admin` account. The password for the `ias_admin` account is defined during the installation of Oracle Application Server.

   After you have logged into Oracle Application Server Control, you can configure OracleAS Web Cache from the Set Up End-User Performance Monitoring page.

6. Select the Access Log Format as `access log:WCLF` for each site from the drop down list. If this format is not in the list, click **Use Required Log Format**.

7. You will return to the Web Cache Administration page in Oracle Application Server Control. Click **Restart** to restart the Web Cache. For more detailed information about configuring these options, refer to the Enterprise Manager Online Help.

8. Close the Oracle Application Server Control browser window and return to the Manage Web Server Data Collection page in the Enterprise Manager Grid Control.

- **OracleAS Web Cache 9.0.4:** To configure OracleAS Web Cache 9.0.4, follow these instructions:

  1. You can configure OracleAS Web Cache by using the Oracle Application Server Web Cache Manager. Navigate to the Forms Application home page in the Enterprise Manager Grid Control.

  2. Click **Monitoring Configuration**.

  3. Click **Manage Web Server Data Collection**.

  4. On the Manage Web Server Data Collection page, select the Web Cache target and click **Configure**. A login dialog box is displayed.

     Tip: If the login dialog box does not appear or if you receive an error message in your browser window, navigate to the Web Cache Home page and click Administer under the Related Links. You will be prompted for the user name and password for Application Server Control. Click **Administration**, scroll down and click End-User Performance Monitoring.

  5. Enter the username and password for the Web Cache administrator account. The first time you log in to the Oracle Application Server Web Cache administrator account, the password is `administrator`.

  6. Configure Oracle Application Server Web Cache to use the Web Cache Log Format (WCLF):

     – Select **Logging and Diagnostics** and then select **Access Logs** in the OracleASWeb Cache Manager navigator frame.

     – In the Cache-Specific Access Log Configuration table, click **Edit Selected** and enable the access log for your selected cache.

     – In the Site-Specific Access Log Configuration table, make sure that the Format style of the selected Site Name is `WCLF` and that it is enabled.

     For more details on changing the `access_log` format, refer to the Enterprise Manager Online Help.

7. Click **Apply Changes** at the top of the Oracle Application Server Web Cache Manager window and restart Oracle Application Server Web Cache by clicking **Restart** on the Oracle Application Server Web Cache Manager Cache Operations page.
8. Close the Oracle Application Server Web Cache Manager window and return to the Manage Web Server Data Collection page in the Grid Control console. You can now enable the collection of end-user performance data.

7.9.2.3 Configuring the Oracle HTTP Server / Apache HTTP Server

You can collect end-user performance data by using Oracle HTTP Server or Apache HTTP Server. Before you use these server, follow these steps:

1. On the Agent Home page, select the Oracle HTTP Server or Apache HTTP Server target type. If you are using a generic third party Apache server, select a Apache HTTP Server target.
2. Add the target of the corresponding type and make sure that the Log file directory and Log file name properties are set in the Monitoring Configuration page.

   The Log file directory and Log file name you specify here will be used by the End-User Performance Mining Engine to upload end-user performance data.

   > **Note:** If the Oracle HTTP Server is installed before the Management Agent has been installed, and is up and running during agent installation, then the target will be discovered automatically. Otherwise you need to manually create the Oracle HTTP Server target and specify the following properties: `Machine name`, `Port number`, `Version of the Apache Server`, `Oracle home path`, `Log file directory` (for EUM), `Log file name` (for EUM) where EUM refers to End-UserPerformance Monitoring.

3. Create a system target and a Forms Application target. Add the Oracle HTTP Server or Apache HTTP Server target to the system target, and make it a key component of the Forms Application target or a part of a key Redundancy Group target. If you are using Oracle HTTP Server, the Redundancy Group is referred to as HTTP Server HA Group.
4. Navigate to the Monitoring Configuration page for the Forms Application target that contains the Oracle HTTP Server or Apache HTTP Server target. Click **Manage Web Server Data Collection**. You will see a table which lists the Web Servers including Oracle HTTP Server, Apache HTTP Server, or OracleAS Web Cache.
5. Select the Oracle HTTP Server or Apache HTTP Server from the table and click **Configure**. Enter the username and password for the host on which the Oracle HTTP Server or Apache HTTP server is installed.
6. After logging in, you will see a table containing the list of sites that are being hosted by the Apache server. These include a list of virtual hosts defined by the user in the Apache Configuration file. The up and the down arrows under the **Monitoring Status** column shows the corresponding site is currently being monitored. For each site, check or uncheck the **Enable Monitoring** checkbox to indicate whether this site is to be monitored. For the site that is to be monitored, enter the log file name in the text box to indicate the location in which the end-user performance data is to be stored. By default, the log file will be created under the

`logs/directory` under Apache root directory. To save the log file in a different directory, enter a file name with the absolute path.

7. Make sure that the log file name you specify here matches the Log file directory and Log file name in Monitoring Configuration page of the Oracle HTTP Server or Apache HTTP Server target.
8. You can also use the one button accelerator to enable all sites or disable all sites all at once.
9. After you have made the configuration changes, click **OK** to go to the Apache Restart page. Restarting the Apache server will finalize all configuration changes, and end-user performance data will be logged by the Apache server.
10. After you have configured the Web server, you must configure the Forms server and enable collection of the End-User Performance data from the Manage Web Server Data Collection Page. For details on configuring the Forms server, refer to the Enterprise Manager Online Help.

7.9.2.4 Starting and Stopping End-User Performance Monitoring

After you have configured the Forms and Web server to enable collection, you can then start collecting end-user performance data.

1. Navigate to the Web Application home page in the Grid Control Console and click **Monitoring Configuration**.
2. Click **Manage Web Server Data Collection**. Enterprise Manager displays the Manage Web Server Data Collection page.
3. In the **Interval (minutes)** column, enter the interval at which Enterprise Manager will collect performance data.
4. Check the **Collection Enabled** checkbox.
5. Click **Apply**, review the changes and confirm by clicking **Apply** again. End-User Performance Monitoring collection is enabled and data will soon be uploaded to the database and shown under the Page Performance page.

To stop collecting end-user performance data:

1. Navigate to the Manage Web Server Data Collection page.
2. Clear the check box in the **Collection Enabled** column of the table and click **Apply**.
3. Click **Apply** again to confirm the changes.

7.10 Configuring OC4J for Request Performance Diagnostics

Enterprise Manager can gather critical request performance data about your Web application and display this performance data. This feature can be instrumental when you are diagnosing application server and back-end performance issues.

Before you can begin collecting request performance data, you must do the following:

- Create a Web application target and associate it with a system that contains the OC4J instances to be monitored.
- Make these OC4J instances as key system components for your Web application and enable the logging and tracing capabilities. If these OC4J instances are a part of an OC4J Cluster, make sure that this OC4J Cluster is a key system component of

your Web application. To enable request performance monitoring, you must configure the specific OC4J instance within the OC4J cluster.

For more information, see the following:

- Selecting OC4J Targets for Request Performance Diagnostics
- Configuring Interactive Transaction Tracing
- Configuring OC4J Tracing for Request Performance Data
- Additional Configuration for Monitoring UIX Applications

7.10.1 Selecting OC4J Targets for Request Performance Diagnostics

Before you configure the OC4J target to collect request performance data, follow the steps given below to add the target to the Web application.

1. Configure the system where the OC4J targets are defined for the Web application target.
2. Navigate to the Web application Home page and click **Monitoring Configuration**.
3. Click **System Configuration**. From the list of system components displayed on this page, select one or more OC4J targets and select the checkbox in the **Key Components** column. The OC4J targets can now be configured and used to collect request performance data.

7.10.2 Configuring Interactive Transaction Tracing

When you use transactions to monitor your Web application, some of the transactions you create often involve application components such as servlets, Java Server Pages (JSPs), Enterprise Java Beans (EJBs), and database connections. Often, the best way to solve a performance problem is to trace these more complex transactions and analyze the time spent processing each application component.

Enterprise Manager provides a mechanism for tracing these transactions. Use the **Service Tests and Beacons** link on the **Monitoring Configuration** page of the Web application target to create your transactions and to trace the transactions as they are processed by the servlets, JSPs, EJBs, or database connections of your application.

However, before you can take advantage of transaction tracing, you must first enable tracing for the OC4J instance used to deploy the application. Each OC4J instance of an OC4J cluster must be configured independently. The OC4J instances of the OC4J clusters selected as key components of the Web application target are displayed on the Manage Web Server Data Collection page.

To enable tracing for an OC4J instance:

1. Navigate to the Web Application Home page and click **Monitoring Configuration**.
2. Click **Manage OC4J Data Collection.**

   Enterprise Manager displays the Manage OC4J Data Collection page.
3. Select the OC4J to configure and click **Enable Logging**.

   Enterprise Manager opens another browser window and displays the Tracing Properties page for the OC4J instance in the Application Server Control.

   If you are prompted to log in to the Application Server Control Console, enter the credentials for the `ias_admin` administrator's account.

4. Select the following options on the Tracing Properties page:
   - **Enable JDBC/SQL Performance Details**
   - **Enable Interactive Trace**

   You can use the default values for most of the tracing properties.

   > **Note:** Turning on the **Enable JDBC/SQL Performance Details** option allows to you drilldown to actual SQL statements but this may require more resources.

5. Click **Apply**.

   If this is the first time you are enabling OC4J tracing for this application server, Enterprise Manager displays a message stating that the `transtrace` application is being deployed. The Application Server Control then prompts you to restart the OC4J instance.

6. Click **Yes** to restart the instance and enable the tracing properties.
7. Return to the Grid Control Console.

   Tracing is now enabled for the selected OC4J instance.

7.10.3 Configuring OC4J Tracing for Request Performance Data

You must configure OC4J instances to enable tracing so that request performance data can be collected. Each OC4J instance of an OC4J cluster must be configured independently. The OC4J instances of the OC4J clusters selected as key components of the Web application target are displayed on the Manage Web Server Data Collection page. To configure the OC4 instances, follow these steps:

1. Navigate to the Web Application home page and click Monitoring Configuration.
2. Click **Manage OC4J Data Collection**.

   Enterprise Manager displays the Manage OC4J Data Collection page.

3. For the OC4J instance that you used to deploy your application, select the check box in the **Collection Enabled** column.
4. In the Interval (minutes) column, enter the interval at which to collect OC4J tracing data.

   The recommended interval setting is 60 minutes.

5. Select the OC4Js to configure and click **Enable Logging**.

   Enterprise Manager opens another browser window and displays the Tracing Properties page for the OC4J instances in the Application Server Control.

   If you are prompted to log in to the Application Server Control Console, enter the credentials for the `ias_admin` administrator's account.

6. Select the following options on the Tracing Properties page:
   - **Enable JDBC/SQL Performance Details**
   - **Enable Historical Trace**

   You can use the default values for most of the tracing properties. However, Oracle recommends that you set the **Frequency to Generate Trace File (seconds)** field to `3600` seconds (equivalent to 60 minutes).

> **Note:** Modifying the value in the **Trace File Directory** field is not supported.

7. Click **Apply**.

   If this is the first time you are enabling OC4J tracing for this application server, Enterprise Manager displays a message stating that the `transtrace` application is being deployed. The Application Server Control then prompts you to restart the OC4J instance.

8. Click **Yes** to restart the instance and enable the tracing properties.

9. Return to the Grid Control Console.

   Request Performance data should begin to appear on the Request Performance page as soon as data for the OC4J instance is collected and uploaded into the Management Repository.

7.10.4 Additional Configuration for Monitoring UIX Applications

If you used Oracle User Interface XML (UIX) to build your application, there is an additional configuration step you must perform before you can monitor the requests of your application.

> **See Also:** Your JDeveloper documentation for information on using UIX to develop Web applications

Before you can monitor the requests of your UIX application, you must do the following:

1. Enable tracing for the OC4J instance you used to deploy your application, as described in "Configuring OC4J Tracing for Request Performance Data" on page 7-38.

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.

7.11 Setting Up Monitoring Templates

A monitoring template for a service contains definitions of one or more service tests, as well as a list of monitoring beacons. A monitoring template can be used to create service tests on any number of service targets, and specify a list of monitoring beacons.

A monitoring template must be created from a service target. Once the template is created, the user can edit the template, create copies, or delete it. Finally, the user can apply the template to other targets, which creates the service tests on the other targets and adds the monitoring beacons.

To create a Monitoring Template, follow the steps given below:

1. Click **Setup** to navigate to the main Setup page in Enterprise Manager.
2. Click the **Monitoring Templates** link in the left panel.
3. Click **Create** to create a monitoring template.
4. In the target selection box, enter or select a service target and click **Continue**.
5. In the Monitoring Template General Page, enter the name of the template that you wish to create.
6. Click **Tests** to add / remove or configure service tests associated with the selected service target. Make the required changes to this page and click **OK** to save the template to the repository.

After you have created the Monitoring Template, use the **Apply** option to apply this template to a service test. You can click **Edit** to modify the template. For more details on these operations, refer to the Online Help.

7.11.1 Configuring Service Tests and Beacons

You can configure the service tests and beacons associated with the template by using the options in the **Tests** page. A service test-based template contains the following elements:

- **Variables**: A variable may occur at multiple locations in the service tests. The Variables table allows you to specify default values for all the variables. These default values will be stored in the template along with the variables. You can specify values other than the default while applying the template to a target. You can perform the following operations:
  - **Add** a variable. The variable can consist of letters, numbers and underscores only.
  - **Rename** a variable. When you rename a variable, all variable references in the service tests will be replaced with the new name.
  - **Remove** variables for properties within service tests. If you remove a non-password variable, all references to the variable in test properties will be replaced with the variable's default value
  - **Replace Text** in test properties with a variable definition.
- **Service Tests:** You can edit the test definition and define variables for various properties. You can select the tests from the original target that are to be part of the template by clicking the **Add / Remove** button. You can specify whether the service test is a key test and if it should be enabled. You can also click **Monitoring Settings** to drill down to this page and define metrics and thresholds for the service tests.
- **Beacons:** Use the **Add / Remove** button to specify which beacons are to be included in the template. You can also specify whether each beacon is a key beacon.

Refer to the Enterprise Manager Online Help for detailed instructions on these operations.

7.12 Configuring Service Levels

A service level rule is defined as an assessment criteria used to determine service quality. It allows you to specify availability and performance criteria that your service must meet during business hours as defined in your Service Level Agreement. For example, e-mail service must be 99.99% available between 8am and 8pm, Monday through Friday.

A service level rule specifies the percentage of time a service meets the performance and availability criteria as defined in the Service Level Rule. By default, a service is expected to meet the specified criteria 85% of the time during defined business hours. You may raise or lower this percentage level according to service expectations. A service level rule is based on the following:

- **Business Hours:** Time range during which the service level should be calculated as specified in your Service Level Agreement.
- **Availability:** Allows you to specify when the service should be considered available. This will only affect the service level calculations and not the actual availability state displayed in the console. You can choose a service to be considered up when it is one or more of the following states:
  - Up: By default the service is considered to be Up or available.
  - Under Blackout: This option allows you to specify service blackout time (planned activity that renders the service as technically unavailable) as available service time.
  - Unknown: This option allows you to specify time that a service is unmonitored because the Management Agent is unavailable be counted as available service time.
- **Performance Criteria**: You can optionally designate poor performance of a service as a Service Level violation. For example, if your Website is up, but it takes 10 seconds to load a single page, your service may be considered unavailable.
- **Business Criteria:** Business criteria are useful in determining in the health of the business processes for a particular service. You can optionally define business metrics that can affect the Service Level. A Service Level violation occurs when a critical alert is generated for a specified business metric.

  Note: The **Business Criteria** column is displayed only if one or more key business indicators are associated with the service. Refer to *Oracle Enterprise Manager Integration Guide*.

- **Actual Service Level:** This is calculated as percentage of time during business hours that your service meets the specified availability, performance, and business criteria.
- **Expected Service Level:** Denotes a minimum acceptable service level that your service must meet over any relevant evaluation period.

You can define only one service level rule for each service. The service level rule will be used to evaluate the **Actual Service Level** over a time period and compare it against the **Expected Service Level**.

7.12.1 Defining Service Level Rules

When you create a service, the default service rule is applied to the service. However, you must edit the service level rule for each service to accurately define the assessment criteria that is appropriate for your service. To define a service level rule:

1. Click the **Targets** tab and **Services** subtab. The Services main page is displayed.
2. Click the service name link to go to the Service Home page.
3. In the Related Links section, click **Edit Service Level Rule**.
4. On the Edit Service Level Rule page, specify the expected service level and the actual service level and click **OK**. The expected service level specifies the percentage of time a service meets the performance, usage, availability, and business criteria defined in the Service Level Rule. The actual service level defines the baseline criteria used to define service quality and includes business hours, availability, performance criteria, usage criteria, and business criteria.

Note: Any Super Administrator, owner of the service, or Enterprise Manager administrator with OPERATOR_TARGET target privileges can define or update the Service Level Rule.

7.12.2 Viewing Service Level Details

You can view service level information directly from the either of the following:

- **Enterprise Manager Grid Control Console** -From any Service Home page, you can click on the Actual Service Level to drill down to the Service Level Details page. This page displays what Actual Service Level is achieved by the service over the last 24 hours/ 7 days / 31 days, compared to the Expected Service Level. In addition, details on service violation and time of each violation are presented in both graphical and textual formats.
- **Information Publisher** - Information Publisher provides an out-of-box report definition called the Services Dashboard that provides a comprehensive view of any service. From the Report Definition page, click on the **Services Monitoring Dashboard** report definition to generate a comprehensive view of an existing service. By default, the availability, performance, status, usage, business, and Service Level of the service are displayed. The Information Publisher also provides service-specific report elements that allow you to create your own custom report definitions. The following report elements are available:
  - **Service Level Details**: Displays **Actual Service Level** achieved over a time-period and violations that affected it.
  - **Service Level Summary**: Displays service level violations that occurred over selected time-period for a set of services.
  - **Services Monitoring Dashboard**: Displays status, performance, usage, business, and service level information for a set of services.
  - **Services Status Summary**: Information on one or more services' current status, performance, usage, business, and component statuses.

    Refer to the Online Help for more details on the report elements.

7.13 Configuring a Service Using the Command Line Interface

Using the Command Line Interface, you can define service targets, templates and set up alerts. EM CLI is intended for use by enterprise or system administrators writing scripts (shell/batch file, perl, tcl, php, etc.) that provide workflow in the customer's business process. EM CLI can also be used by administrators interactively, and directly from an operating system console. Refer to *Enterprise Manager Command Line Interface Guide* for details.

8

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

8.1 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

8.1.1 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 12-4

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`

8.1.2 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.

8.1.3 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
```

8.1.4 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 8–1 to locate and modify the Agent logging and tracing properties in the `emd.properties` file.
5. Restart the Management Agent.

Table 8–1 Management Agent Log and Trace File Properties

Property	Purpose	Example
`LogFilewithPID`	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.	`LogFilewithPID=true`
`LogFileMaxSize`	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 emagent.log` logging file.	`LogFileMaxSize=4096`
`LogFileMaxRolls`	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.	`LogFileMaxRolls=4`
`TrcFileMaxSize`	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 emagent.trc` logging file.	TrcFileMaxSize=4096
TrcFileMaxRolls	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.	TrcFileMaxRolls=4

8.1.5 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 8–2.

   Note: The values described in Table 8–2 are case-sensitive.

5. Restart the Management Agent.

Table 8–2 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.

8.1.6 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 8–3.

4. Restart the Management Agent.

Table 8–3 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 `emagentfetchlet.log` file.	log4j.appender.emagentlogAppender.MaxFileSize=20000000
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 `emagentfetchlet.trc` 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

8.1.7 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 8–2.

   Note: The the values described in Table 8–2 are case-sensitive.

5. Restart the Management Agent.

8.2 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

8.2.1 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.

8.2.2 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/
```

8.2.3 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 8–4.

4. Restart the Management Service.

Table 8–4 Management Service Log File Properties in the emomslogging.properties File

Property	Purpose	Example
log4j.appender.emlogAppender.MaxFileSize	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 `emoms.log` log file.	log4j.appender.emlogAppender.MaxFileSize=20000000
log4j.appender.emlogAppender.MaxBackupIndex	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.	log4j.appender.emlogAppender.MaxBackupIndex=1
log4j.appender.emtrcAppender.MaxFileSize	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 `emoms.trc` log file.	log4j.appender.emtrcAppender.MaxFileSize=5000000
log4j.appender.emtrcAppender.MaxBackupIndex	This property indicates how many times the Management Services will rollover the trace file to a new file name before deleting tracing data.	log4j.appender.emtrcAppender.MaxBackupIndex=10

8.2.4 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 8–2.

   Note: The values described in Table 8–2 are case-sensitive.

5. Restart the Management Service.

8.2.5 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 8–5 lists the location of the log files for some components.

Table 8–5 Component Log File Location

Component	Location
HTTP Server	ORACLE_HOME/Apache/Apache/logs/error_log.timeORACLE_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.

9

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
- Changing the SYSMAN Password
- Dropping and Recreating the Management Repository
- Troubleshooting Management Repository Creation Errors
- Improving the Login Performance of the Console Home Page

9.1 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

9.2 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
    ```

9.3 Dropping and Recreating the Management Repository

This section provides information about dropping the Management Repository from your existing database and recreating the Management Repository after you install Enterprise Manager.

9.3.1 Dropping the Management Repository

To recreate the Management Repository, you first remove the Enterprise Manager schema from your Management Repository database. You accomplish this task using the `-action drop` argument to the `RepManager` script, which is described in the following procedure.

To remove the Management Repository from your database:

1. Locate the `RepManager` script in the following directory of the 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 efkl34lmn -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

9.3.2 Recreating the Management Repository

The preferred method for creating the Management Repository is to create the Management Repository during the Enterprise Manager installation procedure, which is performed using Oracle Universal Installer.

> **See Also:** *Oracle Enterprise Manager 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

9.3.2.1 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 9-1.
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.

9.3.2.2 Using a Connect Descriptor to Identify the Management Repository Database

Alternatively, you can use a connect descriptor to identify the database on the `RepManager` command line. The connect descriptor identifies the host, port, and name of the database using a standard Oracle database syntax.

For example, you can use the connect descriptor as follows to create the Management Repository:

```
$PROMPT> ./RepManager -connect "(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)
(HOST=host1)(PORT=1521)) (CONNECT_DATA=(SERVICE_NAME=servicename)))"
-sys_password efkl34lmn -action create
```

> **See Also:** "Establishing a Connection and Testing the Network" in the *Oracle Database Net Services Administrator's Guide* for more information about connecting to a database using a connect descriptor

The ability to use a connect string allows you to provide an address list as part of the connection string. The following example shows how you can provide an address list consisting of two listeners as part of the `RepManager` command line. If a listener on one host becomes unavailable, the second listener can still accept incoming requests:

```
$PROMPT> ./RepManager -connect "(DESCRIPTION=
(ADDRESS_LIST=
(ADDRESS=(PROTOCOL=TCP)(HOST=host1)(PORT=1521)
(ADDRESS=(PROTOCOL=TCP)(HOST=host2)(PORT=1521)
(CONNECT_DATE=(SERVICE_NAME=servicename)))"
-sys_password efkl34lmn -action create
```

> **See Also:** *Oracle High Availability Architecture and Best Practices*
>
> "Configuring the Management Services" on page 3-14

9.4 Troubleshooting Management Repository Creation Errors

Oracle Universal Installer creates the Management Repository using a configuration step at the end of the installation process. If the repository configuration tool fails, note the exact error messages displayed in the configuration tools window, wait until the other configuration tools have finished, exit from Universal Installer, and then use the following sections to troubleshoot the problem.

9.4.1 Package Body Does Not Exist Error While Creating the Management Repository

If the creation of your Management Repository is interrupted, you may receive the following when you attempt to create or drop the Management Repository at a later time:

```
SQL> ERROR:
ORA-00604: error occurred at recursive SQL level 1
ORA-04068: existing state of packages has been discarded
ORA-04067: not executed, package body "SYSMAN.MGMT_USER" does not exist
ORA-06508: PL/SQL: could not find program unit being called
ORA-06512: at "SYSMAN.SETEMUSERCONTEXT", line 5
ORA-06512: at "SYSMAN.CLEAR_EMCONTEXT_ON_LOGOFF", line 4
ORA-06512: at line 4
```

To fix this problem, see "General Troubleshooting Techniques for Creating the Management Repository" on page 9-6.

9.4.2 Server Connection Hung Error While Creating the Management Repository

If you receive an error such as the following when you try to connect to the Management Repository database, you are likely using an unsupported version of the Oracle Database:

```
Server Connection Hung
```

To remedy the problem, upgrade your database to the supported version as described in *Oracle Enterprise Manager Grid Control Installation and Basic Configuration*.

9.4.3 General Troubleshooting Techniques for Creating the Management Repository

If you encounter an error while creating the Management Repository, drop the repository by running the `-drop` argument to the `RepManager` script.

See Also: "Dropping the Management Repository" on page 9-3

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;
   ```

9.5 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.

10

Using Enterprise Manager For Grid Automation With Deployment Procedures

Deployment procedures are Oracle's latest contribution in automating operations around the grid. This chapter introduces the concept of deployment procedures to system administrators and integrators. The chapter spells out the advantages and features of deployment procedures and discusses a sampling of use cases that these deployment procedures are designed to solve.

Deployment procedures are out-of-box best practices that comprise enumeration of a set of steps that are orchestrated by Oracle Enterprise Manager. Oracle ships a set of "best practice" deployment procedures to accomplish provisioning and patching related tasks. Deployment procedures can be extended and customized for customer needs. The deployment procedure to patch a single instance database differs from the one to patch a Data Guard or a RAC environment. Deployment procedures can vary from one customer to another, or from a test installation to a production installation.

Deployment procedures take into account and resolve the reality that environments are often different, with each having complexities across different tiers with multiple dependencies. The situation is further compounded by existing operational practices. For example, in a typical data center, deployment procedures can involve a design time activity (typically performed by a lead administrator) and a runtime activity (typically performed by the operator).

Deployment procedures have been introduced in Grid Control 10.2.0.3 and are licensed under the Oracle Enterprise Manager Provisioning Pack.

10.1 Key Advantages of Deployment Procedures

The main advantage of deployment procedures lies in the fact that they can provide an extremely flexible framework for data center automation. While a vendor like Oracle often has specific best practice recommendations for patching and provisioning, the reality is that each data center has unique ways of achieving them. Deployment procedures are nothing more than a framework to achieve synergy between Oracle's out of box best practices and customers' own methods. Custom scripts can easily be plugged into deployment procedures for handling special tasks. The following properties of deployment procedures increase their value:

1. Extensible

   The objective of deployment procedures is to have as many best practice methods out of box as possible. In an ideal case the customer should be able to run the deployment procedures as-is against a set of targets. Oracle-shipped best practices deployment procedures cannot be modified. The customer can create a copy of the

Oracle shipped deployment procedure and modify the same to insert or delete steps and error handling modes.

2. Reusable

   Deployment procedures are reusable. The steps of the deployment procedure can be based against directives that are stored in the Software Library. The deployment procedures can also be exported and imported across environments. This implies that the deployment procedures once developed for a test environment need not be recreated for production environment.

3. Hot-pluggable

   The out-of-box deployment procedures are metadata driven so new sets of procedures can be added to the Oracle Enterprise Manager environment without any additional outage.

4. Automatable

   The runtime for all the deployment procedures can be automated using EMCLI and associated verbs, such as Oracle patching, OS patching and so forth. For more information on these verbs, see the *Oracle® Enterprise Manager Command Line Interface 10g Release 3 (10.2.0.3) for Windows or UNIX*.

10.1.1 Deployment Procedures Shipped In Oracle Enterprise Manager

In version 10.2.0.3 of Oracle Enterprise Manager, the out-of-box deployment procedures include the following:

- Application Server Deployment
- Oracle Clusterware/Oracle Real Applications Clusters (RAC) Provisioning
- Delete/Scale Down Oracle Real Applications Clusters
- One Click Extend Cluster Database
- Patch Oracle RAC Database -- All Nodes
- Patch Oracle RAC Database -- Rolling
- Patch Oracle Clusterware (Rolling Upgrade)
- Patch Oracle Database
- Patch Application Server
- Patch Solaris Hosts
- Patch Linux Hosts
- Patch Windows Hosts
- Patch Standalone Oracle Automatic Storage Management

10.2 Deployment Procedure Requirements

Deployment procedures can be used for targets managed by agents with versions of 10.2.0.2.0 or higher. Certain features in the Deployment Procedures Manager will not work with a 10.2.0.1 (or lower version) agent, even if the agent is uploading to a 10.2.0.2.0 Management Server.

Deployment procedures can be found under the Deployment Procedure Manager on the Deployments page. Choose a deployment procedure and click Run to execute it.

To use deployment procedures, you should first set up the Software Library before applying the 10.2.0.3 patchset. This ensures that the deployment procedures are installed out of the box. If you fail to set up the Software Library beforehand, you will have to set it up later and then manually deploy the files after the installation. For more information see Section 10.8.1, "Known Issues". For instructions on how to set up a software library, see Section 18.7, "Setting Up and Configuring a Software Library With Oracle Enterprise Manager".

10.3 Use-Cases for Deployment Procedures

The following sections provide examples of some use cases for deployment procedures.

10.3.1 Using Deployment Procedures to Apply Security-Related Critical Patch Updates to Oracle Databases

The out-of-box deployment procedure Patch Oracle Database can be used to apply critical patch updates (CPU) to several single instance databases simultaneously. In the following example, patch 5049080, a critical patch, is applied to 10.2.0.1 databases.

This involves the following steps:

1. Patch Download and Upload step (optional)

   Upload to Patch Cache is required only if there is no connection between the Management Server and Oracle Metalink. In 10.2.0.3 the Software Library is integrated with the Patch Cache. While uploading a patch to the Patch cache, you can choose the “Create Oracle Software Update Component”option. This will ensure that the patch is visible in the Software Library.

2. Runtime (Deployment time) steps:

   There are 3 inputs that must be provided at runtime - the patch number(s), the targets and the credentials. Finally the procedure must be scheduled for immediate or deferred execution.

   1. Select the procedure “XYZ Corp’s CPU patching procedure” and run it.

      You can choose to select the patch from Metalink or from the Software Library.

   2. Deployment Procedure by default will run the default SQL (catcpu.sql for CPUs) if present or you can provide custom SQL script to run.

   3. The targets are then chosen from a list of targets that are automatically populated on the screen based on the version of the product the patch applies to.

   4. Credentials follow. You can choose to use the ORACLE_HOME credentials from the repository.

   5. Finally the procedure is scheduled and submitted. The procedure can then be monitored by using “Procedure Completion Status” link. It can be retried from a failed step, if required.

10.3.2 Using Deployment Procedures for Single-Click Extend of Real Application Clusters

The out-of-box "Extend Real Application Clusters" deployment procedure can be used for extending an existing Real Application Cluster to one or many nodes. In the following example, an existing cluster is extended to one additional node.

Use the following steps:

1. Runtime (Deployment time) steps:
   a. Run the out-of-box Extend Real Application Clusters procedure.
   b. Choose the existing cluster to be extended and mention the details of the new node where the cluster will be extended.
   c. Mention the credential information for the new nodes and a schedule for extension operation. After you finish, click the Submit button to execute the deployment procedure.

   Once you submit the procedure, you can monitor it using the Procedure Completion Status link.

10.3.3 Using Deployment Procedures for Delete/Scale Down of Real Application Clusters

You can use the out-of-box Delete/Scale Down Real Application Clusters deployment procedure to delete or scale down an existing Real Application Cluster.

Use the following steps:

1. Runtime (Deployment) steps:
   a. Run the out-of-box Delete/Scale Down Real Application Clusters procedure.
   b. From the list of available nodes, select one, multiple, or all nodes for deletion from the cluster.
   c. Mention the credential information for the new nodes and a schedule for extension operation. When you finish, click Submit to execute the deployment procedure.

Once you submit the procedure, you can monitor it using the Procedure Completion Status feature.

10.4 Customizable Deployment Procedures

The out-of-box deployment procedures can be used as starting templates to create similar procedures (using the “Create Like” functionality), which can then be customized. You can edit the deployment procedure to insert or delete a step or a phase, or to enable or disable a step or a phase. Deployment procedures also allow different error handling methods depending upon the case. For example, in a patching operation where hosts are patched in parallel, it may be wise to simply skip the host on which a failure occurs. However, failure on a device creation could render the remaining provisioning operation ineffective. Therefore it may be necessary to abort the entire procedure for failure of such a step.

10.4.1 Customization Examples

The following sections describe three examples that illustrate how deployment procedures can be customized.

10.4.1.1 Insert Custom Step to Backup the Database Before Patching

A data center is notified by Grid Control that its Oracle Database installations are affected by Oracle's latest Critical Patch Update (CPU). The Security administrator studies the impact and hands it over to the lead DBA who first applies it to his test systems. In the process he wants to backup the database before applying the patch. He uses the Create Like feature of the out-of-box Oracle Database Patching Deployment procedure and inserts a custom step before the Apply Patch step, associating his script to take a backup, which he has uploaded to the Software Library. As a result, on the execution of the Deployment Procedure the backup of the database is performed each time before applying the patch.

10.4.1.2 Manual Step

XYZ Corporation has a process of ensuring that users are logged off from their application before the database is shutdown. The DBA checks with key users that they have indeed logged off before proceeding with the database shutdown. This can be achieved by introducing a manual step before the "Stop Database" step. The procedure would pause on the completion of the manual step. Only when the DBA chooses to continue would the procedure advance.

10.4.1.3 Application Service Shutdown and Startup Handling

Deployment procedures can be used to perform operations that are outside the scope of out-of-box procedures. Examples include stopping and starting an ERP application or registering a newly provisioned service with the load balancer. Each of these steps can run in the context of any valid operating system user and can make use of a Pluggable Authentication Module like "pbrun" (Powerbroker). They can also run in superuser mode using "sudo".

10.4.1.4 Set Notification for the Deployment Procedure Run

To receive notifications on the deployment procedure status during the design phase on a Create Like or the out-of-the-box deployment procedure, check the Enable Notifications check box and provide details for the Notification Tag Name, which you can use to associate to other Enterprise Manager Notification methods, and the status for which you would require notifications such as Success, Failure, or Action Required. The administrator is notified when the procedure based on the status selected for notification is run.

The above use case would require Configuration of Mail Server for Notification methods from Enterprise Manager to be working. For instructions about how to configure Notifications, see Section 13.1, "Setting Up Notifications".

To leverage the Enterprise Manager Notification methods, create a notification rule similar to "PAF Status Notification", which you can get to by navigating to the Preferences page and selecting Notification rule, and configure it such that the job name is prepended with the notification job tag that you specified for your procedures. Repeat the process to create notification rules to activate notification for other procedures associated with their notification tag name.

10.4.2 Importing or Exporting Deployment Procedures

In version 10.2.0.3 there is an out-of-box PARDeploy utility to export deployment procedures created by a user and distribute them across various Oracle Enterprise Manager deployments.

Provisioning Archive files can contain deployment procedures and/or components and directives from Software Library. Oracle provides PAR files that contain Oracle best-practices deployment procedures and the directives required to run them.

PARDeploy tool is located at $ORACLE_HOME/bin directory. Please be sure the $ORACLE_HOME environment variable is set to the oracle home of the OMS and the Software Library path is configured before you run PARDeploy. When you run $ORACLE_HOME/bin/PARDeploy, it will print out the following usage information.

Usage:

