

Oracle® Enterprise Manager

Administrator's Guide

Release 9.2.0

March 2002

Part No. A96670-01

ORACLE®

Copyright © 1996, 2002 Oracle Corporation. All rights reserved.

The Programs (which include both the software and documentation) contain proprietary information of Oracle Corporation; 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. Oracle Corporation does not warrant that this document is 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, without the express written permission of Oracle Corporation.

If the Programs are delivered to the U.S. Government or anyone licensing or using the programs on behalf of the U.S. Government, the following notice is applicable:

Restricted Rights Notice Programs delivered subject to the DOD FAR Supplement are "commercial computer software" and use, duplication, and disclosure of the Programs, including documentation, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement. Otherwise, Programs delivered subject to the Federal Acquisition Regulations are "restricted computer software" and use, duplication, and disclosure of the Programs shall be subject to the restrictions in FAR 52.227-19, Commercial Computer Software - Restricted Rights (June, 1987). Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065.

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

Oracle is a registered trademark, and Oracle Store, Oracle8i, Oracle9i, PL/SQL, SQL*Net, and SQL*Plus are trademarks or registered trademarks of Oracle Corporation. Other names may be trademarks of their respective owners.

Contents

Send Us Your Comments	xvii
Preface.....	xix
Part I Oracle Enterprise Manager Console	
1 Console	
Enterprise Manager Console	1-2
Components.....	1-2
Navigator.....	1-2
Detail Pane	1-2
Console Menus	1-3
Context-Sensitive Menus.....	1-4
File Menu.....	1-4
Object Menu	1-4
Configuration Menu	1-5
Tools Menu.....	1-6
Help Menu	1-7
Toolbars/Tool Drawers.....	1-7
Managing Enterprise Manager Administrators	1-8
Types of Enterprise Manager Administrators	1-8
Managing Administrator Accounts	1-9
Creating or Editing an Administrator Account	1-10
Granting Access to Targets	1-11

Reassigning Object Ownership.....	1-13
Changing Administrator Preferences.....	1-13
General.....	1-13
Notification	1-14
Filters	1-15
Email.....	1-16
Paging	1-18
Paging Status Codes for Numeric Pages	1-20
Email/Paging Message Body Definition	1-20
Schedule.....	1-22
Access	1-23
Preferred Credentials	1-25
Setting Administrator Credentials	1-27
Saving Logon Information as a Preferred Credential	1-28
Saving the Contents of a List.....	1-28
Configuring Enhanced Notifications (Paging/Email).....	1-29
Management Regions	1-29
Defining a New Management Region	1-29
Adding a Management Server to a Region.....	1-30
Adding Discovered Nodes to a Management Region	1-31
Removing a Management Region	1-32
Managing HTTP Servers	1-32
Discovering HTTP Servers.....	1-32
Starting Up or Shutting Down a HTTP Server.....	1-33
Determining the Status of HTTP Servers	1-33
Paging/Email Blackout	1-34
Specifying Total Paging/Email Blackout	1-34
Defining a Paging/Email Blackout	1-34
Making a Copy of an Existing Blackout Schedule.....	1-36
Turning Blackout Schedules On and Off.....	1-36
Deleting Paging/Email Blackout Periods	1-37
Viewing Blackout Periods	1-37
Target-level Blackouts.....	1-37

2 The Standalone Console

Choosing to Launch the Console Standalone	2-2
Starting the Standalone Console	2-2
Adding Databases to the Tree in the Standalone Console	2-4
Connecting to a Database in the Standalone Console	2-6
Connecting to the Database As a Different User	2-7
Viewing If You Are Connected As SYSDBA	2-7
Removing a Database from the Tree	2-7
Changing from Using the Console to the Standalone Console	2-8
Editing Local Preferred Credentials in the Standalone Console	2-9
Changing from Using the Standalone Console to the Console	2-11

3 Navigator

Navigator Pane	3-2
Populating the Navigator Tree.....	3-2
Expanding Objects in the Navigator.....	3-2
Launching Tools.....	3-4
Navigator Menu	3-4
Discovering Targets	3-6
Adding a Service	3-8
Refreshing a Target	3-8
Removing a Target.....	3-8
Pinging the Intelligent Agent.....	3-8
Determining Node Properties	3-9
Manipulating Objects in the Navigator	3-9
Administering Objects.....	3-9
Copying Navigator Objects	3-9
Removing a Node from the Navigator	3-10
Node Removal Failure	3-10

4 Groups

Group View Pane	4-2
Managing Groups	4-3
Automatic Group Refresh.....	4-3

Manual Refresh	4-4
Automatic Refresh.....	4-4
Creating a Group.....	4-4
Group General Page	4-6
Group Access Page.....	4-7
Manipulating Group Views	4-8
Monitoring Status	4-8
Expanding Objects.....	4-9
Groups	4-10
Databases and Other Discovered Targets.....	4-10
Launching Applications from a Group.....	4-10
Adding Objects to a Group	4-10
Deleting Objects from the Group	4-10
Modifying a Group	4-11
Removing Groups.....	4-11
Viewing Reports for Targets within a Group	4-12

5 Jobs

Job Process	5-2
Job Tasks	5-2
Writing SQL*Plus Scripts	5-3
Job Credentials.....	5-4
Submitting Jobs	5-5
Cancelling Jobs.....	5-6
Job Detail View.....	5-6
Active Job Page	5-7
Context-sensitive Menu options	5-9
History Page.....	5-9
Refreshing the History Page.....	5-10
Clearing the History Page.....	5-10
Displaying Job Output.....	5-10
Context Menu options	5-10
Job Menu	5-11
Job Library.....	5-12
Creating, Modifying, or Viewing a Job.....	5-12

Creating a New Job	5-13
General Page.....	5-13
Job Task Page.....	5-15
Job Parameters Page.....	5-17
Job Schedule Page.....	5-18
Job Access Page	5-20
Skipped Job Notification.....	5-22
Job Progress Page.....	5-22
Job Output Dialog Box.....	5-23
Modifying Active Jobs	5-24
Alternative Method of Modifying an Active Job	5-25
Viewing Job Details	5-25
Example: Creating a Job.....	5-26
Required Administrator Permissions	5-28
Oracle Job Tasks	5-29
Oracle Database Tasks	5-29
Enterprise Manager Wizard Database Tasks	5-29
Operating System or Node Tasks	5-33
Tcl Script Examples	5-34
Listener Tasks.....	5-35
HTTP Servers	5-36
Job Tasks Run through Wizards.....	5-36