```
parFile <file> -force(optional)
```

```
PARDeploy -action <deploy|view> -parDir <dir> -force(optional)
```

```
PARDeploy -action export -guid <procedure guid> -file
```

```
<file> -displayName <name> -description <desc>
-metadataOnly(optional) PARDeploy -check
```

- *-force <force>* -- Force the Software Library entities to be created
- *-check <check>* -- Check to determine whether Software Library is configured
- *-file <file>* -- PAR file
- *-action <deploy | view | export>* -- Deploy, view, or export par file
- *displayName <displayName>* -- PAR file name
- *parDir <dir>* -- Directory where par files are located
- *metadataOnly <metadataOnly>* -- Flag for metadata-only exports
- *guid <guid>* -- Procedure GUID to export
- *parFile <file>* -- Path of PAR file
- *description <description>* -- PAR file description

PAR files provided by Oracle are located at `$ORACLE_HOME/sysman/prov/paf`. To deploy asprov.par in that directory, you can run the following command.

```
$ORACLE_HOME/bin/PARDeploy -action deploy -parFile $ORACLE_
HOME/sysman/prov/paf/asprov.par -force
```

You can also deploy all of the PAR files in that directory by running the following command:

```
$ORACLE_HOME/bin/PARDeploy -action deploy -parDir $ORACLE_
HOME/sysman/prov/paf/ -force
```

To create a PAR file that contains a specific procedure, use the export action. You must have the GUID of the deployment procedure, which you can acquire by logging into Enterprise Manager and moving to the Deployments tab. Then you can click on the Deployments procedure link in the Deployment Procedure Manager section. In the table, find the deployment procedure that you want to export and click on the name. In the procedure view page, note the URL address on the Navigation toolbar. The format should be similar to this:

```
http://<OMS
host>:<port>/em/console/paf/procedureView?guid=<value of GUID>
```

You will then need to use the GUID value to run PARDeploy tool. For example, if the GUID of the deployment procedure that you want to export is FAC05DD31E3791C3E030579D23106C67, then run the following command:

```
$ORACLE_HOME/bin/PARDeploy -action export -guid
FAC05DD31E3791C3E030579D23106C67 -file exportedDP.par
-displayName "User exported DP" -description "Deployment
Procedure to be copied to other OMS"
```

After you run this command, a new PAR file named exportedDP.par will be created in the directory where you run the command. Then you can import this PAR file to another OMS. You can use this PARDeploy tool to import or deploy it to an OMS, or you can also log in to the second Enterprise Manager, go to the Deployment Procedures main page and click on the Upload button to upload the PAR file.

Please note that when a procedure is exported using PARDeploy, any directives or components referred by the procedure will also be exported. However, only the latest revision of these directives or components will be exported. If you do not want to export components or directives, you can specify the metadataOnly flag when running PARDeploy.

10.5 Executing Deployment Procedures Using Pluggable Authentication Modules (PAM) and SUDO

Deployment procedures allow execution in sudo and other pluggable authentication modules (PAM) contexts. The sudo and the PAM command to be used throughout the procedure can be specified while editing the procedure. The individual directive steps can be chosen to run in normal mode, privileged mode (which requires sudo setup at the target for superuser execution for the command) or as a different user than the preferred credentials.

The use case is described below. The operating system users on the target box are shown below. User names in actual environments may vary.

- oracle: the DATABASE HOME owner. This is a locked account
- emd: The AGENT HOME owner.
- foo: The user who is performing the patching operation

Requirements for the users are:

- All three users belong to the same group, the OS DBA group.
- User foo can sudo to user oracle using foo's password

The goal in these examples is to patch the DATABASE HOME using foo's credentials.

10.5.1 Design Time Experience Example (One Time)

This process must be completed only once to customize the deployment procedure.

1. User selects the Oracle shipped Patch Oracle Database Deployment Procedure and does a Create Like for that deployment procedure and provides it a name.
2. User selects run-as sudo for all steps in the edited deployment procedure
3. User selects the sudo command to be run. In this example it may be "sudo –u oracle"
4. User gives the staging location for the directives in "Staging Area Path". The default is %oracle_spool%/EMStagedPatches/Others. Example: It could be %emd_root%/EMStagedPatches/Others This directory needs to have 770 permissions. This could also be any other directory which can be written onto by user "oracle".

5. User saves the modified deployment procedure (Patch Oracle Database with sudo).

10.5.2 Run Time Experience (Every Time)

This process must be completed for every run of the deployment procedure. User selects the modified patching deployment procedure (Patch Oracle Database with sudo) and runs it.

1. User selects the Database Updates to be patched.
2. User provides the location for staging the patch. The default for this is: User gives the PA Staged location as: %oracle_spool%/EMStagedPatches/Others Example: It could be %emd_root%/EMStagedPatches/Others This directory needs to have 770 permissions and it should exist before the DP is run. This could also be any other directory which can be written onto by user "oracle".
3. Now user gives foo's credentials (Patching User credentials).
4. Schedule and submit the job.

10.6 Deployment Procedure Variables

Oracle Enterprise Manager exposes several variables that can be used with deployment procedures. These variables can be used by Oracle customers to customize their specific tasks like startup and shutdown using their own directives.

Per ORACLE_HOME:

oraHome - Directory of the ORACLE_HOME

replacedStageDir - The absolute staging location of patches

instances - Selected database targets from the ORACLE_HOME

all_instances_home - All database targets running from the ORACLE_HOME

dbSIDs - All sids from the ORACLE_HOME

dbListeners - All listeners from the ORACLE_HOME

isPatchset - Yes/no of if patchset is being applied to the ORACLE_HOME

runRootScript - Yes/no indicating if rootscript needs to be run

Global:

stageDir - the staging directory to use provided like %oracle_home%....

patchIDs - list of patch ids selected

patchSrcs - indicating whether the patches came from Metalink or software lib

patchData – Uniform resource Names (URNs) of the patches

patchReleases - corresponding release of the patches

10.7 Using EMCLI to Execute Deployment Procedures

Oracle provides out-of-the-box templates for creating run time data for deployment procedures used in the most common use cases. These are known as *runtimedata templates*. You can access them under the `emcli/samples` directory in the OMS oracle home and you can then modify the configuration properties in these templates.

The process to use EMCLI to execute deployment procedures is depicted in Figure 10–1.

Figure 10–1 EMCLI Process to Execute Deployment Procedures

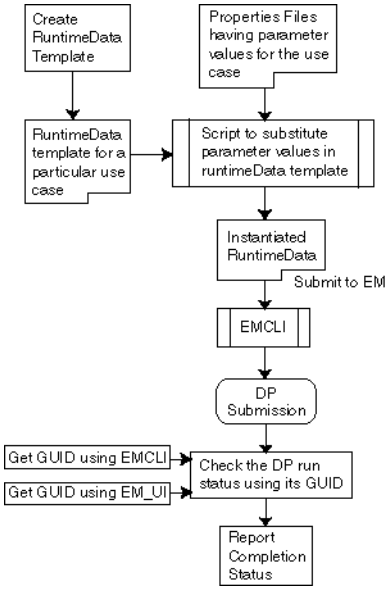

For each run, you can populate a simple properties file with specific information such as hosts, platform for deployment, and other required information.

The provided sample scripts use these two files as inputs and generate an instantiated runtime data response file containing all properties required for provisioning the system. This approach allows you to maintain one template and generate a small properties file for custom changes, thus ensuring that you do not have to be concerned about XML format changes or changes to variables due to customizations created by the system administrator for that data center.

Set up EMCLI on your system and submit the procedure with this instantiated runtime data as the input, along with the procedure to be executed. EMCLI provides an instanceGUID for the procedure being executed. You can monitor the status of this instance using its instanceGUI from CLI or from the Oracle Enterprise Manager user interface.

10.7.1 Creating the RuntimeData Templates

If you have tested and are satisfied with the provisioning procedures provided by Oracle, you can skip this section. If you have customized the procedures to your environment, you can use these steps to create templates for your system.

10.7.1.1 Step 1: Locate the Successful Test Run of Your Customized Procedure

Each procedure execution is associated with a unique instanceGUID.

1. Log in to Enterprise Manager.
2. Go to the Deployments tab and choose Procedure Completion Status.

   You can access the RAC procedure completion status by clicking on the RAC Provisioning Procedures and then clicking on Procedure Completion Status tab.

You can access the AS procedure completion status by clicking on the Application Server Provisioning Procedures and then clicking on Procedure Completion Status.

You can access the Patching procedure completion status by clicking on the Patch procedure and then clicking on the Procedure Completion Status tab.

3. Retrieve the instanceGUID for the execution.

   The link takes you to the procedure steps. The URL includes the instanceGUID. You can acquire the same information from the following EMCLI command:

   emcli get_intances -type={procedure type}

   Usage: ./emcli get_instances -type={procedure type}
   Description: Displays a list of procedure instances
   Option: -type={procedure type}
   Output Columns: GUID, Procedure Type, Instance Name, Status
   See Also: emcli help get_procedure_types

10.7.1.2 Step 2: Download the Runtime Data For This Run

Use EMCLI to download the runtime data associated with this instance. Use the following command:

emcli get_instance_data_xml -instance={instance_guid}

Usage: ./emcli get_instance_data_xml -instance={instance_guid}
Description: Download Instance Data XML
Option: -instance
Option Description: The instance of the GUID.
Example: emcli get_instance_data_xml -instance=16B15Cb29C3F9E6CE040578C
Output: The Instance Data XML

10.7.1.3 Step 3: Locate and Create Instantiation Variables

For various use cases, you must have the XML file for the RuntimeData containing the parameter values for the specific deployment procedure run. This must be in a format where the parameter values, which are passed at runtime, are distinctly identified by appropriate placeholders.

Below is an example of a RuntimeData template fie and an explanation of its syntax.

```
<?xml version = '1.0' encoding = 'UTF-8'?>
<runtimeData xmlns="http://www.oracle.com/paf/runtimeDataXML">
 <elements>
        <scalar value="^rac_partial_delete_node_PUB_NODELIST_ALL_NODES^"
        classname="java.lang.String" name="PUB_NODELIST_ALL_NODES"/>
       <scalar value="^rac_partial_delete_node_NODE_LIST_FOR_DELETION^"
       classname="java.lang.String" name="NODE_LIST_FOR_DELETION"/>
       <scalar value="^rac_partial_delete_node_HOST_LIST_FOR_DELETION^"
       classname="java.lang.String" name="HOST_LIST_FOR_DELETION"/>
       <scalar value="^rac_partial_delete_node_REMAINING_NODE_LIST^"
       classname="java.lang.String" name="REMAINING_NODE_LIST"/>
       .......
       .......
       .......
  </elements>
</runtimeData>
```

In the example above, the actual values have been replaced with special placeholder variables that will be substituted by the script with values from the properties file. The syntax of the placeholder variable is seen here:

^<Usecase_name>_<Procedure_variable_name>^

You are can choose any variable name, however, you should refer to this best practice naming convention to capture the use case name adn the variable name. Since the variable is contained in the properties file for runtime deployment, it serves as a self-explanatory indicator of what should be filled in.

In the example above, the use case name is rac_partial_delete_node for scaling down RAC clusters. The surrounding carets (^) are a requirement of parsing logic which is used to substitute values at runtime.

The template should now be ready for use at runtime.

10.7.2 Using the RuntimeData XML Instead Of RuntimeData Templates

Although the previous discussion describes a very generic approach for submitting any deployment procedure using EMCLI where the inputs are a Properties file and Runtime Data template file, there is an alternative approach that uses the same generic script to submit deployment procedures for different use cases using a property file as the only input parameter. No Runtime Data template file is required. Instead, the Runtime data XML file of a previous deployment procedure run is used after downloading it using EMCLI (see Step 1: Locate the Successful Test Run of Your Customized Procedure).

In this approach you can replace the old runtime values in the runtime XML of the previous run with the new runtime values, which are relevant to the new run for the different use case. For example, if you have performed a successful partial cluster scale-down, suppose your deployment runtime XMLfor the run contains the following syntax:

```
<scalar value="node1" classname="java.lang.String" name="NODE_
LIST_FOR_DELETION"/>
```

On a different occasion, for performing a complete cluster deletion, you might need to have the variable value as:

```
<scalar value="node1,node2" classname="java.lang.String"
name="NODE_LIST_FOR_DELETION"/>
```

For this you can use a properties file of the form:

<old_value>=<new_value>

In this specific example, the properties file will have:

```
node1=node1,node2
```

The script will use the old runtime data XML from the pervious run along with the properties file and it will generate the new runtime data XML file (name will be passed as a parameter) and this can then be used to submit the deployment procedure using EMCLI for the new usecase.

While this second approach would considerably expedite the process, a word of caution is worth mentioning. Since this approach would simply replace an old string with a new string, you might run into issues if the old string is a substring in multiple strings in the deployment procedure runtime XML. In that case the resulting runtime XML might be erroneous. To circumvent this issue, you should format the properties file properly. The rule-of-thumb is to put the specifics before the generics. In

particular, while specifying node names you should put the FQDNs before the simple node names. A fragment of a properties file in the form of old-value=new-value pairs is shown below and illustrates this point.

```
node2.test.com=node3.example.com
node1.test.com=node4.example.com
node2,node1=node3,node4
node2=node3
node1=node4
...
```

10.7.3 Creating a Properties File for Provisioning Procedures

The properties file contains values for the placeholder variables and are essentially in the form of a group of name-value pairs. In the RAC scale-down procedure, the properties file would appear as shown below:

rac_partial_delete_node _PUB_NODELIST_ALL_NODES=node2,node1
rac_partial_delete_node _NODE_LIST_FOR_DELETION=node2
rac_partial_delete_node _HOST_LIST_FOR_DELETION=node2.example.com
rac_partial_delete_node _REMAINING_NODE_LIST=node1

The following command may be used to create the properties file. Run this with the runtime data template XML containing the placeholder variables, and the property name as parameters. It will create the property file name in the current working directory:

grep -e "\^.\^" <runtimedata template xml file> | awk -F"^"'{print $2"="}' | sort | uniq > <prop file>*

For example, the command could look like this:

```
grep -e "\^.*\^" runtimedata_sh.xml | awk -F"^"'{print $2"="}' |
sort | uniq > runtime.properties
```

You can then populate the properties file with appropriate inputs. An alternative approach using the runtime data XML instead of creating templates as described earlier, you could have a set of old-value=new-value pairs in the properties file as seen below:

```
OraCrs10g_home=OraCrs10g
OraASM10g_home1=OraASM10g
OraRAC10g_home1=OraRAC10g
node2.test.com=stamy02.example.com
node1.test.com=stamy01.example.com
node2,node1=stamy02,stamy01
node2=stamy02
node1=stamy01
```

10.7.3.1 Properties File Variables for Provisioning Procedures

The following table shows a list of properties file variables used for Extend Cluster with CRS, RAC, and Agent Push.

Table 10–1 Properties File Variables for Provisioning Procedures

Variable Name	Description	Example
PUB_NODELIST_ALL_NODES	List of nodes	Node1, Node2

Table 10–1 (Cont.) Properties File Variables for Provisioning Procedures

Variable Name	Description	Example
CLUSTERWARE_ORACLE_HOME_LOC	Clusterware Home location	/u01/app/oracle/crs
CLUSTER_ORACLE_HOME_NAME	Clusterware Home name	OraCrs10g_home
ORADATA_LOC	Shared Oradata Location	/oradbshare/oradata/linorcl/temp01.dbf
OCR_PARTITION_VALUE	OCR File Location	/oradbshare/crs/ocr/oradbshare/crs/ocr
VDSK_PARTITION_VALUE	Vdisk File Location	/oradbshare/crs/vdsk
PRIV_NODELIST_VALUE	Private Node Names	node1-priv, node2-priv
PRIV_IP_LIST_VALUE	Private IPs	10.35.69.70,10.2.0.2
VIRTUAL_NODELIST_VALUE	Virtual Node Names	node1-vip, node2-vip
VIRTUAL_IP_LIST_VALUE	Virtual IPs	130.35.69.135,136.34.56.176
DB_SID	Global Database Name	.LINORCL
DB_ORACLE_HOME_LOC	RAC Database Home Location	/u01/app/oracle/product/10.2.0/rac
DB_ORACLE_HOME_NAME	RAC Database Home Name	OraDb10g_home1
DEVICES	List of shared disks	/dev/sda/1%120%ocfs%/oradbshare%ocfs
NICS	Network Interface	eth0:130.35.68.0:PUBLIC,eth1:10.35.68.0:PRIVATE
CLUSTER_NAME	Clusterware Name	CRS
WORK_DIR_LOC	Temporary Work Directory	/temp
REF_NODENAME	Reference Node on Cluster From Where important actions will be triggered	node1.example.com
LISTENER_ORACLE_HOME	Home location of RAC DB Listener	/u01/app/oracle/product/10.2.0/rac
DB_TARGET	Database target	LINORCL
PUB_NODELIST_ADD_INSTANCE_NODES	Node to be added to the existing cluster	node2
EM_REP_USER_NAME	Database user with SYS credentials	SYSMAN

Table 10–1 (Cont.) Properties File Variables for Provisioning Procedures

Variable Name	Description	Example
EXT_NODENAME	Host name for the node to be added to the existing cluster	node2.example.com
CRS_NODE_NIC_LIST	Node name, private name, and virtual name for the node being added to the cluster	node2:node2-priv:node2-vip
REF_USERNAME	Reference Host Credentials	oracle
REF_PASSWORD	Reference Host Credentials	oracle
EXT_NODENAME_username	Host credentials of target being included in the cluster	oracle
EXT_NODENAME_password	Host credentials of target being included in the cluster	oracle
CLUSTER_NODE_VIP_NAME	Virtual Name	node2-vip
CLUSTER_NODE_PRIVATE_NAME	Private Name	node2-priv
PRIV_NODELIST_PREREQS	NodeLists	node1-priv,node2-priv,node1-vip
VIRTUAL_IP_LIST_PREREQS	Virtual IPs	136.34.56.176
PRIV_IP_LIST_PREREQS	Private IPs	10.35.69.70,10.2.0.2,130.35.69.135

10.7.3.2 Properties File Variables for Patching Procedures

The following table lists the variables that are present in the runtime template XMLs that are populated with a matching variable from the properties file. A '?' is used to represent the number of occurrences of the variable in the template.

Table 10–2 Properties File Variables for Patching Procedures

Variable Name	Description	Example
PA_VAR_targetName_?	Host name of the oracle homes where the DB instance resides. This variable must be repeated for each unique host. Please note the targetName under 'homeBeingPatched' is same as the name of the containing host.	Host1.us.oracle.com

Table 10–2 (Cont.) Properties File Variables for Patching Procedures

Variable Name	Description	Example
PA_VAR_oraHome_?	Oracle Home location to be patched. For each home in the host this variable is repeated. Value of 'credContainerLoc' should have same value as oraHome.	/scratch/OracleHomes/h1
PA_VAR_instance_?_?	Selected list of instances (for example, targets) to be patched in the home.	Target1
PA_VAR_patchID_?	Patch ID of the patch to be installed.	123456
PA_VAR_patchData_?	URN of the patch if it is in the software library. Else it is of the form pPatchID_ release_platform_ productid.	p4926771_9.2.0.6_46_9480 represents patch 49226771 of release 9.2.0.6 on LINUX for Oracle Database.
PA_VAR_patchSrcs_?	Either specified as SWLIB if patch is in software library, or as METALINK.	METALINK
PA_VAR_patchTypes_?	Type of patch. Ex: Patch, PatchSet.	Patch
PA_VAR_patchReleases_?	Version of the patch	10.2.0.3
PA_VAR_patchPlatforms_?	Platform of the patch	LINUX
PA_VAR_stageDir	Location where the patches are staged.	%oracle_ home%EMStagedPatches
PA_VAR_isPatchset	True if Patch is a Patchset, False otherwise.	false
PA_VAR_isNotPatchset	True if the patch is a Patch, False otherwise.	true
PA_VAR_sqlScript	Location of the SQL script to be run. Set to "defaultSqlScript", if the SQL script inside the patch must be run and there is not custom SQL script.	defaultSqlScript
PA_VAR_isUpgradeStepEnabled	True of OPatch in the Oracle Home must be upgraded.	true

The elements 'hostsBeingPatched' and 'homesBeingPatched' do not have to be populated from the properties files. However, variables inside those elements are being populated. The element 'hostsBeingPatched' is the target list declaration that contains the hosts where patching will be done while 'homesBeingPatched' is the inner target list declaration that represents the list of homes within the host.

10.7.3.3 Updating the Runtime Data Template to Fit Various Patching Use Cases

Based on your requirements of patching in multiple hosts and homes, the number of child elements under 'hostsBeingPatched' and 'homesBeingPatched' must be modified in the runtime template data. The same is true of other variables such as PA_VAR_ patchId and PA_VAR_patchSrcs.

Using the new runtime data template, the properties file can be generated using the steps described in Section 10.7.4, "Creating the Properties File for Patching Procedures".

10.7.3.4 Getting the Host and Homes From the Instances For Patching Runtime Data

Use the following query to acquire the host name to which the instances belong:

Select host_name from mgmt_targets where target_name = '<instance_name>'

Use the following query to acquire the home location out of which the instance is running:

Select p.property_value from mgmt_target_properties p, mgmt_targets t where p.target_guid = t.target_guid and p.property_name = 'OracleHome' and t.target_name = '<instance_name>'

In the above syntax, a RAC instance represents the RAC instance target, a Standalone instance represents a database target, and a CRS (Clusterware) instance represents the target corresponding to the CRS node, which in most cases is the same as the host name.

10.7.4 Creating the Properties File for Patching Procedures

The method for creating properties files for patching procedures is identical to the method used for provisioning procedures. The same grep command can be used to generate the properties file by using the template runtime XML for standalone database patching. For the alternative approach using the runtime data XML instead of creating templates, the properties file sample is shown below:

```
PA_VAR_targetName_1=host1.us.oracle.com
PA_VAR_oraHome_1=/scratch/home1
PA_VAR_instances_1_1=target11
PA_VAR_instances_1_2=target12
PA_VAR_targetName_2=host2.us.oracle.com
PA_VAR_oraHome_2=/scratch/home2
PA_VAR_instances_2_1=target21
```

10.7.5 Submitting the Deployment Procedure

After the EMCLI has been set up, follow these steps to submit the deployment procedure:

1. Instantiate the Runtime data

   perl instantiate_runtimedata.pl
   <source RuntimeData template file name>
   <destination RuntimeData file name>
   <Properties file name>

   The destination RuntimeData file name is the file that is generated and contains the substituted property values and is used for the deployment procedure execution. You can pass any desired file name for this parameter. The perl script can be used to generate the runtime data for both provisioning and patching procedures.

2. Retrieve the appropriate procedure GUID

   The GUID associated with a specific procedure can be retrieved using EMCLI or the Enterprise Manager User Interface. To obtain the GUID of a deployment

procedure, click on its link. The View Procedure page displays showing the GUID as the last portion of the URL. Alternatively, you use the following EMCLI commands to retrieve the procedure types and the procedure GUID.

./emcli get_procedure_types
./emcli get_procedures -type=(procedure type)

3. Submit the deployment procedure for execution.

   To submit the deployment procedure for execution, run the following command:

   emcli submit_procedure
   -procedure=<DP GUID>
   -input_file=data:<destination RuntimeData file name>
   [-schedule="<time>"]

   Usage:

   emcli submit_procedure
   -procedure="GUID of the procedure"
   -input_file="data:{file_path}"
   -instance_name="name of the procedure instance"
   -schedule=
   start_time:yyyy/MM/dd HH:mm;
   tz:{java timezone ID}];

   Description: Submit a deployment procedure

   Example:

```
emcli submit_procedure -input_file=data:data.xml -
procedure=16B15Cb29C3F9E6CE040578C -
-schedule="start_time:2007/6/21 21:23;tz:America/New_York"
```

10.8 Known Issues and Troubleshooting

The following section discusses known issues surrounding deployment procedures and describes how to troubleshoot problems that may arise when using deployment procedures.

10.8.1 Known Issues

If you upgrade an existing 10.2.0.1 or 10.2.0.2 version of Enterprise Manager to version 10.2.0.3, then you must manually re-upload any existing shiphomes for RAC or Application Server procedures.

10.8.2 Troubleshooting

If a specific run of deployment procedures fails, then the following log files can provide insight into the reason of the failure. Correct the reason for failure and rerun the deployment procedure:

- <OMS_ORACLE_HOME>sysman/log/pafLogs/<log file for this session>
- <OMS_ORACLE_HOME>sysman/log/emoms.trc
- <OMS_ORACLE_HOME>sysman/log/emoms.log

For faster resolution of any deployment procedure-related issues, plan to provide these files to support when you create a support request.

11

Sizing and Maximizing the Performance of Oracle Enterprise Manager

Oracle Enterprise Manager 10*g* 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

11.1 Oracle Enterprise Manager Grid Control Architecture Overview

The architecture for Oracle Enterprise Manager 10*g* 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 11–1 Overview of Enterprise Manager Architecture Components

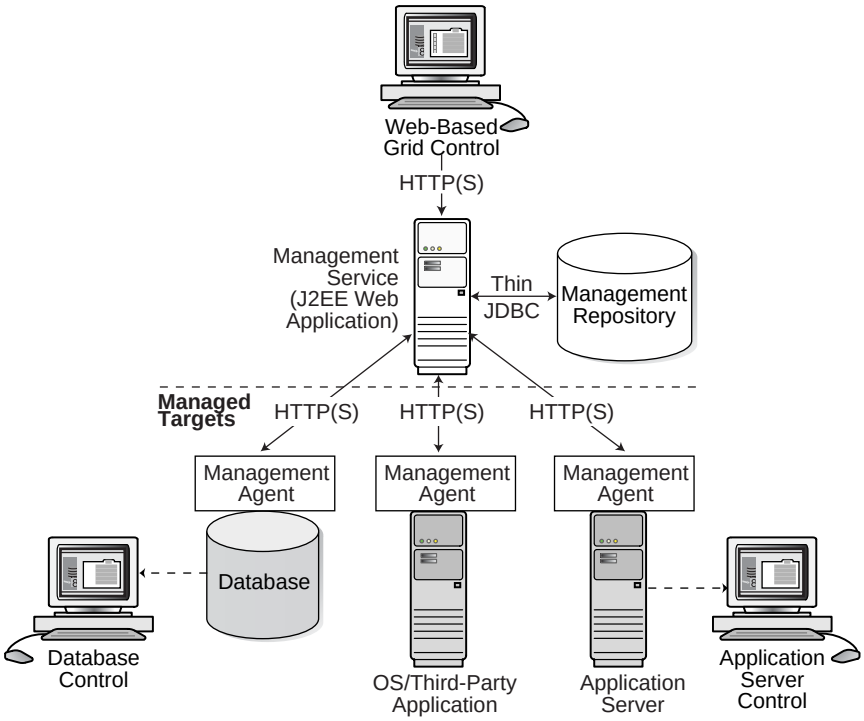

For more information about the Grid Control architecture, see the Oracle Enterprise Manager 10*g* documentation:

- *Oracle Enterprise Manager Grid Control Installation and Basic Configuration*
- *Oracle Enterprise Manager Concepts*

The Oracle Enterprise Manager 10*g* documentation is available at the following location on the Oracle Technology Network (OTN):

`http://otn.oracle.com/documentation/oem.html`

11.2 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 10*g*, choose a rough starting host configuration as listed in Table 11–1.
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.

11.2.1 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 11–1 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 11–2 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 11–3 Total Management Repository Storage

Deployment Size	Minimum Tablespace Sizes*			
	SYSTEM**	MGMT_ TABLESPACE	MGMT_ECM_ DEPOT_TS	TEMP
Small	600 MB	2 GB	1 GB	1 GB
Medium	600 MB	20 GB	1 GB	2 GB
Large	600 MB	200 GB	2 GB	4 GB

*These are strictly minimum values and are intended as rough guidelines only. The actual size of the MGMT_ TABLESPACE could vary widely from deployment to deployment due to variations in target type distribution, user customization, and several other factors. These tablespaces are defined with AUTOEXTEND set to ON by default to help mitigate space constraint issues. On raw file systems Oracle recommends using more than the minimum size to help prevent space constraint issues.
**The SYSTEM and TEMP tablespace sizes are minimums for Enterprise Manager only repositories. If Enterprise Manager is sharing the repository database with other application(s), these minimums may be too low.

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.

11.2.1.1 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 11–2 Network Links Related to Enterprise Manager Components

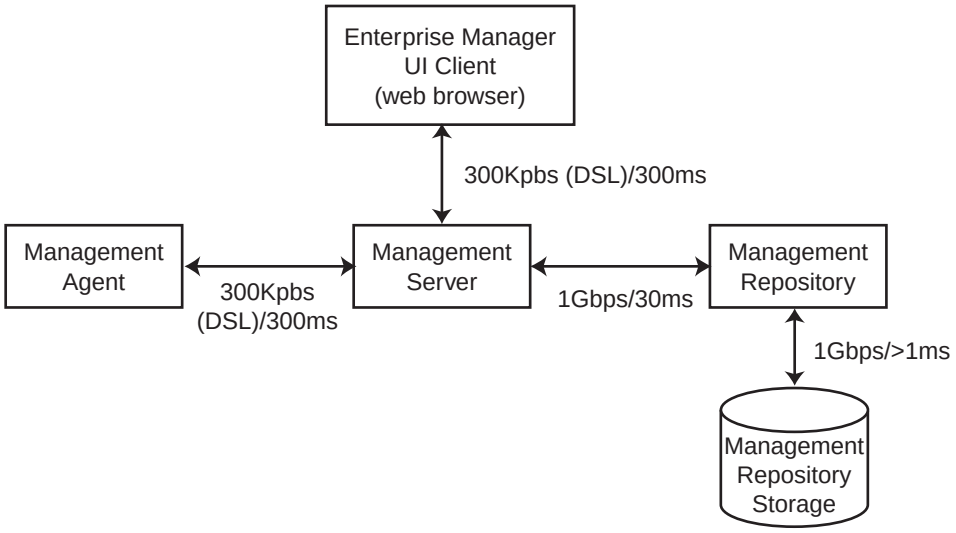

You can see in Figure 11–2 that the bandwidth and latency minimum requirements of network links between Enterprise Manager Grid Control components greatly impact the performance of the Enterprise Manager application.

11.2.2 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.

11.2.3 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.

11.2.3.1 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 10*g* 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);`

11.2.3.2 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_VIOLATION');`

  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_VIOLATION

– MGMT_CURRENT_VIOLATION

– MGMT_SYSTEM_ERROR_LOG

– MGMT_SYSTEM_PERFORMANCE_LOG

– MGMT_METRIC_ERRORS

– MGMT_CURRENT_METRIC_ERRORS

– MGMT_STRING_METRIC_HISTORY

– MGMT_JOB_OUTPUT

These Management Repository tables can be prone to uneven growth patterns and can have large areas of allocated space that are unused due to infrequent, large spikes in data volume.

The best tool for determining whether reorganization is required is the DBMS_SPACE.SPACE_USAGE procedure. The procedure takes advantage of space bitmap structures in the locally managed tablespace MGMT_TABLESPACE to determine the fullness of blocks for segments. The following is an example script for this procedure:

```
variable unformatted_blocks number;
variable unformatted_bytes number;
variable freeBlocks_0_25 number;
variable freeBytes_0_25 number;
variable freeBlocks_25_50 number:
variable freeBytes_25_50 number;
variable freeBlocks_50_75 number;
variable freeBytes_50_75 number;
variable freeBlocks_75_100 number;
variable freeBytes_75_100 number;
variable full_blocks number;
variable full_bytes number;

begin
dbms_space.space_usage('SYSMAN','MGMT_JOB_OUTPUT','TABLE',
      :unformatted_blocks,:unformatted_bytes,
      :freeBlocks_0_25,:freeBytes_0_25,
      :freeBlocks_25_50,:freeBytes_25_50,
      :freeBlocks_50_75,:freeBytes_50_75,
      :freeBlocks_75_100,:freeBytes_75_100,
      :full_blocks,:full_bytes);
end;
/
print unformatted_blocks;
print freeBlocks_0_25;
print freeBlocks_25_50;
```

```
print freeBlocks_50_75;
print freeBlocks_75_100;
print full_blocks;
```

Beginning in version 10.2 of the Oracle RDBMS, you can take advantage of the Segment Advisor functionality to automatically examine SYSMAN segments. Segment Advisor identifies candidates for rebuild and then generates the rebuild statements.

DBMS_SPACE and Segment Advisor provide recommendations and results you can use as a guide in determing an effective course of action. However, there are times when these recommendations can and should be overridden. For example, if you know that segment free space is likely to be reused, do not reorganize a table that may have been identified for reorganization by one of the tools. It is important to keep in mind that these tools only examine the current state of space allocation for a segment. They cannot provide insight into the projected data volume of the Enterprise Manager Grid Control application at your site.

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.

11.2.4 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.

11.2.4.1 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.

11.2.4.2 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.

11.2.4.3 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.

11.2.4.4 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.

11.2.4.5 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.

11.2.4.6 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 Manger 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 11.2.3. If you stay current, you should not need to analyze UI page performance as often, if at all.

11.2.5 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.

11.3 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.

11.3.1 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.

11.3.1.1 Oracle Management Service

As the Management Service is stateless, the task is to restore the binaries and configuration files is 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.

11.3.1.2 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.

11.3.2 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.

11.3.2.1 Management Repository

If restoring the Management Repository to a new host, restore a backup of the database and modify the emoms.properities 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.properities` 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)))
```

11.3.2.2 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.

11.3.2.3 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.

12

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

12.1 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

12.1.1 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.

12.1.2 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.

12.1.3 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 12–1.
4. Save your changes and exit the file.
5. Restart the Management Agent.

Table 12–1 Properties for Controlling the Disk Space Used by the Management Agent

Property	Explanation
UploadMaxBytesXML	Use this property in the `emd.properties` 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 `emd.properties` 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 `UploadMaxDiskUsedPctFloor` property in the `emd.properties` file.

12.1.4 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
```

12.1.5 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

12.1.5.1 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
```

12.1.5.2 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.

12.1.5.3 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/supportedtzs.lst and pick a timezone region with a
standard offset of +5:0 from GMT and update the property 'agentTZRegion' in the
file: ORACLE_HOME/sysman/config/emd.properties
```

This error appears in one of the log files shown in Table 12–2, depending upon which Enterprise Manager product you are using.

Table 12–2 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 8-1 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.

12.1.5.4 Troubleshooting Management Service Time Zone Problems

Section 12.1.5.3 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/nsupportedtzs.lst, update the property
'agentTZRegion' in the file
ORACLE_HOME/sysman/config/emd.properties and restart the agent.
A value which is around an offset of  -05:00 from GMT should be picked.
```

This error appears in one of the log files shown in Table 12–2, 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 `timezlrg.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 12-5 for instructions on changing the time zone assigned to the Management Agent

12.1.6 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
   ```

12.2 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

12.2.1 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.

12.2.1.1 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 12–3.

   Example 12–1 shows sample entries in the `emoms.properties` file.

4. Restart the Management Service.

Table 12–3 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 12-9 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 Services" on page 3-14
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 12–1 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
```

12.2.1.2 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
   ```

12.2.2 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 12-1

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
   ```

12.2.3 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 `emos.properties` file.
5. Restart the Management Service.

13

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

13.1 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.

13.1.1 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 13–1). 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 13–1 shows sample notification method entries.

Example 13–1 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 13–1 Defining a Mail Server

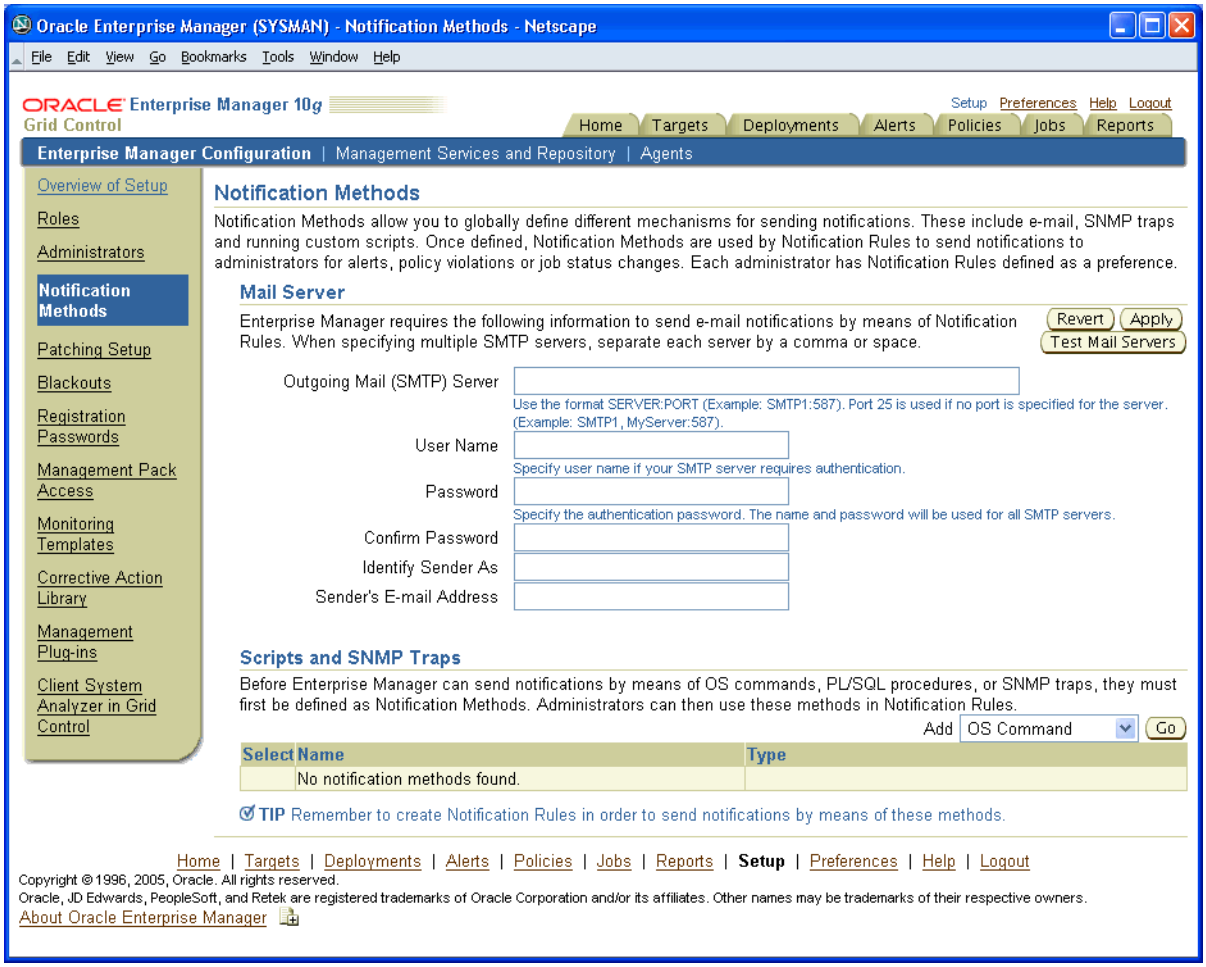

Note: The e-mail address you specify on this page is not the e-mail address to which the notification is sent. You will have to specify the e-mail address (where notifications will be sent) from the Preferences General page.

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 13–2 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 13–2, 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.

13.1.2 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.

13.1.2.1 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 13–3 shows a typical notification that uses the long format.

   The Short Format (Example 13–4) 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 13–3 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 13–4 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
```

13.1.2.2 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.

13.1.2.3 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 13–1 lists all the default notification rules. These are all owned by the SYSMAN user and are public rules.

Table 13–1 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 13–1 (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 13–1 (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.

13.1.3 Setting Up E-mail for Other Administrators

eIf 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 who's 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 desire 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.

13.2 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.

13.2.1 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.

13.2.1.1 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 13-13
- "Passing Corrective Action Execution Status to an OS Command or Script" on page 13-21
- "Passing Job Execution Status to an OS Command or Script" on page 13-26

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 13–5 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 13–5 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 13-11.

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 13–6 shows information required for the notification method.

Example 13–6 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 13-28 or "Assigning Methods to Rules" on page 13-27.

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 13–2 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 13–2 (Cont.) Environment Variables

Environment Variable	Description
SEVERITY	Type of severity. For example, severity for a target's (availability) status metric are: ▪ UP ▪ DOWN ▪ UNREACHABLE CLEAR ▪ UNREACHABLE START ▪ BLACKOUT END ▪ BLACKOUT START Other metrics can have any of the following severities: ▪ 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 13–7 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 13–7 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 13–8 shows an OS script that logs alert information to the file 'alertmsg.txt'. The file is saved to the /u1/results directory.

Example 13–8 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 13–9 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 13–9 HP OpenView Script

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

13.2.1.2 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 13-17

"Passing Corrective Action Status Change Information" on page 13-21

"Passing Job Execution Status Information" on page 13-24

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 13–11, "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 13-15.

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 13–10.

Example 13–10 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 13-28 or "Assigning Methods to Rules" on page 13-27.

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 13–3 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 13–3 (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 13–4 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 13–4 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 13–11 PL/SQL Procedure Using a Severity Code

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

PROCEDURE LOG_CRITICAL_ALERTS(severity IN MGMT_NOTIFY_SEVERITY)
IS
BEGIN
```

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

13.2.1.3 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 13-29.

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 13–12, the SNMP host will receive your SNMP traps.

Example 13–12 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 13–13. Each piece of information is sent as a variable embedded in the SNMP Trap.

Example 13–13 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 13-28 or "Assigning Methods to Rules" on page 13-27.

13.3 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 many want to call an OS script to open a trouble ticket for an in-house support trouble ticket system if a critical corrective action fails to run. In this case you will want to pass status (Problems, Aborted, etc.) to the script to open a trouble ticket and escalate the problem.

13.3.1 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 13–5 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_NAMEn	The name of the *n*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 13–5 (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: ▪ UP ▪ DOWN ▪ UNREACHABLE CLEAR ▪ UNREACHABLE START ▪ BLACKOUT END ▪ BLACKOUT START Other metrics can have any of the following severities: ▪ 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.

13.3.2 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 13–6 lists all corrective action status change attributes that can be passed:

Table 13–6 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 13–7, " 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 13–6 (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 13–4, " 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 13–7 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 13–7 (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 13–14 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;
```

13.4 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 many want to call an OS script to open a trouble ticket for an in-house support trouble ticket system if a

critical job fails to run. In this case you will want to pass status (Problems, Aborted, etc.) to the script to open a trouble ticket and escalate the problem.

13.4.1 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 13–8 lists all corrective action status change attributes that can be passed:

Table 13–8 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 13–9 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 13–9 (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 13–15 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;
```

13.4.2 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 13–10 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_NAMEn	The name of the *n*th target. For example, TARGET_NAME1, TARGET_NAME2.
TARGET_TYPEn	The type of the *n*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.

13.5 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 13–2 Assigning Methods to Rules

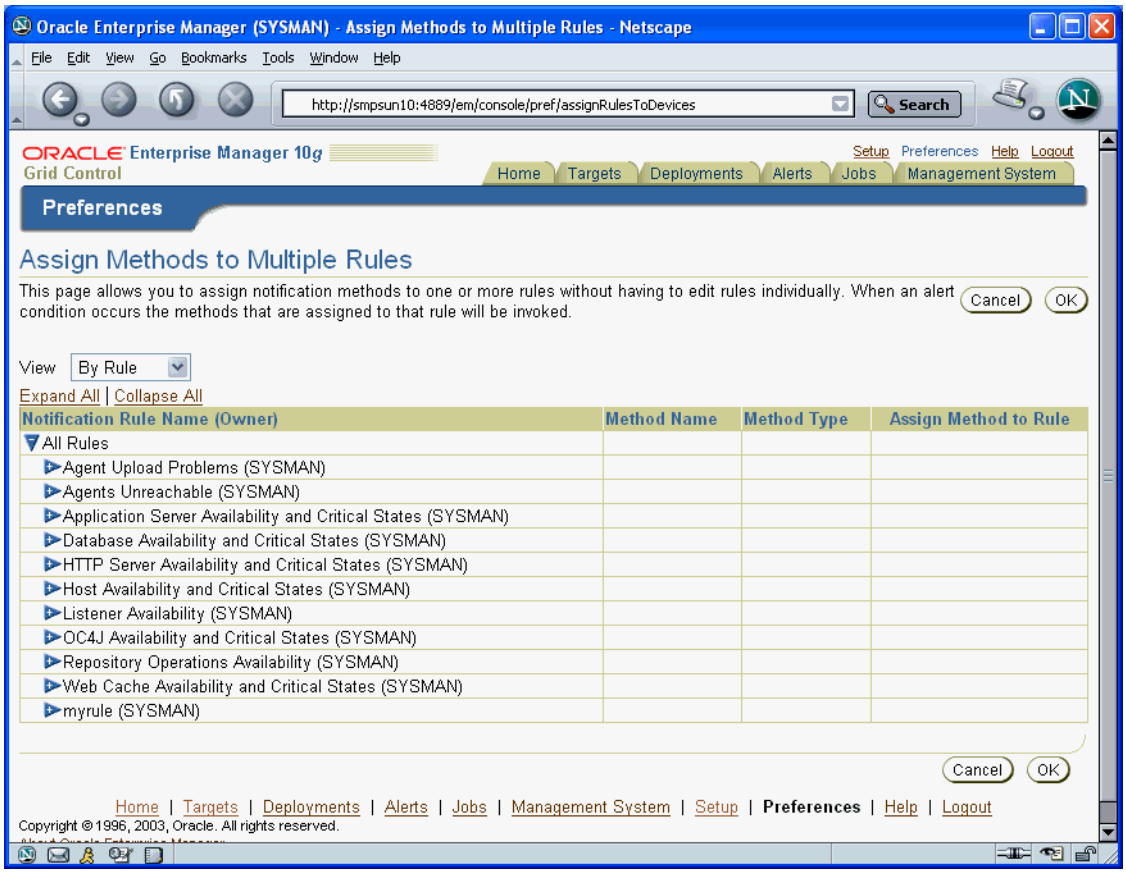

13.6 Assigning Rules to Methods

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 13–3 Assign Rules to Methods

13.7 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.

13.7.1 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.

13.7.2 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.

13.7.2.1 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.

13.7.2.2 MIB Definition

Example 13–16 shows a typical MIB definition used by Enterprise Manager.

Example 13–16 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
```

13.8 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.

13.8.1 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.

13.8.2 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.

13.8.3 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
```

13.8.4 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: rgmemeasmtp.oraclecorp.com: Sending failed;
nested exception is:
javax.mail.MessagingException: 550 5.7.1 <smpemailtest_ie@oracle.com>... Access
denied
```

Always use the Test button to make sure the e-mail gateway configuration is valid. Check that an e-mail is received at the sender's e-mail address

13.8.5 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 (stacb10.us.oracle.com_Management_Service) (SYSMAN,
myrule )
```

No execute permission on executable:

```
Error calling /bin/myscript: java.io.IOException: /bin/myscript: cannot execute
(stacb10.us.oracle.com_Management_Service) (SYSMAN, myrule )
```

Timeout because OS Command ran too long:

```
Timeout occurred running /bin/myscript (stacb10.us.oracle.com_Management_Service)
(SYSMAN, myrule )
```

Any errors such as out of memory or too many processes running on OMS machine will be logged as appropriate.

Always use the Test button to make sure OS Command configuration is valid.

13.8.6 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.

13.8.7 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 plsql_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 plsql_proc: ORA-06553: PLS-306: wrong number or
types of arguments in call to 'PLSQL_PROC' (SYSMAN, myrule)
```

Procedure has bug and is raising an exception.

```
Error calling PL/SQL procedure plsql_proc: ORA-06531: Reference to uninitialized
collection (SYSMAN, myrule)
```

Care should be taken to avoid leaking cursors in your PL/SQL. Any exception due to this condition will result in delivery failure with the message being displayed in the Details section of the alert in the Grid Control console.

Always use the Test button to make sure the PL/SQL configuration is valid.

14

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

14.1 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.

14.2 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

14.2.1 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

14.2.1.1 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 14–1 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.

14.2.1.2 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 14–2 Script Output Information Tags

Tag	Definition
em_result	Use this tag to return script result values. Exactly one em_result tag must be found in STDOUT. If more than one em_result tag is found, the first tag encountered will be used; subsequent em_result tags will be ignored. Example: `print "em_result=200\n"` Returns 200 as the value of the monitored object.

Table 14–2 (Cont.) (Cont.) Script Output Information Tags

Tag	Definition

em_message	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.
	Example:
	`print "em_result=200\nem_message=Disk usage is high\n"`
	Returns 200 as the value of the monitored object in addition to the message "Disk usage is high".
	If you want to include the value of em_result in the message, you can use the placeholder $em_result.
	Example:
	`print "em_message=Disk usage is at $em_result.\n"`
	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:
	`em_message=The value is $em_result`
	Example:
	`print "em_result=200\n"`
	Returns 200 as the value of the monitored object and the generated message "The value is 200"

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.

14.2.1.3 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 14-5. 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.

14.2.2 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 14–1.

Figure 14–1 Create User-Defined Metric Page (OS-Based)

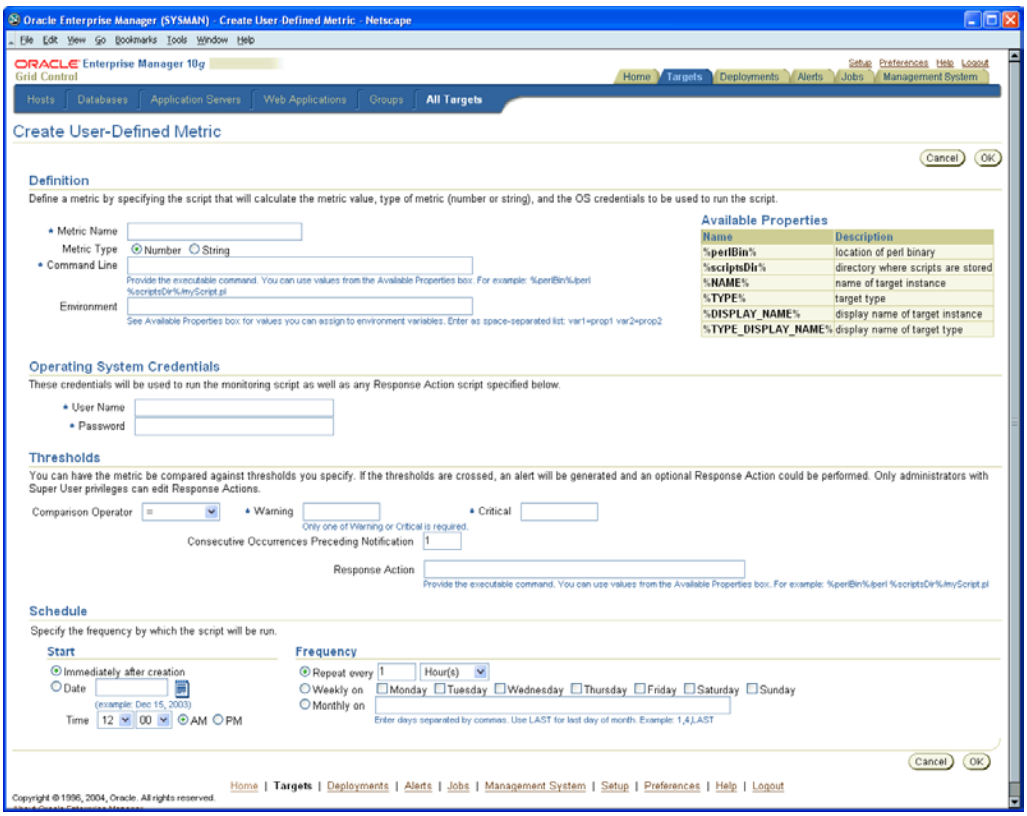

Key elements of this page are described in the following tables.

Table 14–3 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: `/u1/bin/perl /u1/scripts/myScript.pl` 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 14–3 (Cont.) Create User-Defined Metric Page: Definition

User-Interface Element	Description
Environment	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. 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: `var1=%perlBin% var2=%scriptsDir% var3=/u1/orahome10`

Table 14–4 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 14–5 Create User-Defined Metric Page: Threshold

User-Interface Element	Description
Comparison Operator	Select the comparison method Enterprise Manager should use to compare the value returned by the user-defined script to the threshold values. Available Comparison Operators 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	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. Specifically, an alert triggers at the warning severity level if the following comparison is true: <script_value> <comparison_operator> <warning_threshold> and if the consecutive occurrences preceding notification has been reached.

Table 14–5 (Cont.) Create User-Defined Metric Page: Threshold

User-Interface Element	Description
Critical	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. Specifically, an alert triggers at the critical severity level if the following comparison is true: <script_value> <comparison_operator> <critical_threshold> and if the consecutive occurrences preceding notification has been reached.
Consecutive Occurrences Preceding Notification	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.
Response Action	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. For example, the Management Agent executes the response action if: The Alert severity is Warning or Critical AND There is a change in severity (for example, warning -> critical, critical --> warning, clear --> warning or critical) For more information, see Enterprise Manager online help.

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.

14.2.3 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.

14.3 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 14–2.

Figure 14–2 Create User-Defined Metric Page (SQL-Based)

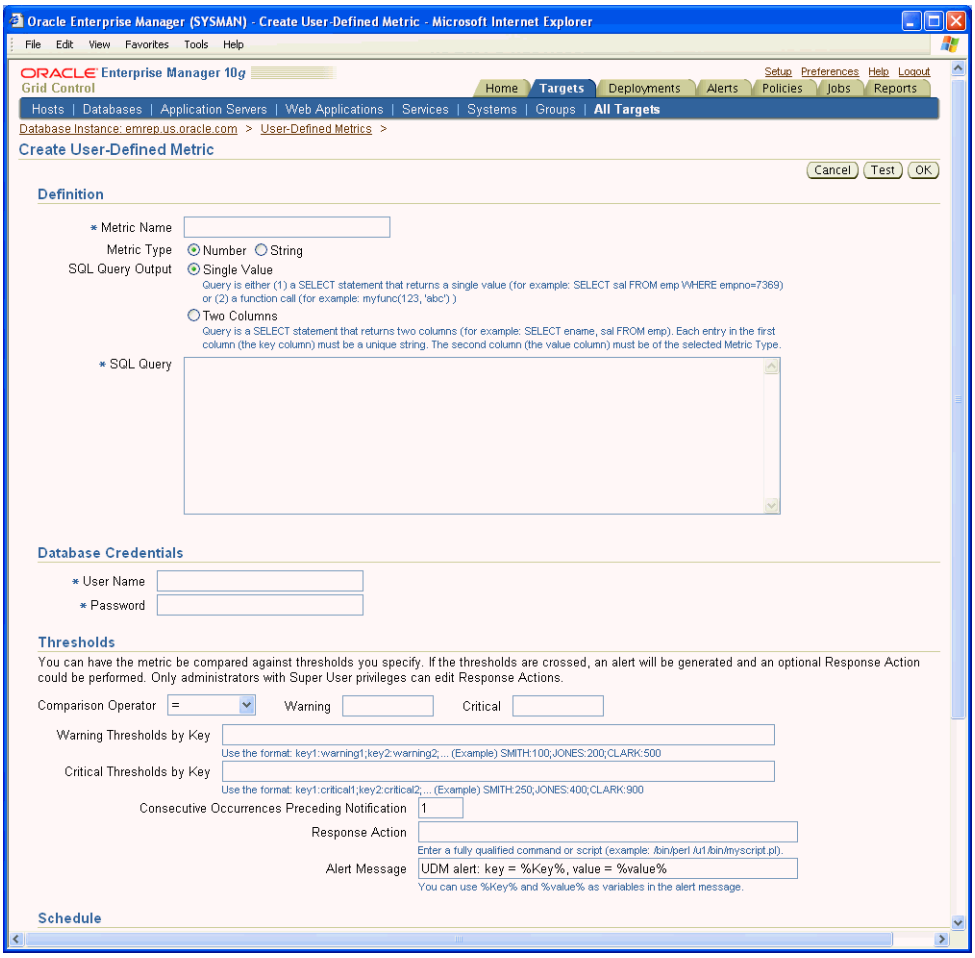

Key elements of this page are described in the following tables.

Table 14–6 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). ■ **Single Value**: Query is one of the following types. A *SELECT statement* returning a single value. Example: SELECT sal FROM emp WHERE empno=7369 A *function call* returning a single value. Example: myfunc(123, 'abc') ■ **Two Columns**: Query is a SELECT statement that returns two columns and possibly multiple rows. Example: SELECT ename, sal FROM emp. 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 14–6 (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 14–7 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 14–8 Create User-Defined Metric Page: Threshold

User-Interface Element	Description
Comparison Operator	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. Available Comparison Operators 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	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. Specifically, an alert triggers at the warning severity level if the following comparison is true: <query_value> <comparison_operator> <warning_threshold> and if the consecutive occurrences preceding notification has been reached.

Table 14–8 (Cont.) Create User-Defined Metric Page: Threshold

User-Interface Element	Description
Critical	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. Specifically, an alert triggers at the critical severity level if the following comparison is true: <query_value> <comparison_operator> <critical_threshold> and if the consecutive occurrences preceding notification has been reached.
Warning Thresholds by Key and Critical Thresholds by Key	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: SELECT ename FROM emp Threshold settings for this example are shown. **Use the format** *key:value* . **Keys are case-sensitive.** ■ 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. All other keys will use the threshold values specified in the Warning and Critical fields.
Consecutive Occurrences Preceding Notification	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.
Response Action	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. For example, the Management Agent executes the response action if: The Alert severity is Warning or Critical AND There is a change in severity (for example, warning -> critical, critical --> warning, clear --> warning or critical) For more information, see Enterprise Manager online help.

Table 14–8 (Cont.) Create User-Defined Metric Page: Threshold

User-Interface Element	Description
Alert Message	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. For example, a payroll system alert for underpayment of salary might be defined as: **Underpaid Employee: %Key% has salary of %value%** 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. If the SQL query returns 1 column, only the %value% variable is applicable in the alert message.

The User-Defined Metric Schedule interface lets you specify the frequency at which the SQL query or function should be run.

14.3.1 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.

14.3.1.1 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')
```