6 Events

Event System Overview.....	6-2
Using Events.....	6-2
Creating Events	6-3
Registering Events.....	6-4
Event Occurrences.....	6-5
Event Notifications.....	6-5
Notifying Administrators	6-6
Interpreting Events	6-6
Correcting Problems.....	6-8
Event Categories and Types.....	6-8
Fault Management Event Tests	6-9

Space Management Event Tests	6-9
Resource Management Event Tests.....	6-10
Performance Management Event Tests	6-10
Unsolicited Event Tests	6-10
Registering Interest in an Unsolicited Event	6-11
Setting the Parameters Property Sheet for Unsolicited Events	6-12
Raising Unsolicited Events.....	6-13
Raising Unsolicited Events through the Enterprise Manager Job System	6-13
Unsolicited Event Caveats.....	6-19
User-Defined Monitoring	6-20
User-Defined SQL Event Test.....	6-20
Support for PL/SQL Functions.....	6-20
User-Defined Event Tests.....	6-22
Creating Your Monitoring Script	6-22
Register the user-defined event in the Console	6-25
User-Defined Event Parameters	6-26
Output	6-28
Bundled User-Defined Event Sample	6-29
Creating and Registering an Event	6-31
Dynamic Modification of Registered Events	6-32
General Behavior.....	6-35
Event Detail View	6-36
Alerts Page.....	6-37
Viewing Alerts.....	6-38
History Page.....	6-38
Registered Page	6-39
Event Menu	6-40
Context-Sensitive Menus.....	6-41
Event Library Dialog	6-41
Editing an Event in the Event Library.....	6-42
Oracle Event Tests.....	6-42
Event Viewer	6-42
Event Viewer: General Page.....	6-43
Event Viewer: Log Page.....	6-44
Event Viewer: Notification Details Page.....	6-44

Responding to Event Occurrences	6-45
Event General Page	6-46
Event Tests Page.....	6-48
Event Parameters Page	6-49
Parameters	6-49
Event Schedule Page	6-50
Event Access Page.....	6-51
Event Fixit Jobs Page	6-54
Event Progress Page	6-55
Administrator Event Notification	6-56
Oracle Event Tests	6-56
Numeric Pager Job/Event Ids	6-57
Event System Features and Requirements	6-57

7 Event Handler

Event Handler Overview	7-2
How the Event Handler Works	7-3
Setting Up the Event Handler	7-4
Quickstart Method (Default)	7-4
Customizing the Event Handler Setup	7-5
Event Handler Configuration Parameters.....	7-7
Blackouts.....	7-7
Filters.....	7-7
Templates	7-10
Command to Execute	7-15
Optional: length of execution time.....	7-15
Summary of Event Handler Configuration Commands	7-17
Enabling the Event Handler	7-17
Disabling the Event Handler	7-17
Viewing Current Event Handler Configuration Settings	7-17
Creating a Configuration File from the Current Event Handler Configuration Registry Entries	7-18
Importing a Configuration File	7-18
Troubleshooting Tips.....	7-18
Sample Filters and Templates:	7-20

Filters:.....	7-20
Templates	7-21
Known Issues	7-23
Differences between UNIX and Windows NT	7-23
Running the Event Handler in a multi-Management Server Environment	7-23
Migrating from Prior Releases	7-24

8 Enterprise Manager Reporting

Enterprise Manager Reporting	8-2
Key Concepts	8-2
What is a Report Definition?	8-2
What is a Report Element?	8-4
Ways to Select Targets for Report Generation.....	8-5
Configuring Enterprise Manager Reporting	8-5
Enterprise Manager Reporting Website	8-6
Creating a Report from an Existing Report Definition	8-9
Editing a Report Definition	8-10
Generating a Report from Enterprise Manager Applications	8-10
Creating a User-defined Report Definition	8-11
The Report Property Sheet.....	8-12
Report General Page.....	8-12
Report Elements Page.....	8-15
Report Parameters Page.....	8-24
Report Publish Page	8-25
Navigating the Enterprise Manager Reporting Website	8-27

9 Enterprise Security Management

Overview of Enterprise Security Manager	9-2
Introduction to Directory Servers	9-3
Entailing and Configuring Your Enterprise Security Environment	9-5
Task 1: Configure an Oracle Internet Directory	9-5
Task 2: Install Oracle Enterprise Manager	9-5
Task 3: Configure Oracle Enterprise Manager for Enterprise User Security	9-6
Task 4: Start Oracle Enterprise Security Manager	9-6
Task 5: Log On To the Directory	9-7

Administering Users	9-8
Oracle Wallets	9-8
Specifying a new User Name	9-9
Specifying a Directory Base	9-10
Specifying a new User Password	9-11
Specifying an Initial Enterprise Role Assignment	9-12
Specifying an Oracle Wallet	9-14
Browsing Users in the Directory	9-15
Administering Oracle Contexts	9-20
Oracle Context Versions	9-20
Specifying Properties of an Oracle Context	9-20
Specifying User Search Bases	9-22
Specifying Oracle Context Administrators	9-23
Accessible Domains	9-27
Managing Database Security	9-28
Registering a Database with an Oracle Context	9-28
Administering Databases	9-31
Managing Database Administrators	9-31
Managing Database Schema Mappings	9-31
Administering Enterprise Domains	9-34
Specifying Database Membership of an Enterprise Domain	9-36
Managing Database Security Options for an Enterprise Domain	9-39
Managing Enterprise Domain Administrators	9-39
Managing Enterprise Domain Database Schema Mappings	9-40
Administering Enterprise Roles	9-41
Creating a new Enterprise Role:	9-42
Removing an Enterprise Role:	9-44
Specifying Database Global Role Membership of an Enterprise Role	9-44
Removing a Database Global Role from an Enterprise Role:	9-46
Adding a Global Role to an Enterprise Role:	9-46
Managing Enterprise Role Grantees	9-48
Removing a User from the List of Enterprise Role Grantees:	9-48
Adding a New User to the list of Enterprise Role Grantees:	9-49
Command Line Tool	9-50

Part II Database Administration Tools

10 Database Administration

Common Features of Database Management Features	10-3
Tree Views.....	10-3
General Information about Databases	10-3
Comprehensive Overview Pages	10-3
Property Sheets	10-3
Multi-Column Lists	10-3
Database Version Awareness	10-3
Database Reports	10-4
Logging of Database Changes	10-4
Showing Object DDL	10-4
Show SQL.....	10-4
Show Dependencies.....	10-4
Right-Mouse Commands.....	10-4
DB Search Capabilities.....	10-5
Database Management Features and Wizards	10-6
Instance Management	10-7
Configuration Operations	10-8
Stored Configurations.....	10-10
Sessions List	10-11
Sessions Folder.....	10-11
Long Running Operations	10-11
Locks	10-11
In-Doubt Transactions	10-12
Resource Consumer Groups	10-12
Resource Plans.....	10-12
Resource Plan Schedule	10-13
Schema Management.....	10-14
Tree List by Schema or Object.....	10-15
Editing an Object.....	10-18
Creating Objects	10-18
Security Management	10-19
User Operations.....	10-19

Role Operations	10-20
Profile Operations	10-20
Storage Management	10-21
Controlfile Operations	10-23
Tablespace Operations	10-23
Datafile Operations	10-23
Rollback Segment Operations	10-23
Redo Log Group Operations	10-24
Archive Log Operations	10-24
Distributed Management	10-24
Warehouse Management	10-25
OLAP Management	10-26
Summary Management	10-27
Workspace Management	10-27
XML Database	10-28
SQL*Plus Worksheet	10-29
SQL Scratchpad	10-30
Wizards	10-32

11 Managing Backup and Recovery

Introduction	11-1
Recovery Manager (RMAN)	11-2
About the RMAN Executable	11-4
About the Target Database	11-4
About Oracle Enterprise Manager	11-4
About the Recovery Catalog Database	11-6
About the Recovery Catalog Schema	11-6
About the Standby Database	11-7
About the RMAN Media Management Interface	11-8
About the Media Management Catalog	11-8
RMAN Backup	11-8
Backing Up a Database	11-9
Backing Up a Database Using a Predefined Strategy	11-9
Backing Up the Database with a Customized Strategy	11-11
Deleting Obsolete Backups and Copies	11-12

Selecting a Full or Incremental Backup	11-12
Choosing an Online or Offline Mode to Back Up Your Database	11-15
Backing Up Individual Files	11-16
Backing Up and Deleting Archived Logs	11-19
Backing Up Archived Logs	11-19
Deleting Archived Logs	11-20
Copying a Datafile	11-22
Overriding a Retention Policy in Backup for Special Cases	11-24
Overriding the Backup Policy	11-25
Overriding the Retention Policy	11-26
Viewing Current Policies	11-26
Restore and Recover	11-28
Recovering the Entire Database	11-29
Restoring the Entire Database	11-33
Recovering and/or Restoring Tablespaces or Datafiles	11-34
Restoring the Control File	11-38
Restoring Archive Log Files	11-40
Recovering Datablocks	11-42
Using a Corruption List for Data Block Recovery	11-44
Using Datafiles for Data Block Recovery	11-44
Using Tablespaces for Data Block Recovery	11-45
Maintenance Operations	11-46
Setting up Backup and Retention Policies in a Target Database	11-46
Configuring a Backup Policy	11-47
Configuring a Retention Policy	11-49
Performing Recovery Catalog Maintenance	11-50
Registering a Database	11-50
Resynchronizing the Catalog	11-51
Resetting the Database	11-51
Configuring the RMAN Environment	11-52
Creating a Backup Configuration	11-52
Specifying Disk Channel Device for Backup Set	11-54
Specifying Tape Channel Device for Backup Set	11-56
Setting Channel Limits	11-57
Specifying Channel Device for an Image Copy	11-58

Setting Up a Proxy Copy for a Tape Backup Set	11-60
Setting the Storage Parameters for the Current Backup Set.....	11-61
Setting Storage Parameters for Datafiles	11-61
Setting Storage Parameters for Archivelogs	11-62
Registering the Recovery Catalog	11-62
Setting Preferred Credentials for Running Backup Jobs.....	11-64
Registering Later Databases with the Recovery Catalog	11-64
Resynchronizing the Recovery Catalog with the Target Database	11-65
Setting Up the Recovery Catalog	11-65
9i Procedure.....	11-66
Pre-8i and 8i Procedures	11-66
Starting Up the Database	11-68
Starting the Database in Mount Mode.....	11-69
Starting the Database in Open Mode.....	11-71
Placing the Database in Mount.....	11-72
Setting the Database in ARCHIVELOG Mode	11-74
RMAN Job Script	11-75
Some Examples	11-76
Backing Up the Recovery Catalog	11-76
Recovery Without a Catalog	11-77
Skipping Tablespaces when Backing Up a Database.....	11-78
Backing Up Often-Used Tablespaces.....	11-78
Specifying the Device Type	11-79
Restarting a Backup.....	11-79
Spreading a Backup Across Multiple Disk Drives.....	11-79
Specifying the Size of Backup Sets.....	11-80
Performing a Non-Cumulative Incremental Backups	11-80
Performing Cumulative Incremental Backups	11-80
Determining How Channels Distribute a Backup Workload	11-81
Backing Up in NOARCHIVELOG Mode.....	11-81
Keeping a Long-Term Backup	11-81
Displaying Backups that are Exempt from the Retention Policy	11-81
Optimizing Database Backup	11-82

A Keyboard Navigation

B Firewalls and Virtual Private Networks

Firewall Communication for Enterprise Manager	B-2
Firewall Between the Console and Management Server	B-2
Firewall Between the Management Server and Agent(s) on Monitored Nodes	B-4
Firewalls and Network Address Translation (NAT)	B-5
Virtual Private Network Configuration for Enterprise Manager	B-6
VPN Connections Between the Enterprise Manager Client and Management Server	B-7
VPN Connections Between the Management Server and Intelligent Agents	B-9
Running the Console in Standalone Mode	B-9
Performance Manager, Capacity Planner, and Firewalls	B-10

C OEMUTIL Utility

Starting OEMUTIL	C-2
Using OEMUTIL	C-2
Performing a single command:	C-2
Performing Commands in Succession:	C-2
OEMUTIL Commands	C-3

D Repository Views

Administrator Views	D-2
Service Views	D-3
Group Views	D-8
Operational System Metrics	D-10
Job Definition Views	D-14
Event Definition Views	D-25
Capacity Planner and Data Gatherer Collections	D-33

Index

Send Us Your Comments

Oracle Enterprise Manager Administrator's Guide, Release 9.2.0

Part No. A96670-01

Oracle Corporation welcomes your comments and suggestions on the quality and usefulness of this document. Your input is an important part of the information used for revision.