14.3.1.2 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
```

14.4 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.

14.4.1 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.

14.4.2 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.

14.4.3 Deploying User-Defined Metrics Across Many Targets Using Monitoring Templates

Monitoring Templates simplify the task of standardizing monitoring settings across your enterprise by allowing you to specify the monitoring settings once and applying them to your monitored targets. You can thus use Monitoring Templates as a way to propagate User-Defined Metrics across a large number of targets.

Assuming you have created the User-Defined Metric on the host or database target, you can use Monitoring Templates to propagate the User-Defined Metric to other hosts or database targets.

To create a Monitoring Template for the User-Defined Metric:

1. Click Setup.
2. From the vertical navigation bar, click Monitoring Templates
3. Click Create. The Copy Target Settings page appears.

4. Specify the host or database on which you defined the User-Defined Metric and click Continue.
5. Fill in the requisite information on the General page.
6. On the Metric Thresholds page, you can choose to keep or remove the other metrics that have been copied over from the target.
7. You can also edit the User-Defined Metric's thresholds, collection schedule, and corrective actions.
8. On the Policies page, you can choose to keep or remove any policy rules that have been copied over from the target.
9. Click OK to save the template settings to the Management Repository.

Once the template containing the User-Defined Metric has been created, you can propagate the User-Defined Metric by applying the template to other hosts or databases.

Important: For OS-based User-Defined Metrics, you will first need to separately deploy the OS Script used by the User-Defined Metric to all destination hosts. The OS Script should reside in the same location across all host targets.

For SQL-based User-Defined Metrics, if the SQL query specified is a function call, then you need to create this function across all databases on which the SQL-based User-Defined Metric will be created.

To apply the monitoring template:

1. On the Monitoring Templates page, select the monitoring template and click Apply.
2. On the Apply Monitoring Template page, add the targets on which the User-Defined Metric should be created.
3. If a two-column SQL-based User-Defined Metric is part of the template, the Metric with Multiple Thresholds option is applied according one of the following situations:
   - **Situation One**: The target to which the template will be applied does not contain the two-column SQL User-Defined Metric defined in the template. In this situation, regardless of which Metric with Multiple Thresholds option is chosen, the User-Defined metric is copied to the target when you apply the template.
   - **Situation Two**: The target to which the template will be applied does contain the two-column SQL User-Defined Metric. The name (case insensitive) and return value (numeric, scalar, or two column) of both the target and template User-Defined Metrics must match. In this situation, you must select one of the Metric with Multiple Thresholds options:
     - *Apply threshold settings for monitored objects common to both template and target*: For only those keys which the target has in common with the template, the target threshold values will be set to the values defined in the template. This option is chosen by default and is recommended for most situations.
     - *Duplicate threshold settings on target*: For keys which are common between target and template, the thresholds will be set to the values defined in the

template. Any extra keys (and their thresholds) that exist in the template but not on the target will be copied to the target in anticipation that these keys will be created in the target at some point in the future. Any extra keys (and their thresholds) that exist on the target but not in the template will be deleted from the target.

Important: If there are other metrics in the monitoring template, refer first to the Enterprise Manager online help for implications of what this option means for these other metrics.

4. Click Continue.
5. On the subsequent page, specify the credentials that should be used when running the User-Defined Metric on the destination targets.
6. Click Finish.
7. When you return back to the Monitoring Templates page, check that "Pending Apply Operations" count for your template is zero. This indicates the number of template apply operations that could be pending. Once they are all complete, the count should be zero.

14.4.4 Deleting User-Defined Metrics Across Many Targets Using Monitoring Templates

Just as templates can be used to deploy User-Defined Metrics across targets, templates can also be used to delete these metrics across targets should these metrics no longer be in use.

To create a Monitoring Template for the User-Defined Metric:

1. Click Setup.
2. From the vertical navigation bar, click Monitoring Templates
3. Click Create. The Copy Target Settings page appears.
4. Specify the host or database on which there is a User-Defined Metric that needs to be deleted and click Continue.
5. Fill in the requisite information on the General page.
6. On the Metric Thresholds page, remove all metrics from the template except the User-Defined Metric to be deleted.
7. Click the pencil icon to access the Edit Advanced Settings page.
8. Check the Mark for Delete option and click Continue. The Mark for Deletion icon now appears next to the User-Defined Metric on the Metric Thresholds page.
9. On the Policies page, remove all policies.
10. Click OK to save the template settings to the Management Repository.

To apply the monitoring template:

1. On the Monitoring Templates page, select the monitoring template and click Apply.
2. On the Apply Monitoring Template page, add the targets on which the User-Defined Metric should be deleted.
3. Click Continue.
4. On the subsequent page, specify the credentials that should be used when running the User-Defined Metric on the destination targets.

5. Click Finish.
6. On the Monitoring Templates page, check that "Pending Apply Operations" count for your template is zero. This indicates the number of template apply operations that could be pending. Once they are all complete, the count should be zero.

Enterprise manager will delete all User-Defined Metrics found on the selected target that match the following criteria:

- Name of the User-Defined Metrics (case insensitive)
- Return value of the User-Defined Metric (numeric or scalar)
- For SQL-based User-Defined Metrics, the output of the query (single value or two columns). The match does not take into consideration the actual script used by the User-Defined Metric. For this reason, even though the script on the target User-Defined Metric may be different from that of the template User-Defined Metric, the target User-Defined Metric will still be deleted.
- Host User-Define Metrics using scripts: You must delete the script from the host on which you want the UDM deleted.
- SQL-based User-Defined Metrics using function calls: You must delete the function from the database on which you want the User-Defined Metric deleted.

15

Using a Software Library

This chapter describes the Software Library feature of Enterprise Manager and contains the following sections:

- Overview of Software Library
- Setting up the Software Library
- Configuring the Software Library For Use With Oracle Enterprise Manager
- Creating and Deleting Entities in the Software Library
- Exporting and Importing Software Library Entities Across Enterprise Manager Deployments
- Exporting and Importing Entities Across Oracle Enterprise Manager Deployments
- Alternate Approach for Transferring a Software Library Across Different Oracle Enterprise Manager Deployments
- Deleting Entities and Purging the Contents of the Software Library
- Software Library Issues

15.1 Overview of Software Library

The Software Library serves as a repository to store certified software images (for example Oracle Database, operating system, Oracle Real Application Clusters, third party softwares) and other related entities. These can then be automatically mass deployed to provision software, software updates and servers using Oracle Enterprise Manger in a reliable and repeatable manner. These provisioning operations, which are unattended and can be scheduled, lead to substantial cost savings.

Software Library can store the following types of entities:

- Components: These entities represent the primary building blocks of the provisioning framework. A component can represent Operating system software, Oracle software or any third party software and applications. Software components are individually maintained within the Software Library and versions, states and maturity levels can be associated with each component.
- Directives: These are constructs used to associate scripts with software components and images. These scripts contain directions on how to interpret and process the contents of a particular component or an image. Directives encapsulate the script, the command line used to invoke the script, and the script configuration properties. Directives are contained within a file that are stored in the Software Library and referenced from the software components that employ them. Versions, states and maturity levels can be associated with each Directive.

- Images, Network Templates, Hardware Templates, and Storage Templates: These entities are associated with the Bare Metal Provisioning application of the Oracle Enterprise Manager and are used to provision software on bare metal and live servers. Briefly, an Image can be described as a collection of components along with the necessary directives that create a deployable configuration for a single or set of target machines. Network templates, Hardware templates and Storage templates are used to define the network, hardware and disk layout configuration of the target machines respectively. Versions, states and maturity levels can be associated with each of these entities.

15.2 Setting up the Software Library

Software library can be created using any mounted filesystem that is readable and writeable from the Oracle Management Service (OMS). If one is using a single OMS in the EM infrastructure then one can even use directory a local directory on the OMS host for creating the Software library. One has to make sure that there is enough space available for the storage of software binaries, and associated scripts for the entities that one wants to create and store.

In case there are multiple OMSes deployed for high availability and fail over, then Software Library should be located in a directory accessible by all OMSes. For multiple OMS environments, the directory can be on a Network File Server accessible from all the OMSes. One has to ensure that there is enough space available on the shared storage to store files that hold the binary data for one's components and other entities.

15.3 Configuring the Software Library For Use With Oracle Enterprise Manager

To configure the Software library follow these steps:

1. Access the Oracle Enterprise Manager's Provisioning Application by navigating to the Deployments tab.
2. Under the Deployments tab go to the Provisioning tab. There are a number of tabs here for creating components, directives etc. One can access some or all of the tabs depending on the privileges assigned to the user.
3. Access the Administration tab. This requires superuser privileges. Refer to section 3.1 "Creating Super Administrator For Enterprise Manager" of the *Oracle Bare Metal Provisioning Best Practices* document. This is available at:

   http://www.oracle.com/technology/products/oem/pdf/bmp_best_practice.pdf
4. In the Software Library Configuration section of the Administration tab, press the Add button.
5. On the Add Software Library Location page, enter the directory location of the software library being created and then click OK.

15.4 Creating and Deleting Entities in the Software Library

The graphical user interface of the Provisioning application has various tabs for creating Components, Directives, Images, Network and Hardware Templates, which are created and stored in the software library. One can create various subdirectories for storing the entities in the Software Library. Same tabs allow a user to delete and edit an entity stored in the Software Library. They also allow a user to view metadata information for an entity.

Refer to section 4.2 "Creating Directives, Components, Network Profiles for RAC-CRS Image" of the *Bare Metal Provisioning Best Practices* document as an example of how to create various entities. This is available at:

http://www.oracle.com/technology/products/oem/pdf/bmp_best_practice.pdf

15.5 Exporting and Importing Software Library Entities Across Enterprise Manager Deployments

Creating an entity (for example, components or directives) involves the following steps:

- Creation of meta-data. As a part of the creation process a user enters information like name, type, version and other information associated with the entity. This information serves as the meta-data of the entity and is entered into the necessary tables in the Management Repository. Metadata for software library entities is stored in the following tables of EM Management Repository:

  MGMT_SWLIB_DIRECTORIES

  MGMT_SWLIB_ENTITIES

  MGMT_SWLIB_MATURITY_STATUS

  MGMT_SWLIB_ENTITY_REVISIONS

  MGMT_SWLIB_ENTITY_REFERENCES

  MGMT_SWLIB_ENTITY_DOCUMENTS

  MGMT_SWLIB_ENTITY_DATA

  MGMT_SWLIB_ENTITY_PARAMETERS

  MGMT_SWLIB_REVISIONS_PARAMETERS

  MGMT_SWLIB_ENTITY_PLATFORMS

  MGMT_SWLIB_DATA_DIRECTORIES

- Optionally storing a file (for example user uploading executable binaries for a operating system component or script for a directive, tar files containing cloned oracle homes created and stored by the Enterprise Manager in the Software library, etc.). This file is stored under a top-level directory of the file system specified while configuring the Software Library.

  An entity is thus possesses two types of information, meta-data and associated files, which are stored in Oracle Enterprise Manager Management Repository (Management Repository) and Software Library respectively.

15.6 Exporting and Importing Entities Across Oracle Enterprise Manager Deployments

Components and Directives can be exported and imported across different Software Libraries used by different EM deployments. The `deploymentLibrabryExport` `deploymentLibrabryImport` scripts used for importing and exporting the software library entities respectively. These scripts are present in the bin directory of the OMS Oracle Home and provide the flexibility to transfer entities from a test to production or between different production environments.

The scripts support the following import/export use cases:

- Exporting all directives and importing the same.
- Exporting all components and importing the same.
- Exporting all entities in the software library (including Network profiles, Images etc.) and importing the same.

15.6.1 How the Importing and Exporting Scripts Funtion

When invoked the export scripts generate an xml file which encapsulates the meta-data and relationships between the entities, and two directories called binaries and documentation. The xml file and the directories are created under the location provided while invoking the script. This location can be mounted filesystem wrtieable by the OMS. The xml and the two directories can then be transferred to a different EM deployment for consumption by the import script.

15.6.1.1 Export Script

The export script can be located at the following location:

- For OMS on Linux:
  `<OMS HOME>/bin/deploymentLibraryExport.sh`
- For OMS on Windows:
  `<OMS HOME>/bin/deploymentLibraryExport.bat`

Sample usage is:

<OMS HOME>/bin//deploymentLibraryExport.sh <conn string> <user> <passwd> <dir loc> <-metadata | -full> [-dir | -comp]

Where the following is true:

<conn string> is of form: "*jdbc:oracle:thin:@dbhost:dbport:sid*" where dbhost, dbport and sid to be replaced appropriately with the hostname, database port and database SID for the Management Repository.

<user> is Management Repository username. This has to be the SYSMAN user or any other user with higher privileges to access the Software Library tables in the Management Repository database.

<passwd> is password for the username supplied for the Management Repository.

<dir loc> is the directory location where entity xml file, binaries directory and documentation directory would be created.

<-metadata | -full> if used with -full, then binary associated with the components also gets exported, else only metadata is exported.

[-dir | -comp] is optional; use -dir to export all directives and -comp to export all components. If nothing is used, all entities will be exported.

15.6.1.2 Import Script

The import script can be located at the following location:

- For OMS on Linux:
  `<OMS HOME>/bin/deploymentLibraryImport.sh`
- For OMS on Windows:
  `<OMS HOME>/bin/deploymentLibraryImport.bat`

Sample usage is:

<OMS HOME>/bin/ /deploymentLibraryImport.sh <conn string> <user> <passwd> <basedir> <filename> [-force]

Where the following is true:

<conn string> is of form: "jdbc:oracle:thin:@dbhost:dbport:sid" where dbhost, dbport and sid to be replaced appropriately with the hostname, database port and database SID for the Management Repository.

<user> is Management Repository username. This has to be the SYSMAN user or any other user with higher privileges to access the Software Library tables in the Management Repository database.

<passwd> is password for the username supplied.

<basedir> is absolute directory location where entity xml file created after the export operation is present

<filename> is the name of the xml file created after the export operation

[-force] is optional; and used when an entity to be imported is already present in the repository. This will create a new revision of the entity in the repository.

Note: If the -force option is not used then the import operation will be aborted if there are existing entities in Software library with the same name and under the same directory structure as the ones being imported. When the import script aborts then none of the entities would be imported.

15.6.2 Use Cases for Import and Export

The following examples are use cases for the Import and Export scripts.

1. **Use Case 1: Exporting all directives**

   Sample usage of the export script for transferring all directives is shown below:

   <OMS HOME>/bin//deploymentLibraryExport.sh <conn string> <user> <passwd> <dir loc> -dir

2. **Use Case 2: Exporting all components**

   Sample usage of the export and import scripts for transferring all directives is shown below.

   For exporting only the meta-data of all the components the export script can be used as follows:

   <OMS HOME>/bin//deploymentLibraryExport.sh <conn string> <user> <passwd> <dir loc> -metadata -comp

   For exporting both meta-data and associated files of all the components the export script can be used as follows:

   <OMS HOME>/bin//deploymentLibraryExport.sh <conn string> <user> <passwd> <dir loc> -full -comp

3. **Use Case 3: Exporting all entities (for example, images and network profiles) along with components and directives)**

   Sample usage of the export and import scripts for transferring all directives is shown below:

For exporting only the meta-data of all the entities the export script can be used as follows:

<OMS HOME>/bin//deploymentLibraryExport.sh <conn string> <user> <passwd> <dir loc> -metadata

For exporting both meta-data and associated files of all the entities the export script can be used as follows:

<OMS HOME>/bin//deploymentLibraryExport.sh <conn string> <user> <passwd> <dir loc> -full

4. **Use Case 4: Importing all entities or all directives or all components**

   Sample usage of the scripts for transferring all directives, all components or all entities is shown below:

   <OMS HOME>/bin/ /deploymentLibraryImport.sh <conn string> <user> <passwd> <basedir> <filename> [-force]

15.7 Alternate Approach for Transferring a Software Library Across Different Oracle Enterprise Manager Deployments

Using Software Library export and import scripts is the Oracle recommended method for transferring the contents of a Software library. However, a Software Library can also be transferred across EM deployments by following the given steps:

1. Archiving the Software Library:
   a. Export and archive the metadata for the entities in the EM Management Repository using database export feature.
   b. Archive the base directory of the Software Library that was configured while setting up the Software Library.
2. Restoring the Software Library:
   a. Import the archived metadata into the Management Repository of the other EM deployment using the database import feature.
   b. Restore the archived base directory by either copying or mounting it to a location accessible by the OMS/OMSes of the other EM deployment.
   c. Configure the Software Library for use with the Enterprise Manager if it is not done already.

15.8 De-Configuring a Software Library

Go to the Administration tab, Software Library section, choose the Software Library entry and then press the Remove button to de-configure a Software Library.

15.9 Deleting Entities and Purging the Contents of the Software Library

The entities can be deleted from the relevant tabs provided by the provisioning application. But it is to be noted that deleting the entity does not purge the files associated with it from the software library file system.

To clean up and delete entities completely from the file system of the Software Library, one needs to run the purgeDeploymentLibrary script. The script is located at the following location:

For OMS on Linux: <OMS HOME>/bin/ purgeDeploymentLibrary.sh

Sample usage of purgeDeploymentLibrary

<OMS HOME>/bin/purgeDeploymentLibrary.sh <conn string> <username> <password> [-job <oms_host>]

Where the following is true:

<conn string> is of form: "jdbc:oracle:thin:@dbhost:dbport:sid" where dbhost, dbport and sid to be replaced appropriately with the hostname, database port and database SID for the Oracle Management Repository.

<username> is Oracle Management Repository username. This has to be the SYSMAN user or any other user with higher privileges to access the Software Library tables in the Management Repository database.

<password> is password for the username supplied.

-job <oms_host> is optional; if provided a job would be submitted (schedule is immediate). User can see the status of the job by navigating to the Jobs tab on Enterprise Manager console. <oms_host> is OMS hostname. In case there are multiple OMSes then host name of any one OMS can be used along with the -job option.

Note: The *-job* option is optional though recommended by Oracle as the purge operation could be a time intensive operation.

15.10 Software Library Issues

Component creation sometimes fails with "Cannot create under the software library, please contact your administrator". This may happen if the Software Library is not configured with the Enterprise Manager. Please follow the Configuring Software Library section. Create the components once the software library is configured.

Metadata for entities is intact in the Management Repository but file system of the Software library where the associated files are stored has been corrupted. Restoring the file system should be fine but one will loose the entities that were created since the last backup. These entities will still show up in the Provisioning UI but errors will be encountered while accessing them or attempting to deploy them.

There is no support for purging Software Library entities from a Windows OMS.

Additional references concerning *Oracle Bare Metal Provisioning Best Practices* are available at the following location:

http://www.oracle.com/technology/products/oem/pdf/bmp_best_practice.pdf

16

Configuring Remedy Connector

Remedy Connector integrates Remedy Help Desk 6.x with Enterprise Manager. Using this connector, you can create a Remedy trouble ticket, update an existing ticket, or close a ticket based on alerts in Enterprise Manager.

This chapter provides the following information for setting up and configuring Remedy Connector:

- Introduction to Remedy Connector
- Prerequisites
- Installing Remedy Connector
- Configuring the Remedy Connector
- Creating Remedy Trouble Tickets
- Enabling SSL for HTTPS
- Navigating Between Remedy and Enterprise Manager
- Out-of-Box Templates
- Reading Ticket Templates
- Customizing Ticket Templates
- Defining New Templates
- Remedy Connector Tips

16.1 Introduction to Remedy Connector

Remedy Connector integrates Enterprise Manager with Remedy Help Desk through either HTTP or HTTPS connection. You can create, update, or close tickets based on only the following types of alerts in Enterprise Manager:

- Metric alerts
- Availability alerts (includes alerts for `Up, Down, Blackout Started, Blackout Ended, Agent Unreachable`, `Agent Unreachable Resolved, Metric Error Detected,` and `Mertic Error Resolved).`

The following sections explain various Remedy Connector concepts that you must understand before you start using Remedy Connector.

16.1.1 Autoticketing

Whenever an alert is triggered in Enterprise Manager, Remedy Connector can automatically open or update a ticket. You can specify the set of alerts for which tickets must be opened and the alert severity for which this should happen.

This can be done in Notification Rules, the user-defined rules that define the criteria by which notifications should be sent for alerts.

See Also: Chapter 13, "Configuring Notifications"

After the ticket is opened, any subsequent update of the alert, for example change in alert severity will cause an annotation to the ticket. Once the alert is cleared (severity is set to `Clear`), then the ticket can optionally be closed.

See Also: "Creating Trouble Ticket Automatically" on page 16-8

16.1.2 Manual Ticketing

From the Enterprise Manager console, you can choose to manually open a Remedy ticket based on an open alert in Enterprise Manager. The Remedy Connector will populate the ticket with details based on the alert and the ticket template selected.

See Also: "Creating a Ticket Manually" on page 16-11

16.1.3 Ticket Templates

Ticket templates are transformation style sheets in XSLT format used to transform Enterprise Manager alerts to ticket format before the requests are sent to Remedy Help Desk.

These template are used to specify how Enterprise Manager alert attributes can be used to populate the fields of a Remedy ticket. A ticket template helps in the mapping of Enterprise Manager Alert fields into Remedy ticket fields.

In autoticketing, a notification method is created for each registered ticket template. The selected notification method determines which ticket template is used when a notification is sent out to the Connector. In the case of manual ticketing, you have to select a ticket template before submitting a request to create ticket.

The Enterprise Manager installation includes some out-of-box ticket templates to facilitate easy usage of this feature.

See Also: "Out-of-Box Templates" on page 16-15

16.1.4 Grace Period

Grace period provides you with a configuration to prevent the creation of large number of tickets for frequently re-occurring alerts. For alerts that occur frequently within a relatively short time interval, it is often desirable to open and maintain one trouble ticket that tracks each occurrence of the alert instead of separate tickets each time.

For recurring alerts, the Grace Period is a time period during which re-occurrences of the same alert will cause an existing ticket for the alert to be updated (or re-opened) instead of causing a new ticket to be opened.

For example, an alert triggers and a ticket is opened for it. If the Grace Period is one hour and the alert is cleared at 10:00 am, then if the same alert retriggers before 11:00 am (one hour grace period), then the ticket that had been originally created for the alert will be updated/reopened rather than creating a new ticket.

16.2 Prerequisites

Before using Remedy Connector, ensure that you meet the following pre-requisites:

- Remedy Helpdesk 6.x is installed and configured.
- Remedy Helpdesk Web services are up and running. See "Web Service Details" on page 16-68.

16.3 Installing Remedy Connector

Remedy Connector is installed as part of the Enterprise Manager base installation. That is, Connector installation is part of the Oracle Management Server (OMS) installation.

After you install Enterprise Manager, when you access the Enterprise Manager console as a Super Administrator, you should see Remedy Connector in the Management Connector Setup page as shown in Figure 16–1(click **Setup** from the Enterprise Manager console and then **Management Connectors** in the left pane).

The default installation is based on default Remedy Web services which do not support any annotation history through the `Worklog` (the history option in the Remedy ticket). For details of Worklog and registering of Worklog template, see "Using Worklog" on page 16-67.

16.4 Configuring the Remedy Connector

1. As Super Administrator, from the Enterprise Manager console, click **Setup.**

   Overview of Setup page appears.

2. Click **Management Connectors** in the left pane.

   The Management Connector Setup page appears. For the Remedy Connector row, the Configured column should be blank (Figure 16–1).

   Note: A check mark instead indicates that the Connector is already configured.

3. Click the **Configure** icon for the Remedy Connector.

   The General tab of the Configure Management Connector page appears (Figure 16–2).

4. Provide the required settings. See "General Settings" for details.

5. Click **OK.**

   The Management Connectors Home page reappears. The row for the Remedy Connector should have a check mark in the Configured column.

6. (Optional) To check for the available ticket templates, click the configure icon again.
7. Click the Ticket Templates tab.

   All out-of-box ticket templates should appear in the table.

If any of the ticket templates are missing, you can register it using the emctl command from `ORACLE_HOME/bin` directory, where `ORACLE_HOME` is the Oracle home directory of OMS.

Run the following command as user with `execute` privilege on emctl and can read the ticket template:

```
emctl register_ticket_template connector <ticketTemplate.xsl>
<server> <port> <database sid/service name for RAC DB>
<username> <password> <connectorTypeName> <connectorName>
<templateName> <description>
```

> **Note:** In the case of multiple OMS installation, you need to run this command only once from any of the OMSs.

Example 16–1