- Did you find any errors?
- Is the information clearly presented?
- Do you need more information? If so, where?
- Are the examples correct? Do you need more examples?
- What features did you like most?

If you find any errors or have any other suggestions for improvement, please indicate the document title and part number, and the chapter, section, and page number (if available). You can send comments to us in the following ways:

- FAX: (650) 506-7200 Attn: Oracle System Management Products
- Postal service:
Oracle Corporation
Oracle System Management Products Documentation Manager
500 Oracle Parkway, 5OP5
Redwood Shores, CA 94065
USA

If you would like a reply, please give your name, address, telephone number, and (optionally) electronic mail address.

If you have problems with the software, please contact your local Oracle Support Services.

Preface

This chapter describes the purpose and organization of this guide. The following information is discussed:

- Purpose of this Guide
- Audience
- How this Guide Is Organized
- How to Use This Guide
- Conventions Used in This Guide
- Documentation Set
- Related Publications
- Documentation Accessibility

Purpose of this Guide

This guide describes Oracle® Enterprise Manager, Release 9.0.2: Oracle's system management Console, common services, and integrated platform graphical tools. Oracle Enterprise Manager also provides an integrated set of standard database administration applications to help automate and simplify the common daily tasks of administrators. These supplemental applications focus on specific areas of database administration, helping administrators with their daily and routine tasks of managing databases and other services and keeping them operational.

From the client interface, the Oracle Enterprise Manager Console, you can do the following:

- Centrally administer, diagnose, and tune multiple databases
- Manage non-database Oracle products and services
- Effectively monitor and respond to the health of your Oracle family of products and third-party services 24 hours a day
- Schedule jobs on multiple nodes at varying time intervals
- Monitor networked services for events
- Customize your display by organizing databases and other targets into logical administrative groups
- Generate custom and predefined reports for monitored targets in your enterprise and optionally publish the reports to a webserver.

While using Enterprise Manager products, you should refer to the online help for specific information on the displayed dialog box, menu, or window. You can display the online help by pressing F1 or selecting a Help button if present.

For an overview of the Oracle Enterprise Manager system, see the *Oracle Enterprise Manager Concepts Guide*.

Audience

This guide is written for those who wish to use Oracle Enterprise Manager to perform system administration tasks.

This guide assumes you are familiar with the administrative tasks you wish to perform. If you are not, refer to the Oracle server documentation set. The Oracle server documentation set contains specific and thorough descriptions of the database administration tasks you can perform with Enterprise Manager applications. In addition, the Oracle server documentation set provides recommendations on how to administer your database optimally. If you have not yet read the introductory chapters of the Oracle server administrator's guide, we recommend that you do so. These chapters describe the specific responsibilities of a database administrator.

You should also be familiar with the operation of your specific Microsoft Windows or UNIX system. See the documentation for your Windows or UNIX system, if necessary.

How this Guide Is Organized

This guide is divided into chapters as described below.

Part I: *Oracle Enterprise Manager Console*

Chapter 1, "Console"

This chapter describes the Enterprise Manager Console, its basic configuration and functions, as well as system administrative functions such as paging, email, and managing administrators.

Chapter 2, "The Standalone Console"

This chapter introduces the option of launching the Console standalone, thus allowing a single administrator to perform simple database schema, instance, storage, security, and other database tasks by connecting directly to the target database(s) without using the Management Server or Intelligent Agents on target machines.

Chapter 3, "Navigator"

This chapter explains the Navigator component of the Enterprise Manager Console, associated menus, and managing objects within the Navigator such as discovering services, managing web servers, and generating reports.

Chapter 4, "Groups"

This chapter explains the Group component of the Enterprise Manager Console and how using Groups can simplify administering your enterprise.

Chapter 5, "Jobs"

This chapter explains how to use the Job Scheduling component of Enterprise Manager to automate administrative tasks.

Chapter 6, "Events"

This chapter explains how to use the Event Management component of Enterprise Manager to monitor managed nodes and services.

Chapter 7, "Event Handler"

This chapter explains how to use the Event Handler to monitor for specific event conditions and have Enterprise Manager respond to them automatically.

Chapter 8, "Enterprise Manager Reporting"

This chapter explains how to use Enterprise Manager's integrated reporting functionality and how to publish generated reports to an HTTP server.

Chapter 9, "Enterprise Security Management"

This chapter explains how to use Enterprise Security Manager to administer enterprise user security for the Advanced Security Option.

Part II: *Database Administration Tools*

Chapter 10, "Database Administration"

This chapter introduces the database management functionality, which is an integrated set of standard database administration applications to help automate and simplify the common daily tasks of administrators. It also describes the user interface elements used in the database administration applications.

Chapter 11, "Managing Backup and Recovery"

This chapter explains how to use the Oracle Enterprise Manager Backup Management wizards to administer your database backup and recovery environment.

How to Use This Guide

The *Oracle Enterprise Manager Administrator's Guide* has been designed to be used closely with the Oracle Server documentation set. While this guide describes how to use Enterprise Manager to perform database administration tasks, the Oracle Server documentation set describes the reasons for and the implications of performing these tasks. Consequently, you should refer to the Oracle Server documentation set while using Enterprise Manager to perform your administrative tasks.

For an overview of the Enterprise Manager system, see the *Oracle Enterprise Manager Concepts Guide*. After reading this manual, you may choose to proceed directly to those chapters that are relevant to the tasks you plan to perform using Enterprise Manager.

Before using the database administration applications, you should read Chapter 10, "Database Administration". This chapter provides an overview of the organization and user interface elements of the applications.

While using the Enterprise Manager products, you should refer to the online help for specific information on the displayed dialog box, menu, or window. You can display the online help by pressing F1 or selecting a Help button if present.

Conventions Used in This Guide

The following sections explain the conventions used in this guide.

Examples

This guide contains code examples. Note that the text of examples appears in a different font than the text of the guide. This is an example of a SELECT statement:

```
SELECT * FROM emp
```

Examples in this guide follow these case conventions:

- Keywords, such as CREATE and NUMBER, appear in uppercase. Keywords have special meanings. When you specify them, they can be in uppercase or lowercase, but they must be used exactly as they appear in the code example.
- Names of database objects and their parts, such as emp and empno, appear in lowercase. However, in the text of this guide, names of database objects and their parts appear in uppercase.
- Parameters act as place holders in examples. They appear in lowercase. Parameters are usually names of schema objects, Oracle datatypes, or

expressions. When you see a parameter in a syntax diagram, you should substitute an object or expression of the appropriate type. Note that parameter names appear in italics in the text of this guide.

Command Syntax

- *Italics* is used for variables, such as *oem_tool*
- | denotes alternative choices
- {*param1* | *param2* | ... } signifies that one of the parameters in {} must be used
- [] identifies optional parameters

Special Text

Special text is provided to alert you to particular information within the body of this guide and within other manuals.

Note: This indicates important information related to Enterprise Manager.

Additional Information: Where necessary, this refers you to your operating system-specific Oracle documentation for additional information.

Attention: This highlights information that is important when performing the described task.

Suggestion: This signifies suggestions and practical hints that can be helpful when using Enterprise Manager.

Warning: This indicates information that you should be aware of before you perform the action described in the current section.

Documentation Set

The Oracle Enterprise Manager Release 9i documentation includes the following:

- The *Oracle Enterprise Manager Readme Release 9i* provides important notes on updates to the software and other late-breaking news, as well as any differences between the product's behavior and how it is documented.
- The *Oracle Enterprise Manager Configuration Guide Release 9i* provides information about configuring the Oracle Enterprise Manager system.
- The *Oracle Enterprise Manager Concepts Guide Release 9i* provides an overview of the Enterprise Manager system.
- The *Oracle Enterprise Manager Administrator's Guide Release 9i* describes the components and features of the Oracle Enterprise Manager system.
- The *Oracle Intelligent Agent User's Guide* describes how to administer the Oracle Intelligent Agent.
- The *Oracle Enterprise Manager Messages Manual Release 9i* contains probable causes and recommended actions for Oracle Enterprise Manager errors.
- The *Oracle Enterprise Manager Event Test Reference Manual* contains detailed descriptions of all available event tests.

In addition to the Oracle Enterprise Manager documentation set, extensive on-line help is provided for components in Oracle Enterprise Manager.

To download free release notes or installation documentation, please visit the Oracle Documentation Center at <http://docs.oracle.com/>

Printed documentation is available for sale in the Oracle Store at <http://oraclestore.oracle.com/>

Related Publications

The *Oracle Enterprise Manager Administrator's Guide* refers to important information in related publications. Depending on the version of the Oracle database, you would refer to the appropriate release. The related books referred to in this guide are listed below:

- For general information about the Oracle9i and how it works, see *Oracle9i Database New Features* and *Oracle9i Database Concepts*.
- For information about administering Oracle9i, see the *Oracle9i Database Administrator's Guide*.
- For information on Oracle's SQL commands and functions, see the *Oracle9i SQL Reference*.
- For information about Oracle messages and codes, refer to *Oracle9i Database Error Messages*.
- For information about the Oracle networking system, see your network-specific documentation.
- For information about the Oracle9i for other platforms, see your platform-specific documentation.
- For information on Oracle 9i Real Application Clusters, refer to the *Oracle9i Real Application Clusters Administration*, which provides essential information for using Oracle9i with Oracle 9i Real Application Clusters and Oracle Enterprise Manager as well as a conceptual and component overview of Oracle 9i Real Application Clusters.
- For information on preparing and planning your Oracle 9i Real Application Clusters installation and configuration in an Oracle9i environment, refer to *Oracle9i Real Application Clusters Setup and Configuration*.

Oracle documentation is available online from the Internet at

<http://tahiti.oracle.com>

In North America, printed documentation is available for sale in the Oracle Store at

<http://oraclestore.oracle.com/>

Customers in Europe, the Middle East, and Africa (EMEA) can purchase documentation from

<http://www.oraclebookshop.com/>

Other customers can contact their Oracle representative to purchase printed documentation.

To download free release notes, installation documentation, white papers, or other collateral, please visit the Oracle Technology Network (OTN). You must register online before using OTN; registration is free and can be done at

<http://otn.oracle.com/admin/account/membership.html>

If you already have a username and password for OTN, then you can go directly to the documentation section of the OTN Web site at

<http://otn.oracle.com/docs/index.htm>

Documentation Accessibility

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

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

Accessibility of Code Examples in Documentation

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

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle Corporation does not own or control. Oracle Corporation neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Support for Hearing and Speech Impaired Customers

Oracle Corporation 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 technical questions, call:

1.800.446.2398

For non-technical questions, call:

1.800.464.2330

Part I

Oracle Enterprise Manager Console

- Chapter 1, "Console"
- Chapter 2, "The Standalone Console"
- Chapter 3, "Navigator"
- Chapter 4, "Groups"
- Chapter 5, "Jobs"
- Chapter 6, "Events"
- Chapter 7, "Event Handler"
- Chapter 8, "Enterprise Manager Reporting"
- Chapter 9, "Enterprise Security Management"

This chapter introduces the Oracle Enterprise Manager Console and provides an overview of its components. The following topics are discussed in this chapter.

- Enterprise Manager Console
- Managing Enterprise Manager Administrators
- Configuring Enhanced Notifications (Paging/Email)
- Managing HTTP Servers
- Paging/Email Blackout

Enterprise Manager Console

The Console is the primary interface used for all Oracle Enterprise Manager operations. It provides menus, toolbars, and the framework to access Oracle tools and utilities in addition to those available through other vendors. The graphical configuration of the Console and the tools available are determined by the optional products installed and user preferences. See Figure 1-1, "Console Window" for an illustration of a Console screen.