```
emctl register_ticket_template connector Remedy_DefaultCategory_LowPriority.xsl
$emHost $dbPort $dbSID sysman $sysmanPwd "Remedy Connector" "Remedy Connector"
"Low Priority Template" "This template creates a ticket with low priority and
default categorization"
```

emctl Parameters

Table 16–1 emctl Parameters

Parameter	Description
`ticketTemplate.xsl`	Fully qualified name of the ticket template file. The file resides in the Connector home directory: `$OMS_HOME/sysman/connector/Remedy_Connector` Oracle recommends that you use intuitive names since there might be notification methods created with the same names and you have to choose one of them when you use autoticketing feature. Use `xsl` as the file extension since the format is XSLT. For example, `Remedy_DefaultCategory_LowPriority.xsl.` If the file is in a different directory, provide complete path for the file.
`server`	Host name of the Enterprise Manager repository.
`port`	Listener port of the repository.
`database sid/ Service Name for RAC DB`	Repository database instance ID or service name if you are using RAC database as the repository.
`username`	Specify `SYSMAN.`
`password`	Password for `SYSMAN.`
`connectorTypeName`	Specify `"Remedy Connector".` The double quotes ("") are mandatory.
`connectorName`	Specify `"Remedy Connector".` The double quotes ("") are mandatory.

Table 16–1 (Cont.) emctl Parameters

Parameter	Description
`templateName`	An intuitive name for the ticket template that will be displayed in Enterprise Manager.
`description`	A short description for the ticket template. This description gets displayed in the Enterprise Manager also.

Figure 16–1 Management Connectors Page

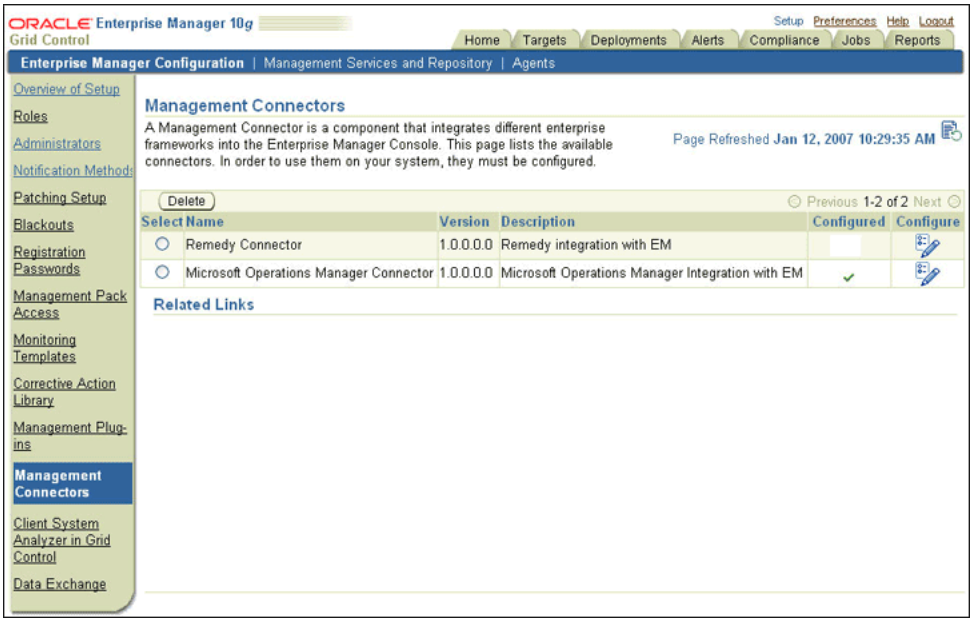

Figure 16–2 Configure Management Connector Page

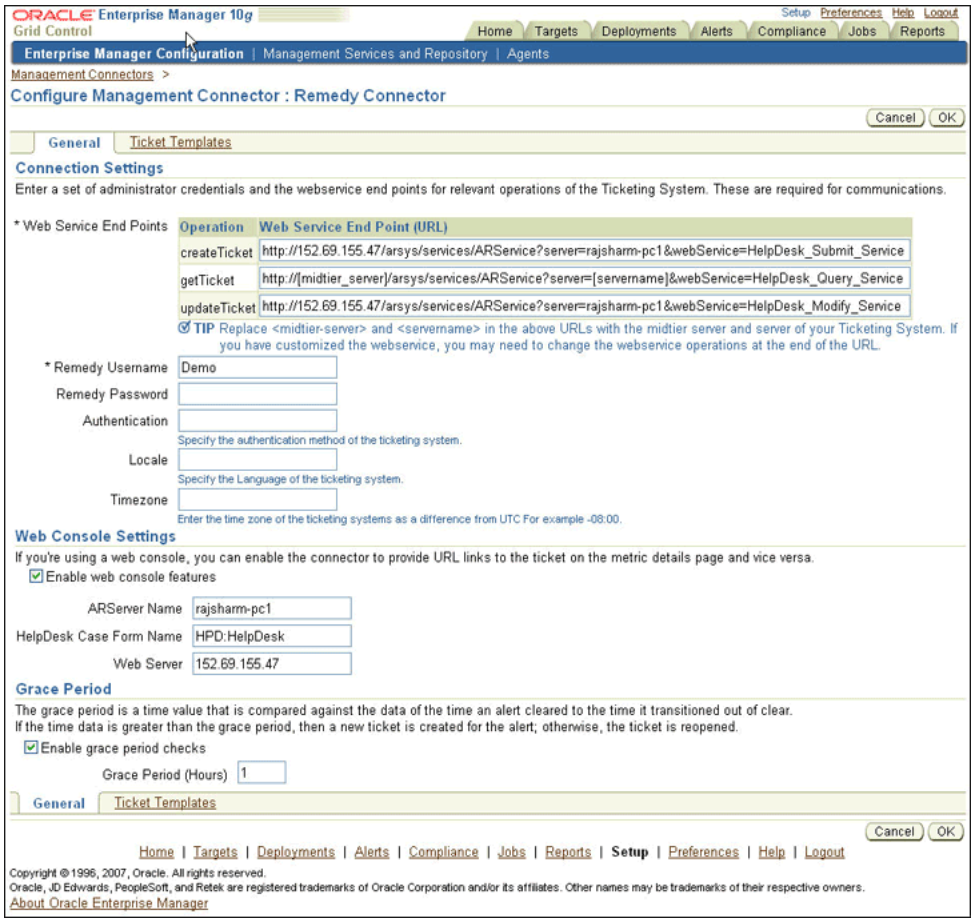

16.4.1 General Settings

The following sections explain how to provide various configuration details.

16.4.1.1 Connection Settings

The Remedy Trouble Ticket connector communicates with the Help Desk through their Web services. Mandatory fields are indicated by a * mark.

- ***Web Service Endpoints** — Endpoints to `createTicket`, `updateTicket`, and `getTicket` Web services exposed by Remedy Help Desk. See "Web Service Details" for additional details.
- ***Remedy Username**— User with the privilege to create, update, and query tickets in Remedy.
- **Remedy Password**— Password associated with the supplied Remedy user.
- **Authentication**— String that a Remedy administrator sets for additional security. Applies only if the Remedy Administrator has configured it on the Remedy AR server. It communicates with the server if there is a secondary authentication server that can be used to verify the Remedy credentials.
- **Locale**— Language of the Remedy system (optional).
- **Time Zone**— Time zone of the Remedy AR System Server (optional).

See Also: Section "Remedy User preferences settings" in the Remedy Remedy AR System Server product manual *Remedy Action Request System 6.3 - Developing AR System Applications: Advanced*

16.4.1.2 Web Console Settings

Web Console settings are required if the users wants the Connector to provide links to Remedy Help Desk tickets created by Enterprise Manager in context of an alert.

To enable this functionality, provide the following Web console settings.

- **Enable web console** — Check this box to enable launching of Remedy ticket page within context from Enterprise Manager.
- **ARServer Name**— The Remedy AR Server name.
- **HelpDesk Case Form Name**—The Remedy form name that the Remedy Web Services (you configured the connector to use) is based on. The Remedy default Help Desk webservices, for example, use the form HPD:HelpDesk.
- **Web Server** — The name or IP address of the server that hosts Remedy Mid-Tier.

16.4.1.3 Grace Period

You can enable and disable the grace period and configure its value. By default, the grace period is disabled. See "Grace Period" on page 16-2 for details.

This setting applies to all alerts processed by Remedy Connector.

16.4.2 Working with Ticket Templates

The following sections provide information about registering, removing, replacing and adding ticket templates.

16.4.2.1 Registering Templates

Ticket templates need to be registered before they are recognized in Enterprise Manager. For autoticketing, a notification method is created for each registered ticket template and a ticket is created and updated based on the ticket template associated with the selected notification method. In the case of manual ticketing, registered ticket templates are available for selection.

All registered ticket templates are displayed in the Configure Management Connector Ticket Templates page. By default, all out-of-box ticket templates are registered. To register additional ticket templates that you create, use the following emctl command:

```
emctl register_ticket_template connector <ticketTemplate.xml>
<server> <port> <database sid> <username> <password>
<connectorTypeName> <connectorName> <templateName> <description>
```

See Also: "emctl Parameters" on page 16-4

16.4.2.2 Viewing Template Code

Click a template name to view the XSLT code for the template.

The ticket templates are in XSLT format. A basic knowledge of XSLT is required to understand the code.

16.4.2.3 Removing Template

> **Important:** If the template you delete has a notification rule associated with it, the notification will fail.

1. Select the template and click **Remove.**
2. At prompt, confirm the removal.
3. Before you exit the page, click **OK** for the deletion to take effect.

> **Note:** Unless you click **OK** before you exit, the template is not deleted. Next time you go to the Ticket Template page, the templates reappear.

Though the ticket template is removed from the Enterprise Manager repository, it is still available on OMS in the Connector home directory. You can re-register the ticket template later if required.

16.4.2.4 Replacing Templates

To replace an existing ticket template, do the following in sequence:

- Delete the ticket template
- Register the new template using emctl

16.4.2.5 Adding New Templates

To add templates other than the out-of-box templates provided by Oracle, users should define new templates and register them using emctl.

> **See Also:** "Defining New Templates" on page 16-63

16.5 Creating Remedy Trouble Tickets

You can have trouble tickets created automatically, or you can create them manually. The following sections explain how to create both types.

16.5.1 Creating Trouble Ticket Automatically

Perform the following steps to create a ticket automatically:

1. Review the Out-of-Box Templates.
2. Select an appropriate ticket template having the desired mapping of Enterprise Manager alert fields to the Remedy ticket fields.
3. If you do not have a ticket template that satisfies your requirement, create one and register it.
4. Create a notification rule using the following steps:

> **Important:** Do not select more than one ticket template for this notification rule.

1. From the Enterprise Manager console, click **Preferences.**
2. In the left pane, under Notification, click **Rules,** then **Create**.
3. In the Create Notification Rule General page, specify the rule name, description and the targets for which this rule should apply.
4. In the Create Notification Rule Availability page select the availability states for which you want to create tickets.
5. In the Create Notification Rule Metrics page, select the metrics and their associated alert severities for which you want to create and update tickets.

   Ensure that you select all relevant alert severities if you want the ticket to be updated when the alert severity changes. For example, to open a ticket for critical alert on CPU Utilization(%) metric, and the ticket to be updated if the CPU Utilization(%) changes to warning or clear severity, then in the notification rule, select `Critical`, `Warning`, or `Clear` severities for the CPU Utilization(%) metric.
6. In Create Notification Rule Methods page, choose the ticket template from the Advanced Notification Methods table (Table 16–3).

   In the table, registered ticket templates appear as Java Callback type notification methods under the same name as the ticket template's file name. This ticket template will be used to open tickets for all availability and metric alerts specified in this notification rule.

   This makes the ticket templates available for use to open tickets.

   See Also: Chapter 13, "Configuring Notifications"

The following process occurs after you create the notification rule for your alerts:

- A notification is sent to the Remedy Connector when a metric alert that matches your rule triggers. The Remedy connector creates/updates a ticket according to the ticket template as set in the notification rule.
- The ticket is created or updated on the Remedy Trouble Ticket system.
- In Enterprise Manager, the alert annotation is updated. A comment is added to the Metric Details page of the alert to indicate that a ticket was created or updated, along with the ticket ID and ticket page URL.

A ticket is updated if there is an existing active ticket for an alert. In Figure 16–4, the first screen shows the ticket in Remedy console and the second screen, the alert as displayed in Enterprise Manager.

Figure 16–3 Notification Methods

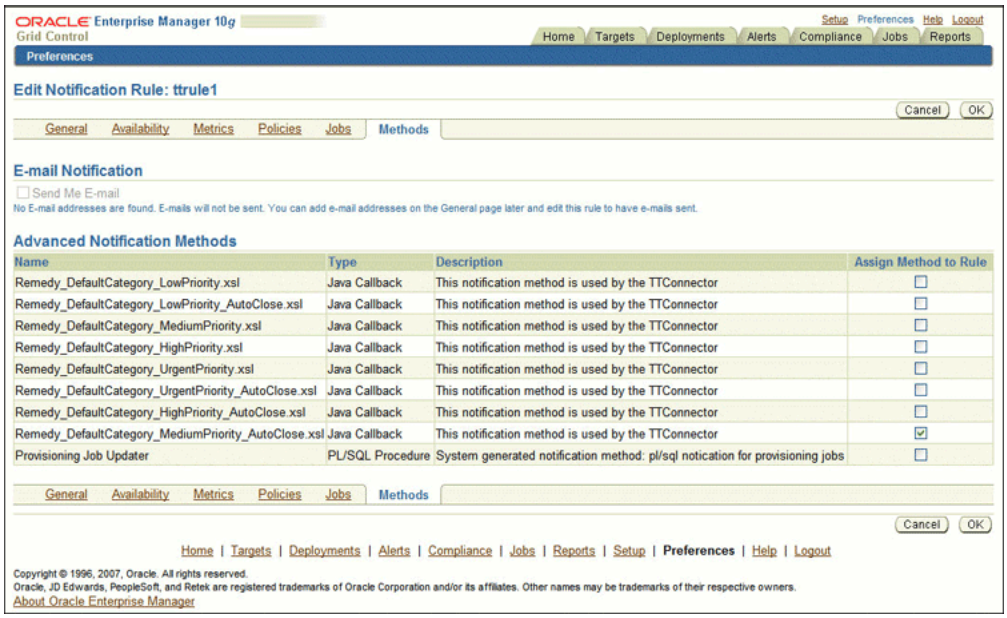

Figure 16–4 Remedy Ticket and the Alert as Displayed in Enterprise Manager

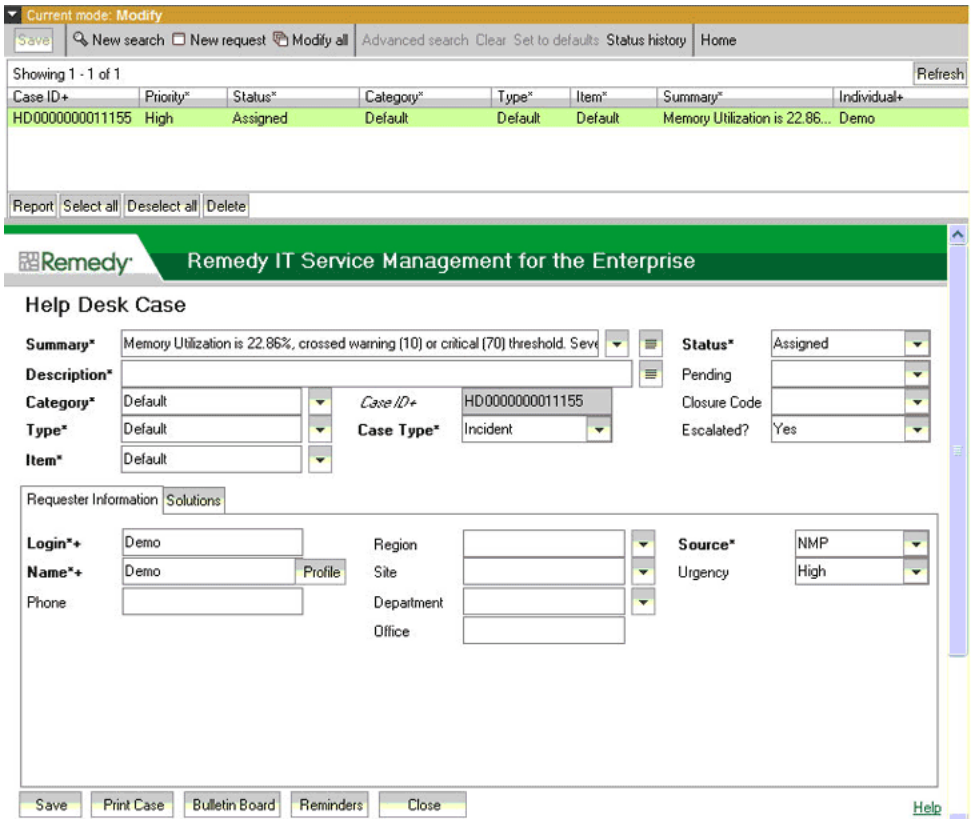

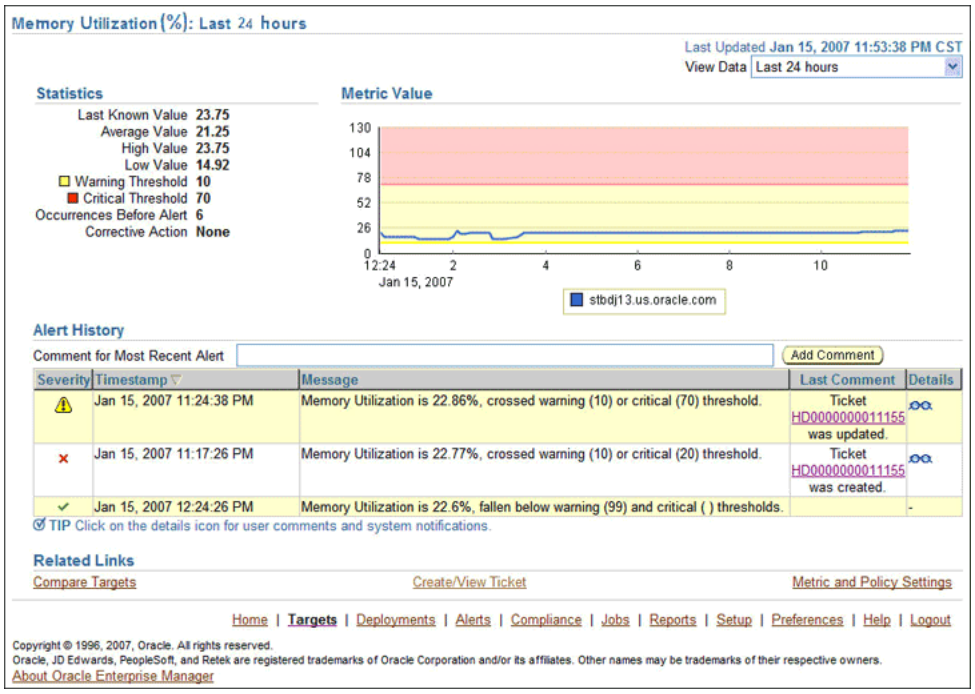

16.5.2 Creating a Ticket Manually

Perform the following steps to create a ticket manually:

1. After a metric alert occurs, go to the associated metric details page for the alert. To access this page, click the alert message in the Enterprise Manager console (Figure 16–5).
2. Click the **Create/View Ticket** link in the Related Links section.

   The Create Ticket page appears if no active ticket exists for the alert.
3. Select a ticket template and then click **Submit** (Figure 16–6).

   If you do not see the desired template, then you can register one using the emctl. "Registering Templates" on page 16-7

   If creating or updating the ticket is successful, the ticket ID appears in the Last Comment column of the Alert History table for the metric alert.

   If the Web console settings are configured and enabled, the ticket ID appears as a link to the ticket page in the Remedy Help Desk. If there is no annotation, then the ticket creation fails and error information is logged in the file `emoms.log`.

   Note: You cannot manually update the ticket using Remedy Connector. You have to manually update the ticket in the Remedy AR server for any subsequent alert change.

Figure 16–5 Metric Details Page

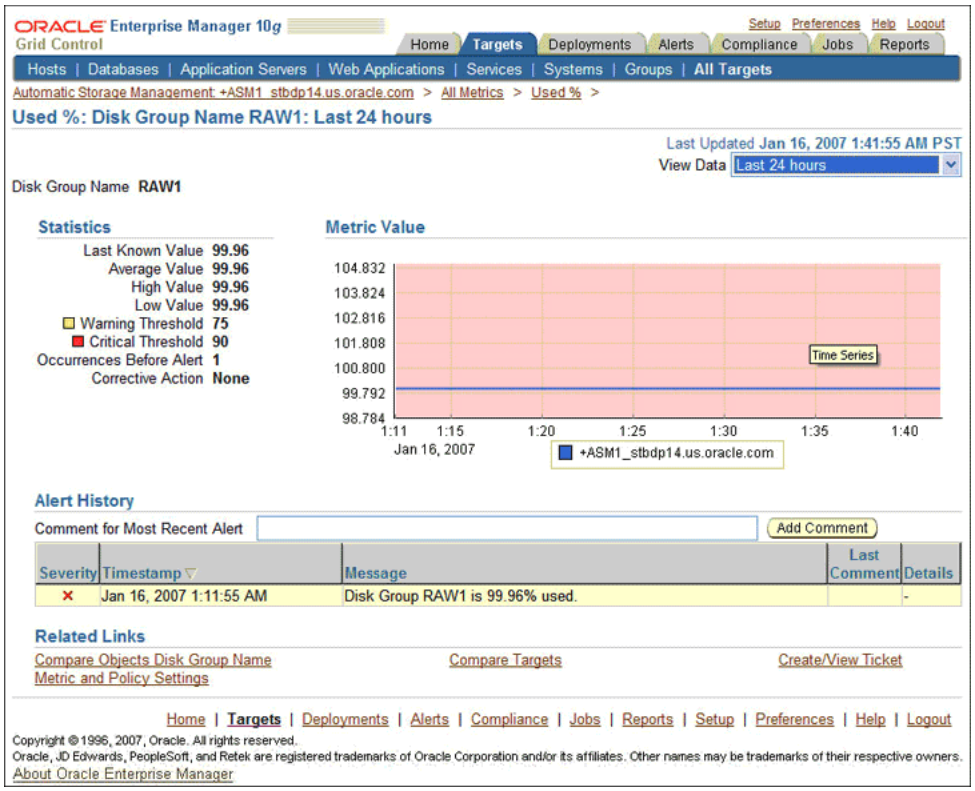

Figure 16–6 Create Ticket Page

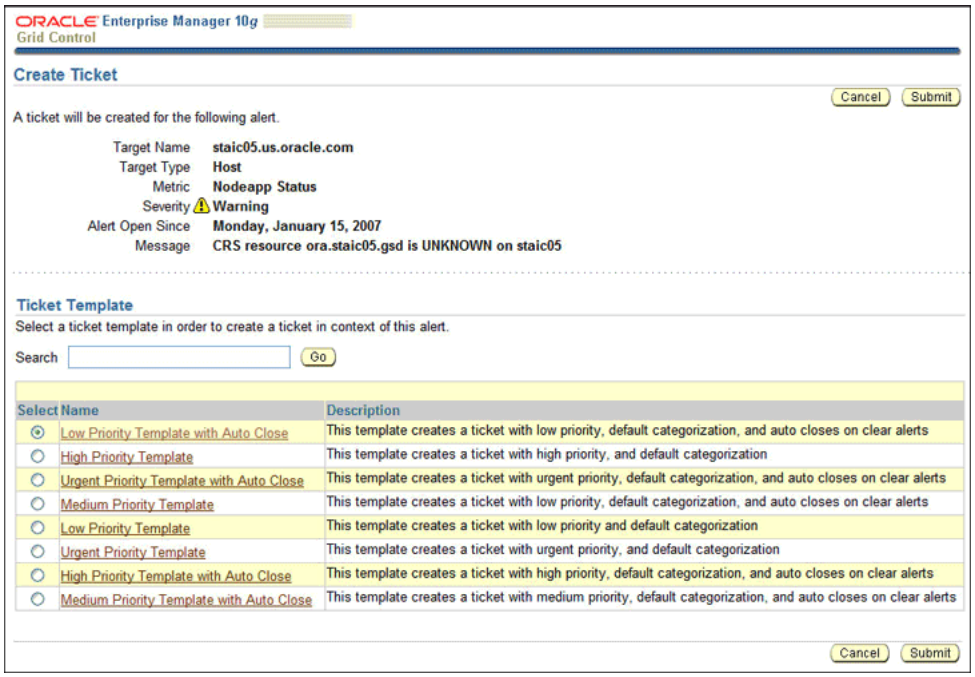

16.6 Enabling SSL for HTTPS

Follow the steps provided in this section if you choose HTTPS as the protocol to establish a connection between Remedy AR server and Enterprise Manager.

16.6.1 Generating a Certificate Request File

Generate a certificate request file for the Remedy AR server and send it to the Certificate authority, for example VeriSign.

> **Note:** The certificate request file is dependent on the Web server used by Remedy.

16.6.2 Importing the Certificate from the Certificate Authority

After you get the certificate, import it to the Web server used by Remedy. The import mechanism varies depending on the Web server used by Remedy Help Desk.

16.6.3 Adding Signed Certificates to Wallet Manager

> **Note:** Oracle Wallet Manager is available at `$ORACLE_HOME/bin` on OMS. See *Oracle Application Server Administrator's Guide* for details.

Do the following on Enterprise Manager:

1. As Super Administrator, create a wallet using `orapki utility` using the following command at OMS host:

   ```
   orapki wallet create -wallet client -auto_login
   ```

   > **Note:** `orapki` is available at `$ORACLE_HOME/bin` on OMS.

2. Add the trusted certificate to the wallet by using the following command:

   ```
   orapki wallet add -wallet client -trusted_cert -cert
   verisignCert.cer
   ```

3. To view the content of the wallet, run the following command:

   ```
   orapki wallet display -wallet client
   ```

   Ensure that `ewallet.p12` is available.

4. In Oracle Wallet Manager, then open the client certificate `ewallet.p12.`
5. Go to Select Trusted Certificates and select **Operations** on the main menu.
6. Select **Export All Trusted Certificates**.
7. Save the file as `certdb.txt.`
8. Place the file `certdb.txt` in the connector home root directory (`$OMS_HOME/sysman/connector`).

   If the file `certdb.txt` already exists in the root directory, open the file and add the contents of your `certdb.txt` to the existing content.

Now this file can be used by the Java SSL for communication between Enterprise Manager and Remedy AR server in HTTPS mode.

> **See Also:** For information on creating wallet, see "Creating and Viewing Oracle Wallets with orapki" in the *Oracle Database Advanced Security Administrator's Guide, 10g Release 2 (10.2).*

16.7 Navigating Between Remedy and Enterprise Manager

The following sections explain how to switch from one console to the other.

16.7.1 Navigating from Remedy to Enterprise Manager

From a ticket page, click the link in the **Description** field to the Alert Details page in the ticket message body (Figure 16–7). This takes you to the Enterprise Manager console login page. After you provide Enterprise Manager username and password, you will be forwarded to the alert related to this ticket.

> **Note:**
> - The Enterprise Manager user whose name you specify should at least have `View` privileges on the target on which the alert was raised.
> - On the Remedy console, if the URL appears as text, then you will need to cut and paste the URL into the browser.

Figure 16–7 Alert Details in Remedy Console

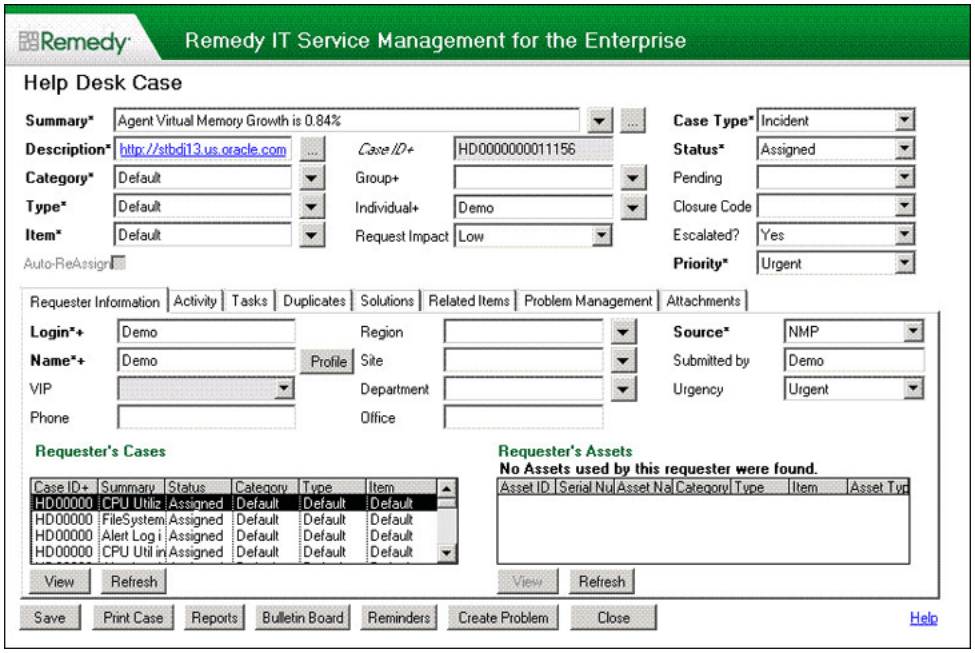

16.7.2 Navigating from the Enterprise Manager to Remedy

1. In the Enterprise Manager console, click the alert message to go to the metric details page for the alert.

2. In the Alert History table, locate the ticket ID link in the Last Comment column.
3. (If not found) Click the icon in the Details column to get more information about the alert.
4. On the page that comes up, locate the ticket ID in the Alert Details table.
5. Click the ticket ID link. You are forwarded to the Remedy Web console login page.
6. Provide valid Remedy account details.

   The ticket page associated with that alert is displayed.

> **Note:** If you do not use Remedy Web console, uncheck Enable web console option in the Web Console Settings section so that ticket ID is shown in plain text. Otherwise it is displayed as a link that does not work.

16.8 Out-of-Box Templates

This section provides details of the out-of-box ticket templates shipped along with the Remedy Connector. The ticket templates specify the mappings between Enterprise Manager alert attributes and Remedy ticket attributes.

All out-of-box templates cause the following actions to occur when a ticket is created for an alert:

- Write alert information to `Description` (Remedy ticket description).
- Set the Remedy ticket summary based on the alert message. On update, the ticket summary field is updated to include the latest alert message information.
- Set the `Category`, `Item`, and `Type` fields in Remedy to default.
- Set the `Priority` (ticket's priority) to the value indicated by the file name of the ticket template. For instance, Remedy_DefaultCategory_HighPriority.xsl sets the ticket priority to `High.`

Following are the out-of-box templates:

- Remedy_DefaultCategory_LowPriority.xsl
- Remedy_DefaultCategory_MediumPriority.xsl
- Remedy_DefaultCategory_HighPriority.xsl
- Remedy_DefaultCategory_UrgentPriority.xsl

Following are the out-of-box templates with the `AutoClose` suffixed to the file names. They set the ticket status to `Close` when the event severity value becomes `Clear`.

- Remedy_DefaultCategory_LowPriority_AutoClose.xsl
- Remedy_DefaultCategory_MediumPriority_AutoClose.xsl
- Remedy_DefaultCategory_HighPriority_AutoClose.xsl
- Remedy_DefaultCategory_UrgentPriority_AutoClose.xsl

Following are the out-of-box templates with `Wlog` suffixed to the file names. They are customized for the Web services with worklog enabled.

On update, the `Description` (Remedy ticket description) is updated with the latest event information and the work log is updated with the latest severity and timestamp information.

- Remedy_DefaultCategory_LowPriority_w_Wlog.xsl
- Remedy_DefaultCategory_MediumPriority_w_Wlog.xsl
- Remedy_DefaultCategory_HighPriority_w_Wlog.xsl
- Remedy_DefaultCategory_UrgentPriority_w_Wlog.xsl
- Remedy_DefaultCategory_LowPriority_AutoClose_w_Wlog.xsl
- Remedy_DefaultCategory_MediumPriority_AutoClose_w_Wlog.xsl
- Remedy_DefaultCategory_HighPriority_AutoClose_w_Wlog.xsl
- Remedy_DefaultCategory_UrgentPriority_AutoClose_w_Wlog.xsl

16.8.1 Reading Ticket Templates

The Table 16–2 illustrates the creation of a ticket using Remedy_DefaultCategory_HighPriority_AutoClose.xsl. This illustration will help you to read a ticket template.

Table 16–2 Ticket Creation (Remedy_DefaultCategory_HighPriority_AutoClose.xsl Mappings)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*

Table 16–2 (Cont.) Ticket Creation (Remedy_DefaultCategory_HighPriority_ AutoClose.xsl Mappings)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values from the alert context.
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. For example, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Escalated		Blank
Hotlist		Blank
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		High*
Region		Blank
Request Urgency		High*

Table 16–2 (Cont.) Ticket Creation (Remedy_DefaultCategory_HighPriority_AutoClose.xsl Mappings)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Requester Login Name	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Requester Name	HDUser	
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	
Type		Default*
Work Log		Blank
Create Time		Blank

Table 16–3 Ticket Updates (Remedy_DefaultCategory_HighPriority_AutoClose.xsl Mappings)

Ticket Attributes	Enterprise Manager Alert Attributes	Value
Status	Severity	▪ If severity is `Clear`, then set the ticket to the status `Closed`. ▪ If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`.
Summary	Message, Severity	
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCategory_HighPriority_AutoClose.xsl Source Code with Annotations

Use the mapping table (Table 16–2) as a reference to read the following XSLT.