Components

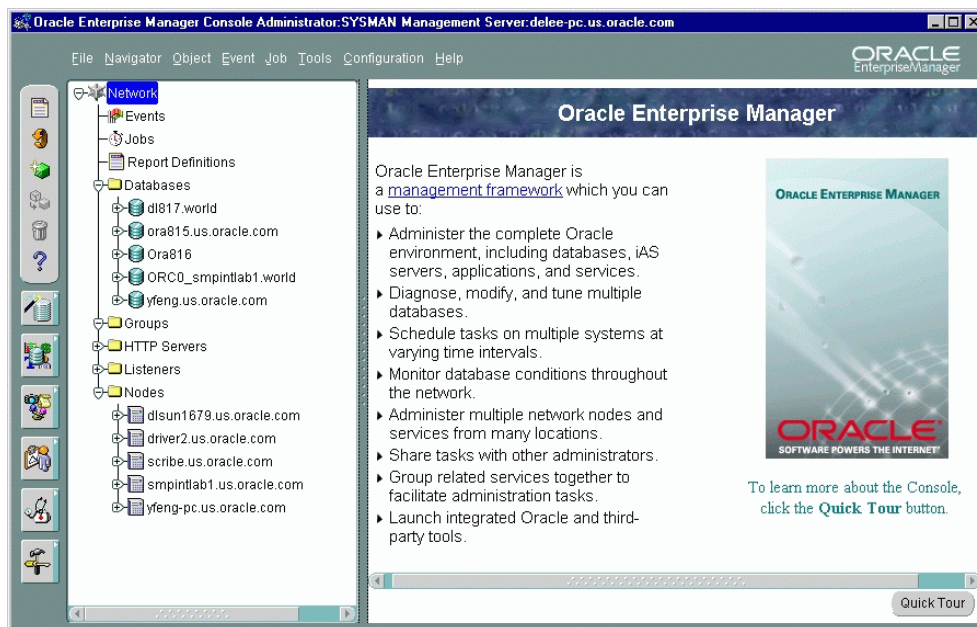
The Console uses a master/detail configuration to provide an integrated, consistent, and efficient way of managing your enterprise environment. When an object in the Navigator (master) is selected, the appropriate interface is displayed on the right-hand side of the Console (the detail pane). Information displayed on the right-hand side of the Console can be a multi-column list, property sheet, or overview page.

Navigator

The Navigator displays a hierarchical list of all the targets in a network, providing a direct view of targets such as databases, groups, listeners, and nodes, plus the objects that they contain. The Navigator shows all the network targets with their relationships to other objects. Objects within the Navigator can be managed via context-sensitive menus. For example, you can also perform many administration tasks from the Navigator, such as creating, editing, or dropping tablespaces. See Chapter 3, "Navigator" for information.

Detail Pane

The Detail pane displays information or functionality relevant to whatever object is selected in the Navigator. Some elements that can be displayed are multi-column lists summarizing attributes or contents of an object, property sheets providing object information and any functionality relating to the object, or comprehensive overview pages with links to related Quick Tours or functional overviews.

Figure 1–1 Console Window

Console Menus

The Console menu bar provides access to the following menus:

Figure 1–2 Console Main Menu

File Navigator Object Event Job Tools Configuration Help

The File, Object, Tools, Configuration, and Help menus are described in this section.

- See "Navigator Menu" on page 3-4 for information about the Navigator menu.
- See "Job Menu" on page 5-11 for information about the Job menu.
- See "Event Menu" on page 6-40 for information about the Event menu.

When using the Console menus, note the following:

- Some menus include other menus. For example, when you select the Database Applications from the Tools menu, a sub-menu containing several options displays.
- Menu options vary depending on the objects selected in the Navigator or the active window in the Console such as the Groups View page. The Tools menu items vary depending on the components that have been installed.

Context-Sensitive Menus

You can click the right mouse button on objects in the Navigator or some windows of the Console to display a context-sensitive. This menu usually contains a subset of the options that are available through a menu in the main menu bar, or functions that are specific to a selected object in the Navigator. For example, if you click the right mouse button on a group in the Navigator, a menu appears with the menu options from the Object menu. Dialogs and property sheets displaying objects in general use the context-sensitive menus for most, if not all operations that can be performed on an individual object.

File Menu

The File menu items allow you to exit the system.

Object Menu

The Object menu provides object-specific functionality for an object selected in the Console Navigator. Though menu options change depending on the object, five base menu options always appear.

Create

Displays the Create Object dialog allowing you to create Navigator objects such as jobs, events, database objects, and report definitions.

Create like

Displays the object's Create/Edit property sheet allowing you to create a new object based on the original object's parameter settings.

View/Edit Details

Displays the object's Edit property sheet allowing you to modify the selected object's parameters.

Delete

Deletes the selected object from the Navigator.

View Published Reports

Displays the Enterprise Manager reporting home page.

Configuration Menu

The Configuration menu provides options to set up administrator and system configurations.

Preferences...

Displays the Edit Administrator Preferences property sheet, which allows the current administrator to change preferences, including access levels and login credentials. See "Changing Administrator Preferences" on page 1-13 for more information.

Add Services to tnsnames.ora

Updates the local network configuration file (tnsnames.ora) with discovered services information.

Font Settings

Changes the font settings for the Console as well as any applications launched from the Console.

Font settings in pop-up menus will not change even if you have specified to change the font settings.

Report Data Purge Options

Displays the Report Data Purge Options dialog. This dialog allows you to set the purge policy for the report data log. The report data log contains information that is collected over time and used in service level reports and enterprise reports. The report data log is stored in one or more tables within the Enterprise Manager repository. All event occurrences and service level response time tasks are logged to the report data log.

SQL Logging

Allows you to log SQL in a permanent log file. SQL Logging allows you to specify the SQL logging file information and options. When you are connected to the Management Server, you can query the log (View SQL Log).

View SQL Log

Available only when you are connected to a Management Server, this menu item displays the SQL Log Viewer, which allows you to view and purge contents of the log.

Configure Paging/Email

Allows you to set up enhanced notification (paging/email) systems for administrators. This menu option is only available to super administrators.

Manage Administrators

Displays the Administrator Manager Accounts property sheet for adding, modifying, and removing administrators. This option is only available to Super Administrators. See "Managing Enterprise Manager Administrators" on page 1-8 for more information.

Grant Access to Targets

Displays the Access to Targets dialog, which allows super administrators to customize Navigator views for all regular administrators. See "Granting Access to Targets" on page 1-11 for more information.

Set Paging/Email Blackout

Displays the Paging/Email Blackout dialog. Paging/Email Blackout allows an administrator with super administrator privileges to suspend paging and email notifications for specified targets and/or services that have been previously discovered in the Navigator. See "Paging/Email Blackout" on page 1-34 for more information.

View Reporting Website Configuration

Displays the current Reporting webserver configuration. An error message is displayed if no webserver has been configured.

Define Management Regions

Displays the Management Regions property sheet. The Management Regions feature allows super administrators to partition the targets of a particular repository and assign them to a subset of the available Management Servers.

Tools Menu

The Tools menu allows you to execute database applications and other utilities that have been installed on your system. The menu options in the Tools menu depend on your installed Oracle Enterprise Manager configuration. The standard complement of options includes:

- Database Wizards
- Application Management
- Change Management Pack
- Database Applications
- Diagnostics Pack
- Tuning Pack

Other optional applications may also be available from this menu.

Help Menu

From the Help menu, you can access the Enterprise Manager online help system. You can also access Enterprise Manager's Quick Tours. Enterprise Manager comes with a variety of Quick Tours, each providing a comprehensive overview for Enterprise Manager and specific integrated applications. Enterprise Manager version information is also available.

Toolbars/Tool Drawers

The Console toolbar, located along the upper-left side of the Console, allows you to access basic manipulation functions for objects in the Navigator and the Console in general. Some functions of the toolbar will be unavailable depending on the object selected in the Navigator. Object functionality are (top to bottom):

- View Enterprise Manager reports published on a web server
- Refresh Navigator views
- Create a new object
- Create a new object based on an existing object
- Remove an object
- Display the master contents page of the help system

Figure 1–3 Console Toolbar and Tool Drawers



The Console tool drawers, located below the toolbar, provides quick and easy access to a wide variety of integrated applications. Placing the cursor over any drawer, or content of that drawer displays the application name.

Managing Enterprise Manager Administrators

Enterprise Manager is a multi-administrator system: every person who is administering systems using Enterprise Manager has their own administrator account which they use to log into the Console.

Types of Enterprise Manager Administrators

The installation of Enterprise Manager creates two super administrator: *reports_user* and *sysman*. The *reports_user* super administrator owns all predefined reports of the reporting system. See "Enterprise Manager Reporting" on page 8-1 for more information. The super administrator *sysman* creates administrators using the Manage Administrators option in the Console Configuration menu. In addition to an administrator name and password, each account can be tagged as a "Super Administrator" account or an account to which the administrator has access to only jobs and/or events. Differences between the two types of accounts are as follows:

- **Super Administrators** automatically have full privileges for all objects in the system. To provide greater security, only Super Administrators can discover, refresh, or remove targets from the Console Navigator.

Most Super Administrators also have a separate account for daily operations but use their Super Administrator account for special operations only available to Super Administrators, such as starting and stopping the Management Server, creating new Enterprise Manager Administrators, configuring paging servers, or checking other administrator's schedules. Using the `sysman` account for daily administration work is not recommended. The Super Administrator account is similar to `root` on UNIX or Administrator on Windows NT and is a user which cannot be deleted or renamed. It is a user that can perform any task and therefore should be used only for setting up the environment. After the necessary DBA accounts are created, the Super Administrator account 'SYSMAN' should not be used anymore.

- **Regular Administrators** can have access to a subset of Console operations and will only see or be able to modify those jobs, events, or groups to which they have been granted access by the Super Administrator or other regular administrators. See "Granting Access to Targets" on page 1-11.

Typically, all administrators share a single Enterprise Manager repository, which allows administrators to share information. Although you can set up multiple repositories, administrators using different repositories will not have access to each other's information; there is no sharing of data between repositories.

Administrative data stored in the repository is filtered based on administrator permissions.

Preferred Credentials must be set up for each administrator account. When an administrator connects through the Navigator, the preferred credentials used are those defined explicitly for that administrator.

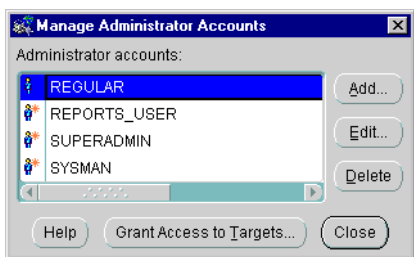
Managing Administrator Accounts

Enterprise Manager administrators are created, edited, and deleted with the Manage Administrators option of the Console Configuration menu. The Manage Administrators option displays the Manage Administrator Accounts dialog (Figure 1-4, "Manage Administrator Accounts"). The Manage Administrators option is only available to Super Administrators. Super Administrators have full privileges for all objects and can create, edit, and delete other administrators. The Super Administrator sets up the administrator name and initial password and determines whether an administrator is a Super Administrator or regular administrator. The Super Administrator also determines whether the administrator has access to the job and event systems.

In this dialog, you can:

- Click the Add button in the Manage Administrator Accounts dialog to display the Create Administrator Account dialog. The Administrator Account dialog allows you to add a new administrator. See "Creating or Editing an Administrator Account" on page 1-10 for more information.
- Click the Delete button to delete an administrator. If the administrator is the owner of any objects, you need to reassign object ownership otherwise the objects are deleted with the administrator. See "Reassigning Object Ownership" on page 1-13 for more information.
- Click Grant Access to Targets to control what targets appear in the Navigator for regular administrators. See "Granting Access to Targets" on page 1-11 for more information.

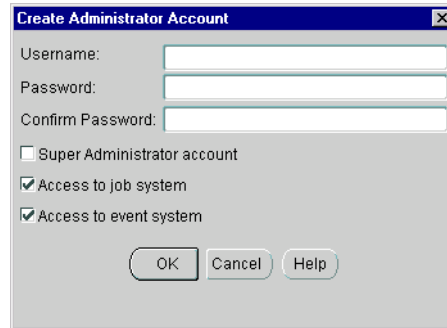
Figure 1–4 *Manage Administrator Accounts*



Creating or Editing an Administrator Account

To add a new administrator, click Add in the Manage Administrator Accounts dialog to display the Create Administrators dialog.