```
<?xml version='1.0' encoding='UTF-8'?>
<xsl:transform version="1.0"
               xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
                xmlns:ns0="http://xmlns.oracle.com/sysman/connector/tt"
                targetNamespace="http://xmlns.oracle.com/sysman/connector/tt"
                elementFormDefault="qualified">
```

```
<!--
This template creates an incident type ticket with default categorization
(Category: Default, Type:Default, Item:Default), and high priority. On update,
the description and message fields are updated, and the ticket is closed if the
associated alert has cleared.
-->

<xsl:template match="ns0:EventModel">
  <xsl:choose>
     <!-- Create the ticket if there is no ticket ID. -->
    <xsl:when test="normalize-space(ns0:TicketId) = ''">
      <urn:Create_Helpdesk_Case xmlns:urn="urn:HelpDesk_Submit_Service">

         <!-- EDIT THE TAG VALUES BELOW TO CHANGE HOW A TICKET IS FILLED
              DURING TICKET CREATION. REFER TO THE REMEDY HELPDESK MANUAL
              FOR DESCRIPTION OF THESE HELPDESK SUPPORT DATAFIELDS -->

      <urn:Case_Type>Incident</urn:Case_Type>
        <urn:Category>Default</urn:Category>
        <urn:Department></urn:Department>
        <urn:Description>
            Ticket created by EM Remedy Connector.
            --------------------------------------
            EM User: <xsl:value-of select="ns0:EMUser"/>

            Event Information:
            Target Type: <xsl:value-of select="ns0:TargetType"/>
            Metric Column: <xsl:value-of select="ns0:MetricColumn"/>
            Metric Name: <xsl:value-of select="ns0:MetricName"/>
            <xsl:choose>
            <xsl:when test="normalize-space(ns0:KeyColumn) != ''">
            Key Column: <xsl:value-of select="ns0:KeyColumn"/>
            Key Values: <xsl:value-of select="ns0:KeyValues"/>
            </xsl:when>
            </xsl:choose>
            Severity: <xsl:value-of select="ns0:Severity"/>
            Collection Time: <xsl:value-of select="ns0:CollectionTime"/>
            Target Host: <xsl:value-of select="ns0:TargetHost"/>
            <xsl:choose>
            <xsl:when test="normalize-space(ns0:NotificationRuleName) != ''">
            Notification Rule: <xsl:value-of select="ns0:NotificationRuleName"/>
            </xsl:when>
            </xsl:choose>
            URL: <xsl:value-of select="ns0:EventPageURL"/>
        </urn:Description>
        <urn:Escalated></urn:Escalated>
        <urn:Hotlist></urn:Hotlist>
        <urn:Item>Default</urn:Item>
        <urn:Office></urn:Office>
        <urn:Orig_Submitter>
          <xsl:value-of select="ns0:HDUser"/>
        </urn:Orig_Submitter>
        <urn:Pending></urn:Pending>
        <urn:Phone_Number></urn:Phone_Number>
        <urn:Priority>High</urn:Priority>
        <urn:Region></urn:Region>
        <urn:Request_Urgency>High</urn:Request_Urgency>
        <urn:Requester_Login_Name>
          <xsl:value-of select="ns0:HDUser"/>
```

```
            <urn:Requester_Login_Name>
              <xsl:value-of select="ns0:HDUser"/>
            </urn:Requester_Login_Name>
            <urn:Requester_Name>
              <xsl:value-of select="ns0:HDUser"/>
            </urn:Requester_Name>
            <urn:Site></urn:Site>
            <urn:Source>NMP</urn:Source>
            <urn:Status>New</urn:Status>
            <urn:Summary>
              <xsl:value-of select="ns0:Message"/>
            </urn:Summary>
            <urn:Type>Default</urn:Type>
            <urn:WorkLog></urn:WorkLog>
            <urn:Create_Time></urn:Create_Time>
          </urn:Create_Helpdesk_Case>
        </xsl:when>
        <!-- Update the ticket otherwise.. -->
        <xsl:otherwise>
          <urn:SetBy_Case_ID xmlns:urn="urn:HelpDesk_Modify_Service">
           <!--
                           UNCOMMENT THE TAGS YOU WISH TO HAVE MODIFIED WHENEVER THE
TICKET IS UPDATED, AND GIVE THEM DESIRED VALUES
           -->
           <!-- <urn:Accounting_Code></urn:Accounting_Code>           -->
           <!--  <urn:Assignee_Login_Name></urn:Assignee_Login_Name>  -->
           <!--  <urn:Case_Type></urn:Case_Type>                      -->
           <!--  <urn:Category></urn:Category>                        -->
           <!--  <urn:Department></urn:Department>                    -->
            <!--  <urn:Description></urn:Description>                          -->
            <!-- <urn:Escalated></urn:Escalated>                      -->
            <!-- <urn:Hotlist></urn:Hotlist>                          -->
            <!-- <urn:Item></urn:Item>                                -->
            <!-- <urn:Office></urn:Office>                            -->
            <!-- <urn:Pending></urn:Pending>                          -->
            <!-- <urn:Phone_Number></urn:Phone_Number>                -->
            <!-- <urn:Priority></urn:Priority>                        -->
            <!-- <urn:Region></urn:Region>                            -->
            <!-- <urn:Request_Urgency></urn:Request_Urgency>          -->
            <!-- <urn:Requester_Login></urn:Requester_Login>          -->
            <!-- <urn:Requester_Name></urn:Requester_Name>            -->
            <!-- <urn:Site></urn:Site>                                -->
            <!-- <urn:Solution_Description></urn:Solution_Description>-->
            <!-- <urn:Solution_Summary></urn:Solution_Summary>        -->
            <!-- <urn:Source></urn:Source>                            -->
            <xsl:choose>
              <xsl:when test="ns0:Severity = 'Clear'">
                <urn:Status>Closed</urn:Status>
              </xsl:when>
              <xsl:when test="ns0:GracePeriodCheckMade = 'Yes'">
                <urn:Status>Assigned</urn:Status>
              </xsl:when>
            </xsl:choose>
            <!-- <urn:Submitted_By></urn:Submitted_By>                -->
            <urn:Summary>
              <xsl:value-of select="ns0:Message"/> Severity:<xsl:value-of
select="ns0:Severity"/>
            </urn:Summary>
            <!-- <urn:Type></urn:Type>                                -->
            <urn:Case_ID>
```

```
            <xsl:value-of select="ns0:TicketId"/>
          </urn:Case_ID>

        </urn:SetBy_Case_ID>
      </xsl:otherwise>
    </xsl:choose>
  </xsl:template>
</xsl:transform>
```

16.8.1.1 Templates Explained

The tables in this section map the fields in all out-of-box ticket templates shipped along with the Remedy Connector.

Remedy_DefaultCategory_LowPriority.xsl

Table 16–4 Ticket Creation (Remedy_DefaultCategory_LowPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values from the alert context.
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. For example, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Escalated		Blank
Hotlist		Blank
Item		"Default"*
Office		Blank

Table 16–4 (Cont.) Ticket Creation (Remedy_DefaultCategory_LowPriority.xsl)

Remedy		
Ticket Attributes	**Enterprise Manager Alert Attributes**	**Value**
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		Low
Region		Blank
Request Urgency		Low
Requester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log		Blank
Create Time		Blank

Table 16–5 Ticket Updates (Remedy_DefaultCategory_LowPriority.xsl)

Ticket Attributes	Enterprise Manager Alert Attributes	Value
Status	Severity	If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Summary	Message, Severity	The alert message in context with the severity appended.

Table 16–5 (Cont.) Ticket Updates (Remedy_DefaultCategory_LowPriority.xsl)

Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCategory_MediumPriority.xsl

Table 16–6 Ticket Creation (Remedy_DefaultCategory_MediumPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values from the alert context.
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. For example, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Escalated		Blank
Hotlist		Blank
Item		"Default"*

Table 16–6 (Cont.) Ticket Creation (Remedy_DefaultCategory_MediumPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		Medium
Region		Blank
Request Urgency		Medium
Requester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log		Blank
Create Time		Blank

Table 16–7 Ticket Updates (Remedy_DefaultCategory_MediumPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Status	Severity	If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Summary	Message, Severity	The alert message in context with the severity appended.

Table 16–7 (Cont.) Ticket Updates (Remedy_DefaultCategory_MediumPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCaterogry_HighPriority.xsl

Table 16–8 Ticket Creation (Remedy_DefaultCaterogry_HighPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values from the alert context.
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Escalated		Blank
Hotlist		Blank

Table 16–8 (Cont.) Ticket Creation (Remedy_DefaultCaterogry_HighPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		High
Region		Blank
Request Urgency		High
Requester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log		Blank
Create Time		Blank

Table 16–9 Ticket Updates (Remedy_DefaultCaterogry_HighPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Status	Severity	If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Summary	Message, Severity	The alert message in context with the severity appended.

Table 16–9 (Cont.) Ticket Updates (Remedy_DefaultCaterogry_HighPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCaterogry_UrgentPriority.xsl

Table 16–10 Ticket Creation (Remedy_DefaultCaterogry_UrgentPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values from the alert context.
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Escalated		Blank
Hotlist		Blank

Table 16–10 (Cont.) Ticket Creation (Remedy_DefaultCaterogry_UrgentPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		Urgent
Region		Blank
Request Urgency		Urgent
Requester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log		Blank
Create Time		Blank

Table 16–11 Ticket Update (Remedy_DefaultCaterogry_UrgentPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Status	Severity	If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Summary	Message, Severity	The alert message in context with the severity appended.

Table 16–11 (Cont.) Ticket Update (Remedy_DefaultCaterogry_UrgentPriority.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Following are the templates with the `AutoClose` suffixed to the file names. They set the ticket status to `Close` when the event severity value becomes `Clear`:

Remedy_DefaultCaterogry_LowPriority_AutoClose.xsl

Table 16–12 Ticket Creation (Remedy_DefaultCaterogry_LowPriority_AutoClose.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values from the alert context.
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	

*Table 16–12 (Cont.) Ticket Creation (Remedy_DefaultCaterogry_LowPriority_*

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Escalated		Blank
Hotlist		Blank
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		Low
Region		Blank
Request Urgency		Low
Requester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log		Blank
Create Time		Blank

Table 16–13 Ticket Update (Remedy_DefaultCaterogry_LowPriority_AutoClose.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Status	Severity	▪ If severity is `Clear`, then set the ticket to the status `Closed`. ▪ If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Summary	Message, Severity	The alert message in context with the severity appended.
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCaterogry_MediumPriority_AutoClose.xsl

Table 16–14 Ticket Creation (Remedy_DefaultCaterogry_MediumPriority_AutoClose.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*

*Table 16–14 (Cont.) Ticket Creation (Remedy_DefaultCaterogry_MediumPriority_*

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing) TargetType, EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing) TargetType MetricColumn (name of the metric, for example, CPU Utilization(%)) MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load) KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name) KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.) Severity CollectionTime TargetHost NotificationRuleName EventPageURL (URL to the metric details page in context of the alert)	Values from the alert context.
Escalated		Blank
Hotlist		Blank
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		Medium

*Table 16–14 (Cont.) Ticket Creation (Remedy_DefaultCaterogry_MediumPriority_*

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Region		Blank
Request Urgency		Medium
Requester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log		Blank
Create Time		Blank

Table 16–15 Ticket Updates (Remedy_DefaultCaterogry_MediumPriority_AutoClose.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Status	Severity	▪ If severity is `Clear`, then set the ticket to the status `Closed`. ▪ If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Summary	Message, Severity	The alert message in context with the severity appended.

*Table 16–15 (Cont.) Ticket Updates (Remedy_DefaultCaterogry_MediumPriority_*

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCaterogry_HighPriority_AutoClose.xsl

Table 16–16 Ticket Creation (Remedy_DefaultCaterogry_HighPriority_AutoClose.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values from the alert context.
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Escalated		Blank
Hotlist		Blank
Item		"Default"*

*Table 16–16 (Cont.) Ticket Creation (Remedy_DefaultCaterogry_HighPriority_*

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		High
Region		Blank
Request Urgency		High
Requester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log		Blank
Create Time		Blank

Table 16–17 Ticket Updates (Remedy_DefaultCaterogry_HighPriority_AutoClose.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Status	Severity	▪ If severity is `Clear`, then set the ticket to the status `Closed`. ▪ If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Summary	Message, Severity	The alert message in context with the severity appended.
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCategory_UrgentPriority_AutoClose.xsl

Table 16–18 Ticket Creation (Remedy_DefaultCategory_UrgentPriority_AutoClose.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*

*Table 16–18 (Cont.) Ticket Creation (Remedy_DefaultCategory_UrgentPriority_*

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values from the alert context.
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Escalated		Blank
Hotlist		Blank
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		Urgent
Region		Blank
Request Urgency		Urgent

*Table 16–18 (Cont.) Ticket Creation (Remedy_DefaultCategory_UrgentPriority_*

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Requester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log		Blank
Create Time		Blank

Table 16–19 Ticket Updates (Remedy_DefaultCategory_UrgentPriority_AutoClose.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Status	Severity	▪ If severity is `Clear`, then set the ticket to the status `Closed`. ▪ If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Summary	Message, Severity	The alert message in context with the severity appended.
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Following are the templates with `Wlog` suffixed to the file names. They are customized for the worklog Web_service.

On update, the `Description` (Remedy ticket description) is updated with the latest event information and the work log is updated with the latest severity and timestamp information.

Remedy_DefaultCaterogry_LowPriority_w_Wlog.xsl

Table 16–20 Ticket Creation (Remedy_DefaultCaterogry_LowPriority_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values from the alert context.
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Escalated		Blank
Hotlist		Blank
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.

*Table 16–20 (Cont.) Ticket Creation (Remedy_DefaultCaterogry_LowPriority_w_*

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Pending		Blank
Phone Number		Blank
Priority		Urgent
Region		Blank
Request Urgency		Urgent
UrgentRequester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log	Severity, CollectionTime	The alert severity and collection time in context.
Create Time		Blank

Table 16–21 Ticket Updates (Remedy_DefaultCaterogry_LowPriority_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values of the alert in context
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Status	Severity	If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Worklog	Severity, CollectionTime	The values in context
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated).	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCategory_MediumPriority_w_Wlog.xsl

Table 16–22 Ticket Creation (Remedy_DefaultCategory_MediumPriority_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing) TargetType MetricColumn (name of the metric, for example, CPU Utilization(%)) MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load) KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name) KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.) Severity CollectionTime TargetHost NotificationRuleName EventPageURL (URL to the metric details page in context of the alert)	Values from the alert context.
Escalated		Blank
Hotlist		Blank
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		Medium

*Table 16–22 (Cont.) Ticket Creation (Remedy_DefaultCategory_MediumPriority_w_*

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Region		Blank
Request Urgency		Medium
UrgentRequester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log	Severity, CollectionTime	The alert severity and collection time in context.
Create Time		Blank

Table 16–23 Ticket Updates (Remedy_DefaultCategory_MediumPriority_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values of the alert in context
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Status	Severity	If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Worklog	Severity, CollectionTime	The values in context
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCategory_HighPriority_w_Wlog.xsl

Table 16–24 Ticket Creation (Remedy_DefaultCategory_HighPriority_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values from the alert context.
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Escalated		Blank
Hotlist		Blank
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		High
Region		Blank

*Table 16–24 (Cont.) Ticket Creation (Remedy_DefaultCategory_HighPriority_w_*

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Request Urgency		High
UrgentRequester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log	Severity, CollectionTime	The alert severity and collection time in context.
Create Time		Blank

Table 16–25 Ticket Updates (Remedy_DefaultCategory_HighPriority_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing) TargetType MetricColumn (name of the metric, for example, CPU Utilization(%)) MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load) KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name) KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.) Severity CollectionTime TargetHost NotificationRuleName EventPageURL (URL to the metric details page in context of the alert)	Values of the alert in context
Status	Severity	If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Worklog	Severity, CollectionTime	The values in context
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCategory_UrgentPriority_w_Wlog.xsl

Table 16–26 Ticket Creation (Remedy_DefaultCategory_UrgentPriority_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing) TargetType MetricColumn (name of the metric, for example, CPU Utilization(%)) MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load) KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name) KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.) Severity CollectionTime TargetHost NotificationRuleName EventPageURL (URL to the metric details page in context of the alert)	Values from the alert context.
Escalated		Blank
Hotlist		Blank
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		Urgent
Region		Blank

*Table 16–26 (Cont.) Ticket Creation (Remedy_DefaultCategory_UrgentPriority_w_*

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Request Urgency		Urgent
UrgentRequester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
	Default*	
Type		
Work Log	Severity, CollectionTime	The alert severity and collection time in context.
Create Time		Blank

Table 16–27 Ticket Updates (Remedy_DefaultCategory_UrgentPriority_w_Wlog.xsl)

Remedy Ticket Attributes	Alert Attributes	Value
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values of the alert in context
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Status	Severity	If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Worklog	Severity, CollectionTime	The values in context
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCategory_LowPriority_AutoClose_w_Wlog.xsl

Table 16–28 Ticket Creation (Remedy_DefaultCategory_LowPriority_AutoClose_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing) TargetType MetricColumn (name of the metric, for example, CPU Utilization(%)) MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load) KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name) KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.) Severity CollectionTime TargetHost NotificationRuleName EventPageURL (URL to the metric details page in context of the alert)	Values from the alert context.
Escalated		Blank
Hotlist		Blank
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		Low
Region		Blank

Table 16–28 (Cont.) Ticket Creation (Remedy_DefaultCategory_LowPriority_AutoClose_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Request Urgency		Low
UrgentRequester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log	Severity, CollectionTime	The alert severity and collection time in context.
Create Time		Blank

Table 16–29 Ticket Updates (Remedy_DefaultCategory_LowPriority_AutoClose_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing) TargetType MetricColumn (name of the metric, for example, CPU Utilization(%)) MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load) KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name) KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.) Severity CollectionTime TargetHost NotificationRuleName EventPageURL (URL to the metric details page in context of the alert)	Values of the alert in context
Status	Severity	▪ If severity is `Clear`, then set the ticket to the status `Closed`. ▪ If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Worklog	Severity, CollectionTime	The values in context
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCategory_MediumPriority_AutoClose_w_Wlog.xsl

Table 16–30 Ticket Creation (Remedy_DefaultCategory_MediumPriority_AutoClose_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values from the alert context.
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Escalated		Blank
Hotlist		Blank
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.

Table 16–30 (Cont.) Ticket Creation (Remedy_DefaultCategory_MediumPriority_ AutoClose_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Pending		Blank
Phone Number		Blank
Priority		Medium
Region		Blank
Request Urgency		Medium
UrgentRequester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log	Severity, CollectionTime	The alert severity and collection time in context.
Create Time		Blank

Table 16–31 Ticket Updates (Remedy_DefaultCategory_MediumPriority_AutoClose_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing) TargetType MetricColumn (name of the metric, for example, CPU Utilization(%)) MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load) KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name) KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.) Severity CollectionTime TargetHost NotificationRuleName EventPageURL (URL to the metric details page in context of the alert)	Values of the alert in context
Status	Severity	▪ If severity is `Clear`, then set the ticket to the status `Closed`. ▪ If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Worklog	Severity, CollectionTime	The values in context

Table 16–31 (Cont.) Ticket Updates (Remedy_DefaultCategory_MediumPriority_AutoClose_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCategory_HighPriority_AutoClose_w_Wlog.xsl

Table 16–32 Ticket Creation (Remedy_DefaultCategory_HighPriority_AutoClose_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing)	Values from the alert context.
	TargetType	
	MetricColumn (name of the metric, for example, CPU Utilization(%))	
	MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load)	
	KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name)	
	KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.)	
	Severity	
	CollectionTime	
	TargetHost	
	NotificationRuleName	
	EventPageURL (URL to the metric details page in context of the alert)	
Escalated		Blank

Table 16–32 (Cont.) Ticket Creation (Remedy_DefaultCategory_HighPriority_AutoClose_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Hotlist		Blank
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank
Phone Number		Blank
Priority		High
Region		Blank
Request Urgency		High
UrgentRequester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log	Severity, CollectionTime	The alert severity and collection time in context.
Create Time		Blank

Table 16–33 Ticket Updates (Remedy_DefaultCategory_HighPriority_AutoClose_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing) TargetType MetricColumn (name of the metric, for example, CPU Utilization(%)) MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load) KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name) KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.) Severity CollectionTime TargetHost NotificationRuleName EventPageURL (URL to the metric details page in context of the alert)	Values of the alert in context
Status	Severity	▪ If severity is `Clear`, then set the ticket to the status `Closed`. ▪ If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Worklog	Severity, CollectionTime	The values in context
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

Remedy_DefaultCategory_UrgentPriority_AutoClose_w_Wlog.xsl

Table 16–34 Ticket Creation (Remedy_DefaultCategory_UrgentPriority_AutoClose_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Case Type		"Incident"*
Category		"Default"*
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing) TargetType MetricColumn (name of the metric, for example, CPU Utilization(%)) MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load) KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name) KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if USERS tablespace triggered at warning or critical severity.) Severity CollectionTime TargetHost NotificationRuleName EventPageURL (URL to the metric details page in context of the alert)	Values from the alert context.
Escalated		Blank
Hotlist		Blank
Item		"Default"*
Office		Blank
Orig Submitter	HDUser	The username that is provided in "Remedy Username" field during the configuration.
Pending		Blank

Table 16–34 (Cont.) Ticket Creation (Remedy_DefaultCategory_UrgentPriority_ AutoClose_w_Wlog.xsl)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Phone Number		Blank
Priority		Urgent
Region		Blank
Request Urgency		Urgent
UrgentRequester Login Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Requester Name	HDUser	The username that is provided in "Remedy Username" field of the Connection Settings configuration.
Site		Blank
Source		NMP* (Network Management Program)
Status		New*
Summary	Message	The alert message in context
Type		Default*
Work Log	Severity, CollectionTime	The alert severity and collection time in context.
Create Time		Blank

Table 16–35 Ticket Updates (Remedy_DefaultCategory_UrgentPriority_AutoClose)

Remedy Ticket Attributes	Enterprise Manager Alert Attributes	Value
Description	EMUser (notification rule owner when the ticket is created through auto-ticketing, and is the EM log-in user when the ticket is created through manual-ticketing) TargetType MetricColumn (name of the metric, for example, CPU Utilization(%)) MetricName (Category of the metric. For CPU Utilization(%) metric, this would be 'Load') KeyColumn** (For metrics that monitor a set of objects, KeyColumn indicates the type of object monitored. For example, for theTablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name) KeyValues** (For metrics that monitor a set of objects, the KeyValues indicate the specific object that triggered the severity. Forexample, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if USERS tablespace triggered at warning or critical severity.) Severity CollectionTime TargetHost NotificationRuleName EventPageURL (URL to the metric details page in context of the alert)	Values of the alert in context
Status	Severity	▪ If severity is `Clear`, then set the ticket to the status `Closed`. ▪ If the grace period test has already been done, and the alert is still within the grace period, then reopen the ticket by setting the ticket to the status `Assigned`; otherwise, leave the status as it is.
Worklog	Severity, CollectionTime	The values in context
Case ID	TicketId (the connector adds this into the alert context before handling the ticketing action. Required by the Remedy Web service to identify the ticket that must be updated)	

In the tables, * mark denotes a literal string and ** indicates `if the attribute applies.`

16.8.2 Customizing Ticket Templates

If the out-of-box ticket templates do not satisfy your requirements, you can modify them. To do this, Oracle recommends that you use one of the existing templates as the base template. Copy this ticket template to a new file, modify, and register the new ticket template.

In most cases, when you modify the ticket template, you might only be changing the mappings. The following examples illustrate this point:

Example 16–2

To create a template to mark the category to `MyCategory,` modify the following attribute in the template:

```
<urn:Category>MyCategory</urn:Category>
```

Example 16–3

If you only want the alert message to show up as ticket summary instead of both message and severity, modify the following attribute:

```
<urn:Summary><xsl:value-of select="ns0:Message"/></urn:Summary>
```

The templates are highly customizable. Oracle recommends that only users with advanced knowledge of XSLT make complex changes.

You can use notification rules as a filter to associate proper ticket templates with alerts. You can have as many tickets templates as you want. One notification rule can have only one ticket template.

16.8.3 Defining New Templates

The out-of-box templates are based on the HPD:HelpDesk form. If the new ticket templates you define are based on the HPD:HelpDesk form, Customizing Ticket Templates apply.

But if you use a custom Remedy Form, for instance HPD:CustomHelpDesk, you will have to define a new ticket template.

Enterprise Manager Attributes

The following table provides the Enterprise Manager fields that you can map when using the default Remedy Help Desk Web services:

Table 16–36 Enterprise Manager Attributes

Data Fields	Description
EMUser	▪ For auto-ticketing, this is the notification rule owner. ▪ For manual ticketing, this is the console user that triggered the ticket creation.
HDUser	Helpdesk user registered with the Connector; this is same as the user name specified for the WS authentication.
TicketID	Identifies the ticket associated with the current alert (this is available after ticket creation).

Table 16–36 (Cont.) Enterprise Manager Attributes

Data Fields	Description
ConnectorID	Identifies the connector that processed the event and issued the ticket creation or ticket update. This is the ID for Remedy Connector.
TargetType	Type of the target that the alert is associated with, for example, `host`.
TargetName	Name of the target that the alert is associated with. For example, `Database1` or `stadc40.us.oracle.com`.
MetricColumn	Name of the metric that triggered the alert. For example CPU Utilization(%).
MetricName	Category of the metric. For example, `Load` for the memory utilization alert.
KeyColumn	For metrics that monitor a set of objects, the KeyColumn indicates the type of object monitored. For example, for the Tablespace Space Used (%) metric that monitors tablespaceobjects, the KeyColumn is 'Tablespace Name'.
KeyValues	Key values associated with a key value base alert. For metrics that monitor a set of objects, the KeyValues indicates the specific object that triggered the severity. For example, for the Tablespace Space Used (%) metric that monitors tablespace objects, the KeyValues is 'USERS' if the USERS tablespace triggered at warning or critical severity.
Message	Description of the alert. For example, CPU Utilization is 100%, crossed warning (80) or critical (95) threshold.
Severity	Severity of the alert: `critical, warning, clear,` or `down.`
CollectionTime	Timestamp of an alert occurrence.
EventPageURL	URL to the alert details page of the alert.
NotificationRuleName	Name of the notification rule that generated the notification during auto-ticketing.
TargetTimezone	Timezone of the target associated with the alert.
GracePeriodCheckMade	Value `Yes` indicates that the alert is cleared since the last update or creation, but is within the configured grace period.
TargetHost	Name of the server hosting the target that generated the alert.

The model that contains the above attributes is described by the following schema:

```
<?xml version="1.0" encoding="US-ASCII" ?>
<xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema"
            xmlns="http://xmlns.oracle.com/sysman/connector/tt"
            targetNamespace="http://xmlns.oracle.com/sysman/connector/tt"
            elementFormDefault="qualified">

  <xsd:element name= "EventModel" type="EMEventModel"/>

  <xsd:complexType name="EMEventModel">
<xsd:sequence>
      <xsd:element name="TicketId" type="xsd:string" minOccurs="0" maxOccurs="1"
/>
      <xsd:element name="ConnectorId" type="xsd:string" minOccurs="1"
maxOccurs="1" />
      <xsd:element name="EventId" type="EventIdType" minOccurs="1" maxOccurs="1"
```

```
/>
<xsd:element name="TargetType" type="xsd:string" minOccurs="1" maxOccurs="1" />
<xsd:element name="TargetName" type="xsd:string" minOccurs="1" maxOccurs="1" />
<xsd:element name="MetricColumn" type="xsd:string" minOccurs="1" maxOccurs="1" />
<xsd:element name="MetricName" type="xsd:string" minOccurs="1" maxOccurs="1" />
<xsd:element name="KeyColumn" type="xsd:string" minOccurs="0" maxOccurs="1" />
<xsd:element name="KeyValues" type="xsd:string" minOccurs="0"
maxOccurs="unbounded" />
<xsd:element name="Message" type="xsd:string" minOccurs="1" maxOccurs="1" />
<xsd:element name="Severity" type="SeverityType" minOccurs="1" maxOccurs="1" />
<xsd:element name="SeverityCode" type="SeverityCodeType" minOccurs="1"
maxOccurs="1" />
<xsd:element name="CollectionTime" type="xsd:dateTime" minOccurs="1" maxOccurs="1"
/>
<xsd:element name="EventPageURL" type="xsd:string" minOccurs="0" maxOccurs="1" />
      <xsd:element name="EMUser" type="xsd:string" minOccurs="1" maxOccurs="1" />
      <xsd:element name="HDUser" type="xsd:string" minOccurs="1" maxOccurs="1" />
      <xsd:element name="NotificationRuleName" type="xsd:string" minOccurs="0"
maxOccurs="1" />
      <xsd:element name="TargetHost" type="xsd:string" minOccurs="1" maxOccurs="1"
/>
      <xsd:element name="GracePeriodCheckMade" type="xsd:string" minOccurs="0"
maxOccurs="1" />
      <xsd:element name="TargetTimezone" type="xsd:string" minOccurs="1"
maxOccurs="1" />
    </xsd:sequence>
</xsd:complexType>

  <xsd:complexType name="EventIdType">
    <xsd:sequence>
      <xsd:element name="TargetId" type="xsd:string" minOccurs="1" maxOccurs="1"/>
      <xsd:element name="MetricId" type="xsd:string" minOccurs="1" maxOccurs="1"/>
      <xsd:element name="KeyId" type="xsd:string" minOccurs="0" maxOccurs="1"/>
    </xsd:sequence>
  </xsd:complexType>

<xsd:simpleType name="SeverityType">
<xsd:restriction base="xsd:string">
<xsd:enumeration value="Clear" />
<xsd:enumeration value="Info" />
<xsd:enumeration value="Warning" />
<xsd:enumeration value="Critical" />
<xsd:enumeration value="Agent Unreachable Clear" />
      <xsd:enumeration value="Blackout End" />
      <xsd:enumeration value="Blackout Start" />
      <xsd:enumeration value="Metric Error End" />
      <xsd:enumeration value="Metric Error Start" />
      <xsd:enumeration value="Unknown" />
</xsd:restriction>
</xsd:simpleType>

<xsd:simpleType name="SeverityCodeType">
<xsd:restriction base="xsd:string">
<xsd:enumeration value="15" />
<xsd:enumeration value="18" />
<xsd:enumeration value="20" />
<xsd:enumeration value="25" />
<xsd:enumeration value="115" />
      <xsd:enumeration value="125" />
      <xsd:enumeration value="215" />
```

```
      <xsd:enumeration value="225" />
      <xsd:enumeration value="315" />
      <xsd:enumeration value="325" />
</xsd:restriction>
</xsd:simpleType>

</xsd:schema>
```

Remedy Attributes

This section lists the Remedy attributes available for mapping to when using the default Remedy Help Desk Web services.

Case Type
Category
Description
Escalated
Escalated
Hotlist
Item
Office
Orig Submitter
Pending
Phone Number
Priority
Region
Request Urgency
Requester Login Name
Requester Name
Site
Source
Status
Summary
Type
Work Log
Create Time

See Also: *Remedy Help Desk for the Enterprise 6.0 User's Guide*

Format for Creating Ticket Templates

To create ticket templates for custom Remedy forms, adhere to the following format:

```
<?xml version='1.0' encoding='UTF-8'?>
<xsl:transform version="1.0"
xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
xmlns:ns0="http://xmlns.oracle.com/sysman/connector/tt"
targetNamespace="http://xmlns.oracle.com/sysman/connector/tt"
elementFormDefault="qualified">

<!--
This template creates an incident type ticket with default categorization
(Category: Default, Type:Default, Item:Default), and low priority. On update,
the description and message fields are updated, and the ticket is closed if the
associated alert has cleared.
-->

<xsl:template match="ns0:EventModel">
```

```
<xsl:choose>
<xsl:when test="normalize-space(ns0:TicketId) = ''">

*[Insert your mappings from EMModel into your custom Create Ticket Webservice SOAP
Document] *

</xsl:when>
<xsl:otherwise>

* [Insert your mappings from EMModel schema into your Custom Update Ticket
Webservice SOAP Document]*

</xsl:otherwise>
</xsl:choose>
</xsl:template>
</xsl:transform>
```

16.9 Remedy Connector Tips

This section provides various tips that might help you to use Remedy Connector effectively.

16.9.1 Recommended Protocol

Oracle recommends that you use HTTPS as the protocol for the communication between Enterprise Manager and Remedy AR server.

Use HTTP only if a secure connection is not required and the data can be transferred in clear text between the two systems.

16.9.2 Supported Alerts

This release supports the following types of alerts:

- Metric alerts
- Availability alerts

16.9.3 Notification Failure

Notification would be blocked for processing if the notification device is down due to any issues, for instance, the Remedy AR server is down, the Remedy configuration on Enterprise Manager is wrong, or ticket is removed in Remedy.

Notification failure on one target impacts all other targets of the same target type for which the rule applies. That is, subsequent notifications are blocked until the issue is fixed or the maximum retrials fail.

Note: The maximum retrial period is one day.

16.9.4 Using Worklog

Worklog is a history option in the Remedy ticket that lets you maintain alert history in the ticket. The Remedy default Web services do not allow modification of this option.

To use worklog, perform the following steps before using the Remedy Connector:

1. In the Remedy AR server, import the Web service definition `HelpDesk_Modify_Service_w_Worklog.def` from the Remedy Connector home directory (`$OMS_HOME/sysman/connector_Remedy_Connector`).

   See Also: Section "Importing Object Definitions" in the Remedy Remedy AR System Server product manual *Remedy Action Request System 6.3 - Developing AR System Applications: Advanced*

2. Configure the Connector to use the `HelpDesk_Modify_Service_w_Worklog` Web service by setting the `Update Ticket` endpoint accordingly.
3. Import all packaged work log templates (select the files with names ending in `Wlog.xsl)`using the emctl command provided in "Registering Templates" on page 16-7.

16.9.5 Web Service Details

This section provides information about the Web services that you require depending on the ticket template you choose.

16.9.5.1 For Default Templates (without Worklog Support)

If you choose default ticket templates, ensure that the following HPD:HelpDesk related Web services are up and running on the Remedy AR server:

- HelpDesk_Modify_Service
- HelpDesk_Query_Service
- HelpDesk_Submit_Service

16.9.5.2 For Worklog Templates

If you choose Worklog templates (*_Wlog), then you have to import

`HelpDesk_Modify_Service_w_Worklog.def`, which is part of the remedyconnector.jar, to register the HelpDesk_Modify_Service_w_Worklog in Remedy.

17

Configuring Microsoft Operations Management Connector

Microsoft Operations Management Connector enables Oracle Enterprise Manager to retrieve alerts from Microsoft Operations Manager (MOM). The retrieved alerts are stored in the Enterprise Manager repository and displayed through the Enterprise Manager console.

This chapter provides the following information on setting up and configuring MOM Connector:

- Introduction to MOM Connector
- Installing MOM Connector
- Prerequisites
- Configuring the MOM Connector
- Enabling SSL Connection for HTTPS
- MOM Connector Tips

17.1 Introduction to MOM Connector

MOM Connector retrieves data through Web Services on HTTP and HTTPS protocols. The retrieval is done through a polling mechanism. The polling interval is user-definable. The polling interval cannot be shorter than 5 minutes. If an interval shorter than 5 minutes is defined, then it defaults to 5 minutes.

The default target_type defined in Enterprise Manager is `mom_managed_host`. The retrieved alerts are stored in the default target instance `generic_mom_managed_host`. You can create more target instances based on your requirements. See "Configuring the MOM Connector" on page 17-2.

17.2 Installing MOM Connector

MOM Connector is installed as part of the Enterprise Manager base installation. After you install Enterprise Manager, you should see the MOM Connector listed on the Management Connectors page.

See Also: *Enterprise Manager Grid Control Installation and Basic Configuration Guide* available at:
`http://www.oracle.com/technology/documentation/oem.html`

17.3 Prerequisites

Before using MOM Connector, ensure that the following pre-requisites are met:

- Microsoft Operations Manager 2005 is installed
- MOM Connector framework Web services are enabled during the installation of MOM 2005
- Internet Information Services (IIS) 6.0 or higher is deployed on the Windows server that hosts MOM

See Also: MOM 2005 Product Documentation at the following URL: `http://www.microsoft.com/mom/techinfo/productdoc/default.mspx`

17.4 Configuring the MOM Connector

1. From the Enterprise Manager console, click **Setup.**
2. Click **Management Connectors** link in the left pane.

   The Management Connector Setup page appears. For the MOM Connector row, the Configured column should be blank (Figure 17–1).

   Note: A check mark in the Configured column indicates that the Connector is already configured.

   ---

3. Click the **Configure** icon for the MOM Connector.

   The General tab of the Configure Management Connector page appears.
4. Provide the required settings. See "General Settings" on page 17-4 for details (Figure 17–2).
5. Click **OK.**
6. (Optional) Go to the Targets tab and specify the details for creating additional target instances. See "Additional Target Instances" on page 17-4 for details.
7. Click **OK.**

   The Management Connector Setup page reappears.

   The MOM Connector row should have a check mark in the Configured column.

Figure 17–1 Management Connectors Page

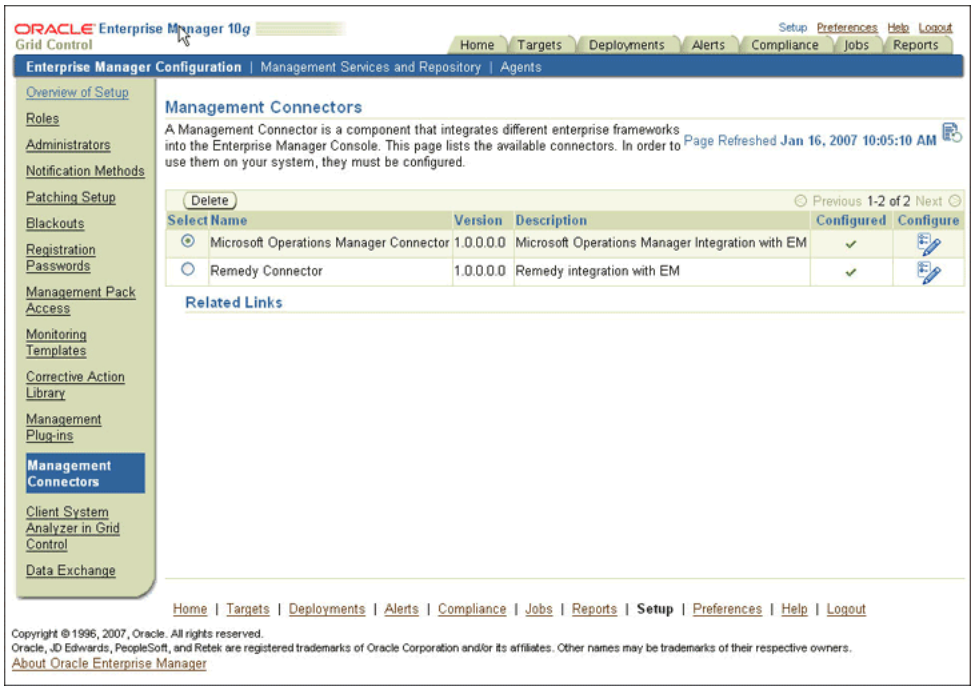

Figure 17–2 Configure Management Connector Page

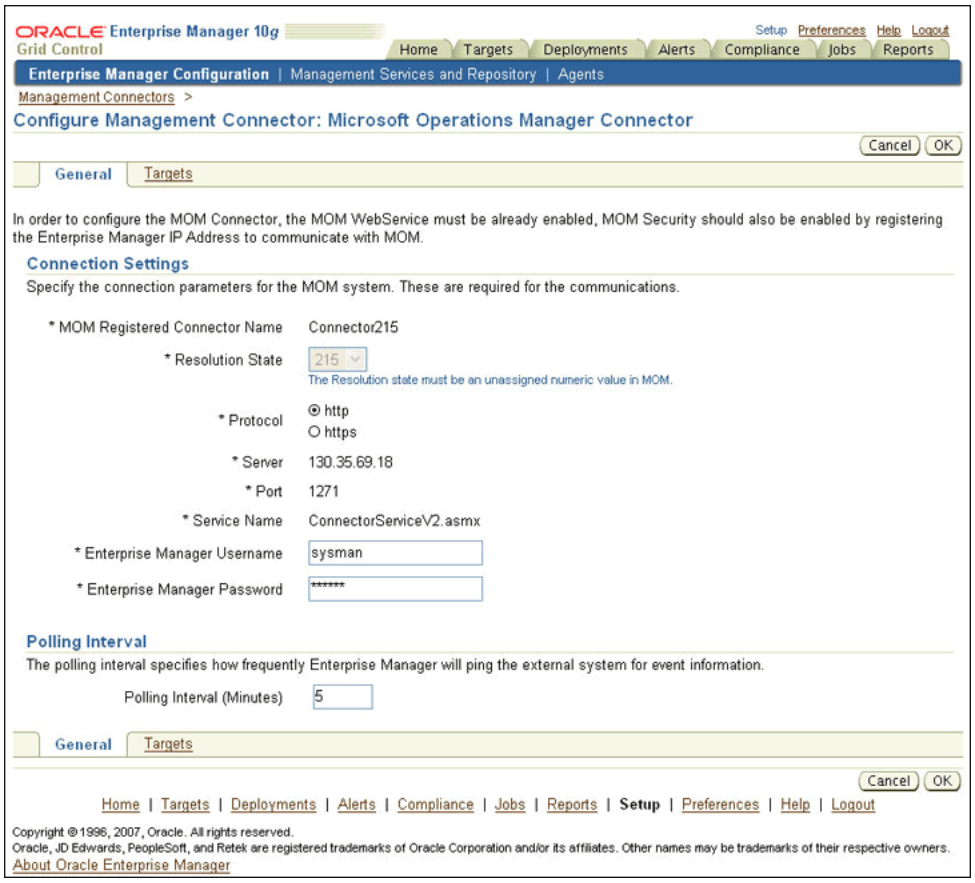

17.4.1 General Settings

This section provides the communication details required for configuring the connector.

- **MOM Registered Connector Name —** Specify the connector name that you want to register with MOM. This is the name that MOM administrator identifies for marking alerts to Enterprise Manager. Ensure that the name you provide is unique in MOM.
- **Resolution State —** Specify a value between 1 and 255. Make sure that you do not specify a value already in use including the standard values such as 85 or 255.

  The default value for MOM Connector is 218.
- **Protocol —** Select either HTTP or HTTPS based on the protocol MOM Web services are running on.
- **Server —** Enter the IP address or the DNS name of the MOM server.
- **Port —** Enter the port number used by the MOM Web server. The default is 1271.

  In the case of HTTPS connection, you must change the port to the HTTPS port enabled in MOM.
- **Service Name —** Enter the MOM Web service name. The default is ConnectorServiceV2.asmx. In most cases, you need not change this.
- **Polling Interval —** Enter the time interval between event collections.
- **Enterprise Manager Username —** Enter the username to be used to insert alerts, for example, `SYSMAN`. The user must have full privilege on the target of type MOM Managed Host.

  Administrator privileges are recommended, for instance you should be a super user such as SYSMAN.
- **Enterprise Manager Password —** Enter a password for the username you specified.

17.4.2 Additional Target Instances

When you deploy or install a MOM Connector, a default target instance `generic_mom_managed_host` is created. In addition, you can create specific target instances.

Alert gets assigned to the respective targets based on the computer name in MOM. If the name specified under `Computer` column in the MOM Operator console matches the target name you specify in the `Target Name` column in the section "Targets Managed by External System", then the alert will be assigned to that target.

If the target does not exist, then the alert will be assigned to the default target instance.

In the Configure Management Connector page, click the **Targets** tab to create specific target instances. Figure 17–3 shows the Configure Management Connector Target page.

Default Connector Target

The default target instance holds external alerts that are not associated with any particular Enterprise Manager target. Oracle recommends that you do not remove this default connector target when you create additional targets.

If the default target is removed accidentally, you can recreate it by clicking the **Enable** button in the Targets tab of the Configure Management Connector page.

If the default target is removed, for alerts that cannot be mapped to any target, error messages are logged in the file `emoms.log`.

Targets Managed by External System

(Optional) In this section, you can create target instances and allow external alerts to be associated with those target instances.

To add a new target,

1. Specify the target name in the Target Name column.
2. In the MOM HOSTNAME column, specify the fully qualified hostname (hostname along with the domain name, for example `smp-mpi2.us.oracle.com`) of the target.
3. Click **OK.**

> **Note:** If you want to create more target instances, then click **Add Rows** to add additional rows and then specify the target information. To remove a target, check the **Select** check box corresponding to the target, click **Remove,** then click **OK.**

Figure 17–3 Configure Management Connector Target Page

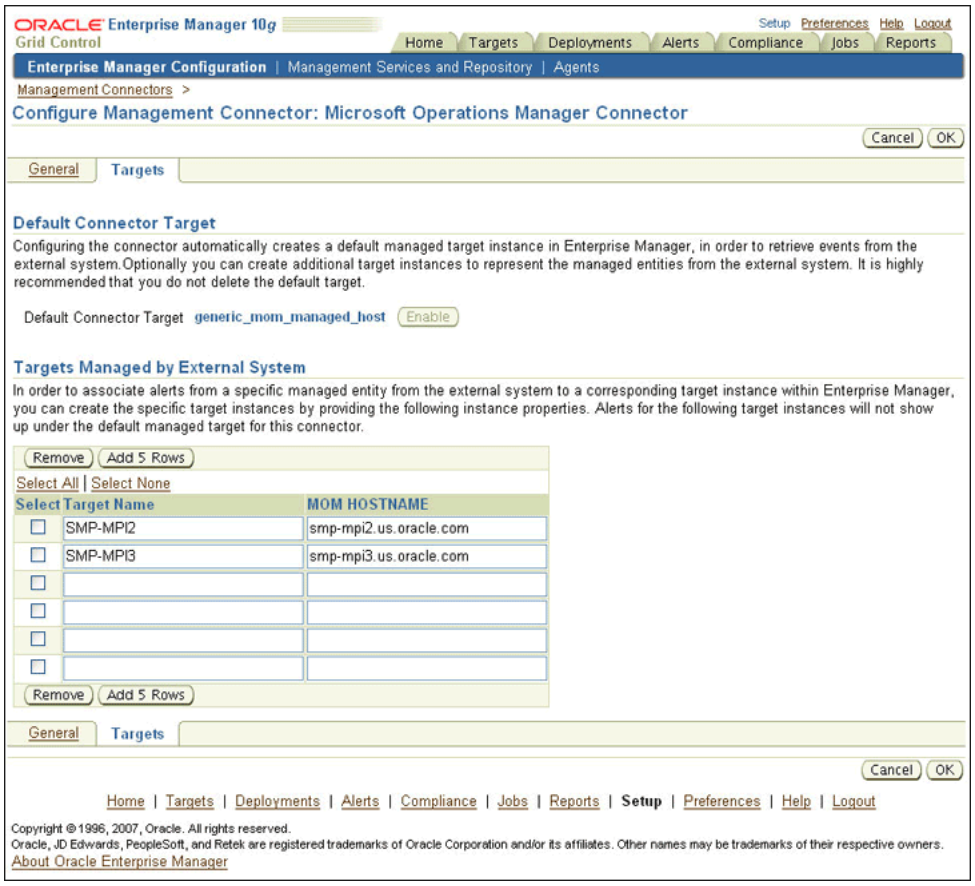

17.4.2.1 Response Status of Targets

The response status of generic_mom_managed_host asserts whether the MOM server is running. Immediately after registration, the status shows `Pending`.

After you configure the MOM Connector and the first event collection job is run (within a polling interval), the status shows either `Up` or `Down`. The response status of individual targets represent whether the host (represented by that target) is up or down.

When the target is created, Enterprise Manager pings the host. If the server is reachable, the status is marked `Up.` Subsequently, if Enterprise Manager receives any alerts (indicating that the host is down), then it marks the target as `Down.`

If the target is down, then alert retrieval is interrupted until the target is up again.

> **Note:** In Enterprise Manager, the response status of targets managed by MOM might not be definitive and can be used only for reference purposes. For accurate information of its managed entities, use MOM.

17.5 Enabling SSL Connection for HTTPS

Follow the steps provided in this section if you choose HTTPS as the protocol to establish a connection between MOM and Enterprise Manager.

17.5.1 Generating a Certificate Request File

Generate a certificate request file from the MOM server by doing the following:

1. On the Windows task bar, go to Start, then click **Run.**
2. Type `inetmgr` in the Open field.

   The Internet Information Services (IIS) Manager screen appears.
3. In the left pane, navigate to Web Sites and select the Microsoft Operations Manager 2005 connector framework.
4. Right-click and select **Properties.**

   The Microsoft Operations Manager 2005 Connector Framework Properties dialog box appears.
5. In the Directory Security tab, go to the Secure Communications section and click **Server Certificate.**
6. Using the wizard, create a new certificate.

   > **Note:** When you specify details for certificate generation, do not use abbreviated forms for city and state names as the wizard does not recognize abbreviations. For example, CA is not accepted for California.

7. At the completion of the wizard, a text file, `certreq.txt` is generated.
8. Send this request file to the Certificate authority, for example VeriSign.

17.5.2 Using the Certificate from the Certificate Authority

After processing your request, the certificate authority sends you the certificate. After you receive the certificate, do the following:

1. Paste the content in to a text file.

   The content looks like the following:


```
-----BEGIN CERTIFICATE-----
MIICdzCCAV8CAQAwDQYJKoZIhvcNAQEEBQAwITELMAkGA1UEBhMCVVMxEjAQBgNVBAMMCXJvb3Rf
dGVzdDAeFw0wNjExMTAxMDI5MzJaFw0xNjExMDcxMDI5MzJaMGYxCzAJBgNVBAYTAlVTMQ4wDAYD
VQQIEwVUZXhhczEPMA0GA1UEBxMGRGFsbGFzMQ8wDQYDVQQKEwZvcmFjbGUxCzAJBgNVBAsTAkVN
MRgwFgYDVQQDEw9zbXAtbXBpMi1vcmFwa2kwgZ8wDQYJKoZIhvcNAQEBBQADgY0AMIGJAoGBAOch
GIHp6MFW78OQw/mSdU0xfVq5u9pgqndnTqoh4aGFg1bTZD6/Azf3Nn8ibtKVJmGp3PLa3xP/gk7S
tjZ/9sM4bvnw0Y4U9xsj0BiDG4JBo35uXAUxDHLReh8F4x45Wtv/SxvE0tjNnESlBMYynLip7P9l
fSzcGKjSViyFW9M9AgMBAAEwDQYJKoZIhvcNAQEEBQADggEBAIktFTvDs7ULf0PclYXsJPeK4vFq
7HZ86omktA9lYS+oA6SaudwDGY5yxcl9O2s78o+EK9e8Wz4wM4dmUg4aSuHVHWs75W86uh7gpEFo
wssH9mtcxkqIbdPVwQoeAUTVOifNaujfXtgClvlvOjkfzvvD7SieRjD9mP2rJ2pRWUbv7xR7oJmt
RXp6t22a+MKMQQR8ofAZV/WxFJcgmBR/JxLA28X+jnzmIH/yqHK/b6Agwwy7PgbJrwPI7WQ/busm
6ASeV8ZgSfAkJ83nWz4NICnH5Y8Dyu8vDtERsOQ8z/WttrBDEmcGikkO9P+o2Y9w1pEJQhh4bKtD
PyO9YLmlrLM=
-----END CERTIFICATE-----
```


2. Save the file as `cert.cer`.
3. On the Windows task bar, go to Start, then click **Run.**
4. Type `inetmgr` in the Open field.

   The Internet Information Services (IIS) Manager screen appears.
5. In the left panel, navigate to Web Sites and select the Microsoft Operations Manager 2005 connector framework.
6. Right-click and select **Properties.**

   The Microsoft Operations Manager 2005 Connector Framework Properties dialog box appears.
7. In the Directory Security tab, go to the Secure Communications section, and click **Server Certificate.**

   Add the certificate file to the server.

17.5.3 Adding Signed Certificates to Wallet Manager

> **Note:** Oracle Wallet Manager is available at `$ORACLE_HOME/bin` on OMS. See *Oracle Application Server Administrator's Guide* for details.

1. Create a wallet using `orapki utility` by using the following command:

   ```
   orapki wallet create -wallet client -auto_login
   ```

   > **Note:** `orapki` is available at `$ORACLE_HOME/bin` on OMS.

2. Add the trusted certificate to the wallet by using the following command:

   ```
   orapki wallet add -wallet client -trusted_cert -cert
   verisignCert.cer
   ```

3. To view the content of the wallet, run the following command:

```
orapki wallet display -wallet client
```

Ensure that a file named `ewallet.p12` is available.

4. In the Oracle Wallet Manager, open the client certificate `ewallet.p12`.
5. Go to Select Trusted Certificates and select **Operations** on the main menu.
6. Select **Export All Trusted Certificates**.
7. Save the file as `certdb.txt.`
8. Place the `certdb.txt` in the connector home root directory (`$OMS_HOME/sysman/connector`).

   If the file `certdb.txt` already exists in the connector home root directory, then open the file and add the contents in your `certdb.txt` to the existing content.

Now this file can be used by the Java SSL for communication between Enterprise Manager and MOM server in HTTP mode.

See Also: For information on creating wallet, see "Creating and Viewing Oracle Wallets with orapki" in the *Oracle Database Advanced Security Administrator's Guide, 10g Release 2 (10.2).*

17.6 MOM Connector Tips

This section provides various tips that might help you in resolving issues you encounter while configuring or using the MOM Connector.

17.6.1 Recommended Protocol

Oracle recommends that you use HTTPS as the protocol for the communication between Enterprise Manager and MOM.

Use HTTP only if a secure connection is not required and the data can be transferred in clear text between the two systems.

In the case of HTTPS connection, you need to change the port to the HTTPS port enabled in MOM. Because Enterprise Manager polls data from MOM to Enterprise Manager, this means configuring MOM in HTTPS mode.

17.6.2 Connector Configuration Fails

If the connector configuration fails, then ensure the following in the MOM Administrator Console:

- Resolution state is not already in use
- Connector name is unique

17.6.3 MOM Connector Fails to Retrieve Alerts

Ensure the following if Enterprise Manager fails to retrieve alerts although MOM server is marked for forwarding alerts:

- Valid resolution state is correctly marked
- Connector is not accidentally disabled on MOM server
- Enterprise Manager username and password that you specified in the MOM Configuration page are valid.

17.6.4 Alert Logging to Additional Target Instance Fails

Even if you define additional target instances, alerts are logged to the default target instance if the target name has characters with case mis-match.

The target name is case-sensitive and therefore should match the case of the target name in the Enterprise Manager.

17.6.5 Adding Targets in the Same Transaction

You cannot delete and add the same target in the same transaction.

17.6.6 Polling Interval

The value of polling interval is in minutes. The minimum value is 5 minutes. If you specify a shorter polling interval, then it will default to 5 minutes.

17.6.7 Alerts per Polling

Processing Capability

Connector can process alerts up to twenty times the value of the polling interval at an instance.

For example, if Schedule Interval is 5, then the maximum number of alerts retrieved is 100 (20*5). This is also the default value.

This number need not be the same as the maximum number of alerts the connector is capable of requesting from MOM.

Modifying the Processing Capability

The maximum number of alerts the connector is capable of requesting from MOM depends on the number you specify in the files generic_request_newalerts.xml and generic_request_updatedalerts.xml.

For example, if you have specified 100 (`<maxCount>100</maxCount>`), then the Connector can request up to 100 alerts at an instance.

To optimize the network usage, Oracle recommends that you set the maximum value for both the processing and the requesting capabilities as the same.

17.6.8 Metric/Target Does Not Exist

Target

If target does not exist, then the alert is assigned to `generic_mom_managed_host`.

Metric

If the metric does not exist, then Enterprise Manager creates one.

Generic MOM Target

If the alert does not match any target and the generic MOM target (generic_mom_managed_host) does not exist, then the alerts are discarded and a message is logged to the log file.

18

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
- Setting Up and Configuring a Software Library With Oracle Enterprise Manager

18.1 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

18.1.1 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 10*g* 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.

18.1.2 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.

18.2 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 18.2.1, "AGENT_HOME Versus AGENT_STATE Directories" on page 18-3.

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.

18.2.1 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.

18.3 Manually Configuring a Database Target for Complete Monitoring

When you first discover an Oracle Database 10*g* 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 10*g* target.

However, when you monitor an Oracle9*i* database or an Oracle8*i* 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 Oracle9*i* 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;
   ```

   Note: The above script should not be run on an Oracle database of version 8.1.7 or prior. Oracle does not support SQL Response Time for 8.1.7 databases or prior.

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 Oracle8*i* 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.

18.4 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.

18.5 Configuring Clusters and Cluster Databases in Grid Control

This section describes how to configure clusters, cluster databases, and discovering instances.

18.5.1 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

18.5.2 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

18.5.3 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

18.5.3.1 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 18.5.2, "Configuring Cluster Databases").

For more information about configurations for Oracle Enterprise Manager Grid Control, see Chapter 3, "Grid Control Common Configurations".

18.6 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.

18.6.1 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.

18.6.1.1 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

18.6.1.2 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.

18.6.2 Configuration Parameters

The Client System Analyzer (CSA) can be further configured by modifying the context parameters in the CSA application's WAR file.

Table 18–1 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 "JInitator".	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 18–1 (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 18–1 (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 `http://java.sun.com/products/plugin/index.html#download`	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 `application/x-java-applet`.	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 18–2 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

18.6.2.1 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.

18.6.3 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 18–1 followed by a description of each tag contained in the file.

Example 18–1 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 18–1 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 18–3, 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 18–3 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 18–3 (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

18.6.4 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.

18.6.5 CSA Deployment Examples

The following sections outline sample use cases for client configurations.

18.6.5.1 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 18–2.
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 18–2 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 18–2 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 18–2, 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 18–2, 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.

18.6.5.2 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 18–2, 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.

18.6.5.3 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 18–3.

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.symsan.eml.ecm.csa.CSAResultInterface` and one that implements `oracle.sysman.eml.ecm.csa.CSACustomInteface`, both of which are shown in Example 18–3. 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 18–3 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 assocaited 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 in them in the MGMT_ECM_CSA_CUSTOM table and then writes a rule that checks these values.

18.6.5.4 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.

18.6.5.5 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.

18.6.5.5.1 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`.

18.6.5.5.2 Example 2: Inventory

In Example 18–4, 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 18–4 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".

18.6.5.5.3 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 18–5 displays the context parameter values for the CSA servlet filter on the sales portal.

Example 18–5 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 18–6 represents the context parameter definitions to map rules to collection tags.

Example 18–6 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>
```

18.7 Setting Up and Configuring a Software Library With Oracle Enterprise Manager

The following sections describe how to set up and configure a software library using Oracle Enterprise Manager.

18.7.1 Setting Up a Software Library

The software library should be located in a directory accessible by all Oracle Management Servers (OMS). If there is only one OMS, the directory can be local. For multiple OMS environments, the directory can ge on a Network File Server accessible from all Oracle Management Servers. You should ensure that there is sufficient space available on the shared storage to store file that hold the binary data for all of the components.

If you create operating system components, TAR files containing all the RPMs for a Linux installation will be stored in the software library. If you create Oracle database components, TAR files containing all the files from a reference Oracle home directory or contents from the installable media will be stored in the software library.

You should ensure that the shared storage is accessible through NFS mount points to all Oracle Management Servers in the environment.

18.7.2 Configuring a Software Library

The graphical user interface of the Provisioning application shows various tabs for Components, Directives, and Images. You can access all of the tabs depending on the privileges assigned to you. For example, if you have superuser privileges, you can access the Administration tab. The Administration tab contains different sections that you can use to configure various elements in the environment.

To configure a software library, follow these steps:

1. In the Software Library Configuration section of the Administration tab of the Provisioning application, press the **Add** button.
2. On the Add Software Library Location page, enter the directory location of the software library you are creating and then click **OK**.

18.7.3 Deleting or Cleaning Up a Software Library

To delete the software library , select the software library from the Administration tab and click Remove to delete it.

To delete the components of a software library, select the components from the software library and choose Delete. This will remove the components. To clean up the delete components completely from the File System , you must run the following command:

```
<OMS HOME>/bin/purgeDeploymentLibrary <conn string> <username>
<password> [-job <oms_host>
```

The options for this command are listed below:

<conn string> is of form: "jdbc:oracle:thin:@dbhost:dbport:sid" where dbhost, dbport and sid should be replaced appropriately

<username> is repository username

<password> is repository passwd

[-job <oms_host>] is optional; if provided a job would be submitted (schedule is immediate). The user can view the status of the job by navigating to the Jobs page on the Enterprise Manager console.

<oms_host> is OMS hostname

Index

A

accessibility
  enabling accessibility mode, 1-20
  enabling accessiblity features, 1-20
  providing textual descriptions of charts, 1-21
Additional Management Agent installation type, 1-4
advanced configuration
  introduction, 1-1
  types of tasks, 1-1
Agent Registration Password, 5-6
  changing, 5-13
Agent Upload Problems
  default notification rule, 13-7
AGENT_HOME
  definition, 1-4, 1-5
AGENT_HOME/bin, 1-5
AGENT_HOME/network/admin, 5-18
AGENT_HOME/sysman, 1-5
AGENT_HOME/sysman/admin/scripts/db/config/response.pl, 18-5
AGENT_HOME/sysman/config, 1-5
AGENT_HOME/sysman/log, 1-5
Agents Unreachable
  default notification rule, 13-7
alerts, 11-7
Application Performance Management, 5-35, 6-11
Application Server Availability and Critical/Warning States
  default notification rule, 13-7
Application Server Control
  directory structure, 1-5
  introduction, 1-5
  Ports page, 6-12
  starting and stopping, 2-7
  starting and stopping on Windows systems, 2-7
Application Service Level Management
  using to monitor the Management Service, 3-9
archive logging
  for Management Repository database, 9-1
assistive technology, 1-20
asynchronous I/O, 11-14

B

Backup, 11-16
Bad SQL
  configuring the database to show Bad SQL, 18-4
baselines, 11-5
Beacons, 11-14
  configuring firewalls to allow ICMP traffic, 6-11
  monitoring Web Applications over HTTPS, 5-35
blackouts
  controlling with emctl, 2-15
  examples, 2-17
buffer cache, 11-12

C

capacity
  predicting, 11-2
Certificate dialog box
  Internet Explorer, 5-33, 5-37
charts
  providing textual descriptions for accessibility, 1-21
collection directory, 18-2
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, 3-10
  managing multiple hosts, 3-4
  using multiple Management Services, 3-6
Configuring Services, 7-1
  Availability, 7-4, 7-6
    Beacons, 7-5
    Key Beacons, 7-5
    Local Beacon, 7-5
    Service Test-Based, 7-5, 7-6
    System-Based, 7-5, 7-6
  Command Line Interface, 7-43
  Create, 7-4
  End-User Performance Monitoring, 7-2, 7-14
    Access Log Format, 7-18
    chronos_setup.sh, 7-20, 7-22
    EUM, 7-16

Manage Web Server Data Collection, 7-17, 7-19
Set URLs, 7-17
Standalone Web Cache, 7-23
Unprocessed Samples, 7-25
URL Pattern, 7-19
Web Cache Log Format, 7-20
Web Cache Manager, 7-19
Interactive Transaction Tracing, 7-2
J2EE Server Activity, 7-2
Metrics
Performance, 7-5
Usage, 7-8
Usage Metrics, 7-5
Monitoring Settings, 7-13
Beacon Overrides, 7-13
Collection Settings, 7-13
Data Granularity, 7-13
Frequency, 7-13
Monitoring Templates, 7-39
Beacons, 7-40
Service Tests, 7-40
Service Tests and Beacons, 7-40
Variables, 7-40
Oracle Application Server 10g (9.0.4), 7-3
Performance, 7-5
Performance Metrics, 7-7
Aggregation Function, 7-7
Recording Transactions, 7-2, 7-12, 7-28
Request Performance, 7-3, 7-36
Correlate Requests, 7-3
Database Connections, 7-37
Enterprise Java Beans (EJBs), 7-37
JDeveloper, 7-39
Manage OC4J Data Collection, 7-37
OC4J Cluster, 7-36
OC4J Instances, 7-36
OC4J Tracing, 7-38
Oracle Application Server 10g (9.0.4), 7-3
Oracle User Interface XML (UIX), 7-39
Service Tests and Beacons, 7-37
Tracing Properties, 7-38
Root Cause Analysis, 7-3, 7-7, 7-11
Component Tests, 7-12
Topology, 7-11
Topology Viewer, 7-3
Service Level Rules, 7-41
Actual Service Level, 7-41
Availability, 7-41
Business Hours, 7-41
Expected Service Level, 7-41
Information Publisher, 7-42
Performance Criteria, 7-41
Services Dashboard, 7-42
Service Tests and Beacons, 7-9
Configuring Dedicated Beacons, 7-11
SSL Certificate, 7-10
Tests, 7-9
Web Proxy, 7-11
Service-Test Based Availability
Key Service Tests, 7-6
System, 7-3
Key Components, 7-4
System-Based Availability
Key Components, 7-6
Test Performance, 7-2
Thresholds
Critical, 7-5
Warning, 7-5
Time Zone, 7-4
Types
Aggregate Service, 7-14
Types of Service
Generic Service, 7-4
Types of Services
Aggregate Service, 7-4
OCS Service, 7-4
Web Application, 7-4
connect descriptor
using to identify the Management Repository database, 9-3, 9-5
creating a monitoring script, 14-2

D

data collections
how Enterprise Manager stores, 18-1
restoring default, 18-2
understanding default and custom, 18-1
Data Guard
configuring Enterprise Manager availability, 9-1
Database Availability and Critical/Warning States
default notification rule, 13-7
Database Configuration Assistant
See DBCA
Database Control
configuring after installation, 1-8, 1-11
configuring during installation, 1-8
configuring with DBCA, 1-9
configuring with EMCA, 1-11
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, 1-10
starting on Windows, 1-10
DBSNMP database user, 2-14
setting the password for, 2-14
DBSNMP user, 18-5
default_collection directory, 18-1
deployment procedure
variables, 10-8
deployment procedures
advantages, 10-1
customizable, 10-4

definition, 10-1
importing or exporting, 10-5
requirements, 10-2
submitting, 10-16
troubleshooting, 10-17
use cases, 10-3
using EMCLI to execute, 10-8
using for grid automation, 10-1
directory structure
introduction to, 1-1
Disaster Recovery, 11-17
disk mirroring and stripping
Management Repository guideline, 9-1
disk space management
controlling Management Agent disk space, 12-3
controlling the contents of trace files, 8-4
controlling the size and number of log and trace files, 8-3, 8-5, 8-6
controlling the size of log and trace files, 8-8
dontProxyFor
description of property, 6-8
property in emoms.properties, 6-8
dropping the Management Repository, 9-3

E

E2E monitoring, 11-14
EM Website
using to monitor the Management Service, 3-9
Web Application target, 3-9
em_message, 14-4
em_result, 14-3
emagent.log, 8-1
emagentlogging.properties, 8-5
log4j.rootCategory property, 8-6
MaxBackupIndex property, 8-5
MaxFileSize property, 8-5
emagent.nohup, 8-2
emagent.trc, 8-2
e-mail notifications
upper limits, 13-4
EMCA
command-line arguments, 1-12
configuring Database Control for Real Application Clusters, 1-16
configuring Database Control with, 1-11
reconfiguring Database Control after changing the listener port, 1-19
sample EMCA input file, 1-16
specifying port assignments, 1-18
troubleshooting problems with the Database Control, 1-19
troubleshooting tips, 1-19
using an input file for EMCA parameters, 1-15
emctl, 2-1
controlling blackouts, 2-15
listing targets on a managed host, 2-15
location in AGENT_HOME, 1-5
security commands, 5-6
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, 5-20, 5-27
emctl config oms sso, 5-19
emctl getemhome, 1-8
emctl istop, 2-3
emctl reload, 2-13
emctl secure agent, 5-9
sample output, 5-10
emctl secure lock, 5-12
emctl secure oms, 5-6, 5-15
sample output, 5-7
emctl secure setpwd, 5-14
emctl secure unlock, 5-12
emctl start agent, 2-2
emctl start blackout, 2-17
emctl start dbconsole, 2-8
emctl start iasconsole, 2-7
emctl start oms, 2-5
emctl status agent, 2-2
emctl status blackout, 2-17
emctl status oms, 2-5
emctl stop agent, 2-2
emctl stop blackout, 2-17
emctl stop dbconsole, 2-9
emctl stop iasconsole, 2-7
emctl stop oms, 2-5
emctl upload, 2-13
EMD_URL
property in the emd.properties file, 12-2
emd.properties, 8-3, 12-2, 12-3
EMD_URL, 12-2
emdWalletDest, 12-2
emdWalletSrcUrl, 12-2
location, 1-5
LogFileMaxRolls, 8-3
REPOSITORY_PROXYHOST, 6-5
REPOSITORY_PROXYPORT, 6-5
REPOSITORY_URL, 3-3, 3-5, 12-2
TrcFileMaxrolls, 8-3
TrcFileMaxSize, 8-3
UploadMaxBytesXML, 12-3
UploadMaxDiskUsedPct, 12-3
emdRepPort
property in the emoms.properties file, 12-9
emdRepPwd
property in the emoms.properties file, 12-9
emdRepServer
property in the emoms.properties file, 12-9
emdRepSID
property in the emoms.properties file, 12-9
emdRepUser
property in the emoms.properties file, 12-9
emdWalletDest
property in emd.properties, 12-2
emdWalletSrcUrl
property in emd.properties, 12-2

em.notification.emails_per_minute
property in emoms.properties, 13-4
em.notification.os_cmd_timeout
property in emoms.properties, 13-12
emoms.log, 8-6, 8-7
emomslogging.properties, 8-7, 8-8
MaxBackupIndex, 8-7
MaxFileSize, 8-7
emoms.properties, 9-2, 12-8
configuring the JDBC connection to the Management Repository, 3-4, 3-5
dontProxyFor property, 6-8
emdRepPort, 12-9
emdRepPwd, 12-9
emdRepServer, 12-9
emdRepSID, 12-9
emdRepUser, 12-9
em.notification.emails_per_connection, 13-3
property in emoms.properties, 13-3
em.notification.emails_per_minute, 13-4
em.notification.os_cmd_timeout, 13-12
maxInactiveTime, 18-7
oracle.net.crypto_checksum_client, 5-17
oracle.net.crypto_checksum_types_client, 5-17
oracle.net.encryption_client, 5-17
oracle.net.encryption_types_client, 5-17
oracle.sysman.eml.mntr.emdRepPwd, 9-2
oracle.sysman.eml.mntr.emdRepPwdEncrypted, 9-2
oracle.sysman.emRep.dbConn.enableEncryption, 5-17
oracle.sysman.emSDK.sec.DirectoryAuthenticationType, 5-23
oracle.sysman.emSDK.svlt.ConsoleServerPort, 12-10
proxyHost property, 6-8
proxyPort property, 6-8
sample Management Repository properties, 12-9
emoms.trc, 8-6
emwd watchdog script
in the AGENT_HOME/bin directory, 12-4
End-User Performance Monitoring
Web Server
Apache HTTP Server 2.0, 7-2, 7-3, 7-15, 7-16
Oracle Application Server Web Cache, 7-2, 7-3, 7-17
Oracle HTTP Server, 7-15
Oracle HTTP Server Based on Apache 2.0, 7-2, 7-3, 7-15
OracleAS Web Cache, 7-15
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, 5-4
compared with Oracle HTTP Server security features, 5-4
configuring, 5-4
enabling for Management Repository, 5-15
enabling for multiple Management Services, 5-11
enabling for the Management Agent, 5-9
in a firewall environment, 6-2
overview of steps required, 5-5
restricting HTTP access, 5-11
types of secure connections, 5-5
Enterprise User Security
configuring Enterprise Manager for, 5-23

F

fetchlet
log and trace files, 8-4
firewalls
between browser and the Grid Control, 6-3
between Grid Control and a managed database target, 6-10
between Management Service and Management Agents, 6-10
between Management Service and Management Repository, 6-9
configuring for ICMP traffic, 6-11
configuring for UDP traffic, 6-11
configuring the Management Agent for, 6-4
configuring the Management Service for, 6-6
configuring to allow incoming data from Management Service, 6-9
configuring to allow incoming traffic to Management Agent, 6-6
considerations before configuring, 6-1
considerations when using with multiple Management Services, 6-10

G

getemhome
emctl command, 1-8
Grid Control
architecture overview, 11-1
components, 11-1
configuring notifications, 13-1
deploying on a single host, 3-2
sizing, 11-2
starting, 2-10
starting all components of, 2-10
stopping, 2-11
stopping all components of, 2-11
guidelines
for deploying the Management Repository, 9-1

H

Host Availability and Critical/Warning States
default notification rule, 13-8
hostname
argument to the emctl secure oms command, 5-15
HTTP 500 - Internal server error, 2-6

HTTP access
restricting, 5-11
HTTP Server Availability and Critical/Warning States
default notification rule, 13-8
http_em.conf, 12-10
httpd.conf
configuring for use with a server load balancer, 3-14
Oracle HTTP Server configuration file, 3-14
HTTPS, 5-5
Hyper-Threading, 11-11

I

ICMP, 6-11
initialization parameter
adjusting when using multiple Management Services, 3-6
Internet Control Message Protocol, 6-11
Internet Explorer
Certificate dialog box, 5-33, 5-37
security alert dialog box, 5-32
Security Information dialog box, 5-35
introduction to advanced configuration, 1-1
I/O Channels
monitoring, 11-13
istop
emctl command, 2-3

J

J2EE, 5-4
directory in Oracle Management Service home, 1-4
Java Message Service (JMS), 1-19
javax.net.ssl.SSLException
SSL handshake failed, 5-36
job_queue_processes, 18-6

L

Listener Availability
default notification rule, 13-8
Listener port
obtaining, 6-10
load balancing
connections between the Grid Control Console and Management Service, 3-12
connections between the Management Agent and Management Service, 3-11
Loader, 11-11
Loader backlog (files)
on the Grid Control Management System tab, 3-9
loader threads, 11-12
log files
controlling the content of, 8-4
controlling the size and number of, 8-6
controlling the size of, 8-3
fetchlet log files, 8-4
locating and configuring, 8-1
locating Management Agent, 8-2
locating Management Service, 8-6
Management Agent, 8-1
Oracle Management Service, 8-6
rollover files, 8-2
log4j.appender.emagentlogAppender.MaxBackupIndex, 8-5
log4j.appender.emagentlogAppender.MaxFileSize, 8-5
log4j.appender.emagenttrcAppender.MaxBackupIndex, 8-5
log4j.appender.emagenttrcAppender.MaxFileSize, 8-5
log4j.appender.emlogAppender.MaxBackupIndex, 8-7
log4j.appender.emlogAppender. MaxFileSize, 8-7
log4j.appender.emtrcAppender.MaxBackupIndex, 8-7
log4j.appender.emtrcAppender. MaxFileSize, 8-7
log4j.rootCategory property in emagentlogging.properties, 8-6
log4j.rootCategory=WARN, emlogAppender, emtrcAppender, 8-8
LogFileMaxRolls property in emd.properties, 8-3
Login Timeout Value
modifying the default, 18-6
LVM (Logical Volume Manager), 9-1

M

Management Agent, 11-1, 11-4, 11-5
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, 6-6
configuring to use a proxy server, 6-5
configuring trust points, 12-7
reinstalling, 11-18
starting and stopping on UNIX, 2-1
starting and stopping on Windows, 2-2
Management Information Base (MIB), 13-29
definition, 13-29
MIB variable descriptions, 13-30
Management Options page
in DBCA, 1-9
in Oracle Universal Installer, 1-8
Management Repository
See Oracle Management Repository
Management Server, 11-10
Management Servers
adding, 11-13
Management Service, 11-1, 11-4
See Oracle Management Service
starting and stopping on Windows systems, 2-6
using a server load balancer, 3-11
master agent
Oracle Peer SNMP Master Agent service, 2-3
MaxBackupIndex
property in emomslogging.properties, 8-7
MaxBackupIndex property in

emagentlogging.properties, 8-5
MaxFileSize
property in emomslogging.properties, 8-7
MaxFileSize property in
emagentlogging.properties, 8-5
maxInactiveTime
property in emoms.properties, 18-7
MIB
See Management Information Base (MIB)
Microsoft Operations Management Connector, 17-1
see also MOM Connector, 17-1
MOM Connector
configuring, 17-2
installing, 17-1
overview, 17-1
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, 14-2
monitoring templates, 14-16

N

Netscape Navigator
New Site Certificate dialog box, 5-34
network/admin, 5-15, 5-16, 5-18
New Site Certificate dialog box
Netscape Navigator, 5-34
Notification backlog
on the Grid Control Management System tab, 3-9
Notification Methods, 13-10
notification methods
based on a PL/SQL Procedure, 13-15
based on an SNMP trap, 13-19
based on operating system commands, 13-11
notification rules
definition, 13-6
out-of-box, 13-6
out-of-the-box notification rules, 13-5
subscribing to, 13-6
notification schedules, 13-5
notification system
e-mail errors, 13-39
notification system errors, 13-37
notification system, trace messages, 13-37
notifications
assigning methods to rules, 13-27
assigning rules to methods, 13-28
configuring, 13-1
defining multiple mail servers, 13-3
long e-mail notifications, 13-4
mail server settings, 13-2
mail server settings in emoms.properties, 13-3
management information base (MIB), 13-29
notification schedules, 13-5
sample Operating System command script, 13-14
setting up, 13-1
short email notifications, 13-5
using custom notfication methods, 13-11
Notifiction Rules
Custom, 13-9

O

OC4J Availability and Critical/Warning States
default notification rule, 13-8
OEM_MONITOR, 18-5
Operating System command
sample notification method for, 13-12
sample script, 13-14
Operating System scripts, 13-11
while creating notification methods, 13-11
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, 5-7, 5-20
opmnctl status, 2-6
opmnctl stopall, 2-4, 5-6, 5-19
opmnctl stopproc ias-component=WebCache, 2-6
ORA-12645
Parameter does not exist, 5-16
Oracle Advanced Security, 5-5, 5-15, 6-9
enabling for Management Repository, 5-18
enabling for the Management Agent, 5-18
Oracle Application Server
Enterprise Manager directories installed with, 1-5
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, 2-6
starting and stopping, 2-5
starting and stopping with opmnctl, 2-6
using with Grid Control, 2-5
Web Cache Manager, 7-19
Oracle Database 10g
Enterprise Manager directories installed with, 1-6
Oracle Enterprise Manager
components, 11-4
directory structure, 1-1
log files, 8-1
rollup process, 11-12
security model, 5-1
starting and stopping Enterprise Manager
components, 2-1
Oracle Enterprise Manager 10g Grid Control
See Grid Control
Oracle Enterprise Manger
tuning, 11-10
Oracle HTTP Server
configuring for use with a server load
balancer, 3-13
Oracle Identity Management, 5-3
Oracle Internet Directory, 5-21
Oracle Management Agent

about the log and trace files, 8-1
changing the port, 12-2
configuring when protected by a firewall, 6-4
controlling disk space used by, 12-3
controlling the content of trace files, 8-4
controlling the size of log and trace files, 8-3
directory structure, 1-4
directory structure on Windows, 1-5
enabling security for, 5-9, 5-18
fetchlet log and trace files, 8-4
installing with Grid Control, 1-4
location of log and trace files, 8-2
log and trace files, 8-1
log and trace rollover files, 8-2
reconfiguring to use a new Management Service, 12-1
starting and stopping, 2-1
Watchdog process, 12-4
Oracle Management Repository, 11-2
changing the Management Repository password, 12-9
configuration parameters in the emoms.properties file, 12-9
configuring for high availability, 3-17
deploying on a remote host, 3-4
deployment guidelines, 9-1
dropping, 9-3
enabling Oracle Advanced Security, 5-18
enabling security for, 5-15
identifying with a connect descriptor, 9-3, 9-5
recreating, 9-3, 9-4
reloading data, 2-13
restoring, 11-17
starting the Management Repository database, 2-11
troubleshooting, 9-6
uploading data, 2-13
using raw devices with a RAC database, 3-18
Oracle Management Service, 11-17
about the log and trace files, 8-6
adjusting the PROCESSES initialization parameter, 3-6
bin directory, 1-3
components installed with, 1-2
configuring for use with a proxy server, 6-7
configuring to allow incoming data from Management Agent, 6-9
configuring to use a new Repository, 12-8
configuring when protected by a firewall, 6-6
determining when to use multiple Management Services, 3-8
enabling security for, 5-6
enabling security for multiple Management Services, 5-11
home directory, 1-2
j2ee directory, 1-4
location the log and trace files, 8-6
log and trace files, 8-6
modifying monitoring credentials, 2-13
monitoring the load, 3-9
monitoring the response time, 3-9
monitoring with Application Service Level Management, 3-9
opmn directory, 1-4
reconfiguring, 12-8
reconfiguring to use a new port, 12-9
restoring, 11-17
starting, stopping, and checking, 2-4
sysman directory, 1-4
tips for monitoring the load and response time, 3-8
using multiple management services, 3-6
Oracle Net firewall proxy access, 6-9, 6-10
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 Technology Network (OTN), 7-23
Oracle Universal Installer
Management Options page, 1-8
ORACLE_HOME/bin, 1-3
ORACLE_HOME/hostname_sid/, 1-6
ORACLE_HOME/install, 6-12
ORACLE_HOME/j2ee, 1-4
ORACLE_HOME/network/admin, 5-15, 5-16, 5-18
ORACLE_HOME/oc4j/j2ee, 1-7
ORACLE_HOME/oc4j/j2ee/OC4J_DBConsole, 1-7
ORACLE_HOME/opmn, 1-4
ORACLE_HOME/opmn/bin, 2-6
ORACLE_HOME/sysman, 1-4, 1-6
Oracle8i database
configuring for monitoring, 18-4
Oracle9i
configuring for monitoring, 18-4
oracle.net.crypto_checksum_client
property in emoms.properties, 5-17
oracle.net.crypto_checksum_types_client
property in emoms.properties, 5-17
oracle.net.encryption_client
property in emoms.properties, 5-17
oracle.net.encryption_types_client
property in emoms.properties, 5-17
oracle.sysman.eml.mntr.emdRepPwd
property in emoms.properties, 9-2
oracle.sysman.eml.mntr.emdRepPwdEncrypted
property in emoms.properties, 9-2
oracle.sysman.emRep.dbConn.enableEncryption
entry in emoms.properties, 5-17
oracle.sysman.emSDK.sec.DirectoryAuthenticationType
property in emoms.properties, 5-23
oracle.sysman.emSDK.svlt.ConsoleServerPort
property in emoms.properties file, 12-10
OS scripts
See Operating System scripts
OTN (Oracle Technology Network), 7-23
OUIinventories.add, 18-3

P

password
  changing the Management Repository password, 12-9
  changing the SYSMAN password, 9-2
peer encapsulator service
  SNMP, 2-3
Performance Metrics
  Beacon Aggregation Function
    Average, 7-7, 7-8
    Maximum, 7-7
    Minimum, 7-7, 7-8
    Sum, 7-7, 7-8
  System Aggregation Function
    Maximum, 7-8
PERFSTAT, 18-5
PL/SQL procedures, 13-11
  sample, 13-18
  while creating a notification method, 13-15
  while creating notfication methods, 13-11
Pluggable Authentication Modules, 10-7
portlist.ini, 6-12
ports
  4888, 6-7, 6-9
  4889, 6-9, 12-9
  changing the Management Agent port, 12-2
  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, 6-12
  portlist.ini, 6-12
  reconfiguring Database Control after changing the listener port, 1-19
  reconfiguring the port used by the Management Service, 12-9
  specifying Database Control ports, 1-18
  viewing a summary of ports assigned during installation, 6-12
PROCESSES, 3-6
ProcessManager
  service used to control the Management Service on Windows systems, 2-7
proxy server
  configuring Management Agent for, 6-5
  configuring the Management Service for, 6-7
proxyHost
  property in emoms.properties, 6-8
proxyPort
  property in emoms.properties, 6-8
Public Key Infrastructure (PKI), 5-5, 5-36

R

RAID-capable disk
  Management Repository guideline, 9-1
raw devices
  when used with the Management Repository, 3-18
Recovery, 11-16
Remedy Connector, 16-1
  configuring, 16-3
  installing, 16-3
  overview, 16-1
Remote Method Invocation (RMI), 1-18
RepManager script, 9-3, 9-4
Repository Operations Availability
  default notification rule, 13-8
REPOSITORY_PROXYHOST
  property in emd.properties, 6-5
REPOSITORY_PROXYPORT
  property in emd.properties, 6-5
REPOSITORY_URL
  property in emd.properties, 3-3, 3-5
  property in the emd.properties file, 12-2
rollover files, 8-2
rollup process, 11-12
Root Cause Analysis
  Mode
    Automatic, 7-11
    Manual, 7-11
root password
  See also SYSMAN
  when enabling security for the Management Service, 5-6

S

scalability
  determining when to use multiple Management Services, 3-8
screen readers, 1-20
script registration, UDM, 14-5
script results, returning, 14-2
security
  about Enterprise Manager security, 5-1
  authorization and access enforcement, 5-2
  classes of users and their privileges, 5-2
  Enterprise Manager security model, 5-1
  leveraging Oracle Application Server security services, 5-3
  leveraging Oracle Identity Management Infrastructure, 5-3
  overview of steps required to enable Enterprise Manager Framework Security, 5-5
  See also Enterprise Manager Framework Security
security alert dialog box
  Internet Explorer, 5-32
security certificate alerts
  responding to, 5-32
security features
  See Enterprise Manager Framework Security
Security Information dialog box
  Internet Explorer, 5-35
self-monitoring
  feature of the Management Agent, 12-4
Server Connection Hung
  error while creating the repository, 9-6
Server Load Balancer, 5-15

server load balancer, 3-11
configuring a virtual pool, 3-13
configuring a virtual service, 3-13
modifying the httpd.conf file, 3-14
modifying the ssl.conf file, 3-14
using with Management Services, 3-11
using with the Grid Control Console, 3-12
ServerName directive
in the httpd.conf file, 3-14
in the ssl.conf file, 3-14
Service Tests and Beacons
Tests
DNS, 7-9
FTP, 7-9
SOAP, 7-9
Web Transaction, 7-9
Services control panel
using to start and stop the Management Agent, 2-8, 2-9
using to start the Management Service, 2-6
session timeout
modifying, 18-6
setupinfo.txt, 2-6
Single Sign-On
bypassing the Single Sign-On logon page, 5-23
configuring Enterprise Manager with, 5-19
registering Enterprise Manager as a partner application, 5-22
registering Single Sign-On users as Enterprise Manager administrators, 5-21
using Single Sign-On to authorize Enterprise Manager users, 5-19
SNMP
Oracle Peer SNMP Master Agent service, 2-3
Oracle SNMP Peer Encapsulator service, 2-3
SNMP traps, 13-11, 13-19
sample, 13-20
Software Library, 15-1
configuring, 15-2
creating and deleting entities, 15-2
de-configuring, 15-6
export script, 15-4
exporting and importing entities, 15-3
import script, 15-4
purging, 15-7
setting up, 15-2
SQLNET.CRYPTO_SEED
entry in sqlnet.ora, 5-18
SQLNET.ENCRYPTION_SERVER
entry in sqlnet.ora, 5-18
sqlnet.ora, 5-15, 5-16
SQLNET.CRYPTO_SEED, 5-18
SQLNET.ENCRYPTION_SERVER, 5-18
ssl.conf
configuring for use with a server load balancer, 3-14
Oracle HTTP Server configuration file, 3-14
starting and stopping
Enterprise Manager components, 2-1
state directory
in the Management Agent home, 12-2
Statspack, 18-4
Status Codes, Corrective Actions, 13-24, 13-26
SYSMAN
changing the SYSMAN password, 9-2
checking for existence of, 9-6
entering SYSMAN password when enabling security, 5-6
sysman/admin/default_collection, 18-2
sysman/emd/collection, 18-2
system errors, notification, 13-37

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, 11-5, 11-7
Top SQL Report
configuring the database to show the Top SQL Report, 18-4
trace files
component tracing levels, 8-4
controlling the content of, 8-4
controlling the contents of Management Service, 8-8
controlling the size and number of, 8-6
controlling the size of, 8-3
fetchlet trace files, 8-4
locating Management Agent, 8-2
locating Management Service, 8-6
Management Agent, 8-1
Oracle Management Service, 8-6
rollover files, 8-2
TrcFileMaxRolls property in emd.properties, 8-3
TrcFileMaxSize property in emd.properties, 8-3
Troubleshooting
when using EMCA, 1-19
troubleshooting
general techniques while creating the Management Repository, 9-6
problems starting or configuring the Database Control, 1-19
while creating the Management Repository, 9-5
with EMCA, 1-19
troubleshooting, notifications, 13-36
trust points
Management Agent Configuration, 12-7

U

UDP, 6-11
uix-config.xml, 1-20
upload directory

in the Management Agent home, 12-2, 12-3
UploadMaxBytesXML
property in the emd.properties file, 12-3
UploadMaxDiskUsedPct
property in the emd.properties file, 12-3
Usage Metrics
Aggregation Function
Average, 7-8, 7-9
Maximum, 7-8, 7-9
Minimum, 7-8, 7-9
Sum, 7-8, 7-9
User Datagram Protocol, 6-11
User-defined metric page
Command Line, 14-6
Comparison Operator, 14-7, 14-12
Consecutive Occurrences Preceding Notification, 14-8, 14-13
Critical, 14-8, 14-13
Environment, 14-7
Metric Name, 14-6, 14-11
Metric Type, 14-6, 14-11
Operating System User Name and Password, 14-7
Response Action, 14-8, 14-13
Warning, 14-7, 14-12
User-defined metric page, Central Console, 14-5, 14-10
user-defined metric, example, 14-8
User-defined metrics, 14-1
user-defined metrics, 14-16

V

virtual pool
when configuring a server load balancer, 3-13
virtual service
when configuring a server load balancer, 3-13

W

watchdog process
for the Management Agent, 12-4
Web Application
Source
Step, 7-7
Step Group, 7-7
Transaction, 7-7
Web Application target
using to monitor the Management Service response time, 3-9
Web Applications
monitoring over HTTPS, 5-35
Web Cache
See Oracle Application Server Web Cache
Web Cache Availability and Critical/Warning States
default notification rule, 13-9
web.xml, 1-21