In the Create Administrators dialog, enter a unique administrator name and password for the Enterprise Manager administrator. Note that the administrator name (account login used by Enterprise Manager) is not the same as a database username (account login used by the database).

Figure 1-5 Create Administrator Account

The screenshot shows a dialog box titled "Create Administrator Account". It has three text input fields labeled "Username:", "Password:", and "Confirm Password:". Below these fields are three checkboxes: "Super Administrator account" (unchecked), "Access to job system" (checked), and "Access to event system" (checked). At the bottom of the dialog are three buttons: "OK", "Cancel", and "Help".

Check the access available to the administrator:

Super Administrator

Allows the administrator to add other administrators, as well as access to all created objects, e.g., jobs, events, and groups.

Access to Job System

Allows the administrator access to the Job system, but does not automatically give access privileges for specific jobs.

Access to Event System

Allows the administrator access to the Event system.

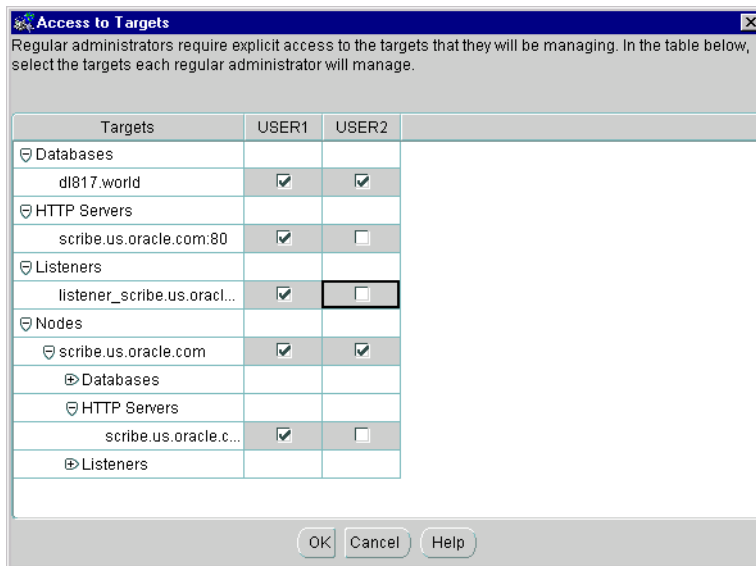
After an administrator has been created, the administrator can log in to the Enterprise Manager Console and set up administrator preferences.

If you click the Edit button, the Edit Administrator Preferences window displays. You can then edit the preferences for the selected administrator.

Granting Access to Targets

Enterprise Manager Super Administrators can control what regular administrators see in the Console Navigator through the Access to Targets dialog. This dialog can be displayed by either choosing Grant Access to Targets from the Console's Configuration menu, or clicking Grant Access to Targets from the Manage Administrator Accounts dialog. See "Managing Administrator Accounts" on page 1-9.

Figure 1–6 Access to Targets Dialog



The ability to control what targets and services appear to different users helps simplify maintaining your enterprise. For example, if you are managing 100 targets in your enterprise, and Administrator A's responsibilities require that he manage a only 10 targets, the Super Administrator can specify that only the 10 managed targets appear in the Navigator when Administrator A logs into the Enterprise Manager Console. Again, only Super Administrators can customize Navigator views.

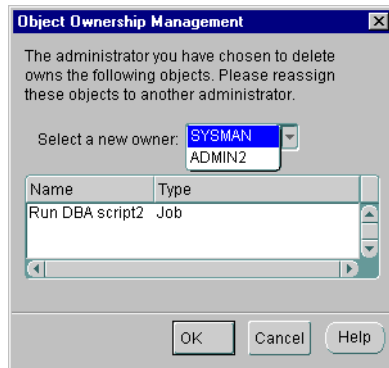
To customize the Navigator for regular administrators:

1. From the Console's Configuration menu, choose Manage Administrators. The Manage Administrators Accounts dialog appears.
2. Click Grant Access to Targets. The Access to Targets dialog appears. The Targets column can be expanded to display lower level objects in the Navigator.
3. Under each administrator, click all checkboxes that correspond to the objects you want to appear in the Console Navigator when that user logs in.
4. Click OK when finished to close the dialog.
5. Click Close to close the Manage Administrator Accounts dialog.

Reassigning Object Ownership

If you attempt to delete an administrator who owns any object, the Object Ownership Management dialog appears. See Figure 1–7, "Object Ownership Management". The dialog allows you to reassign the objects to another Enterprise Manager administrator. If you do not reassign the objects, the objects are deleted along with the administrator.

Figure 1–7 Object Ownership Management



Changing Administrator Preferences

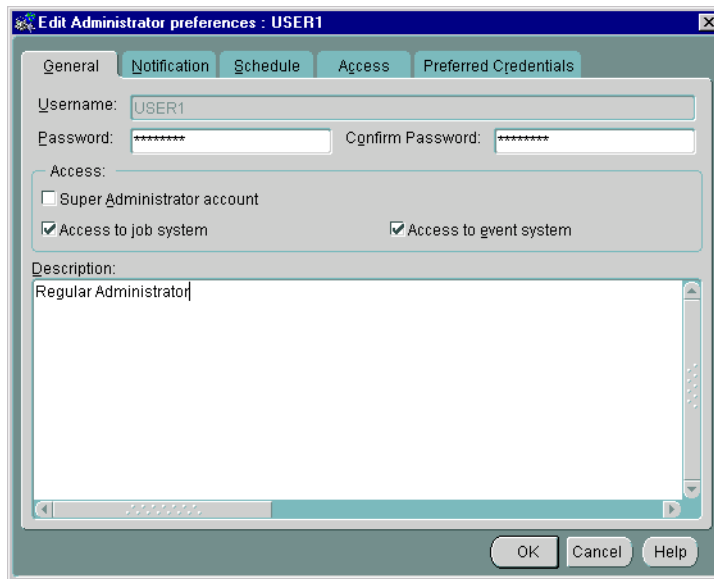
The Edit Administrator Preferences property sheet is displayed with the Preferences option of the Configuration menu and allows you to modify the administrator preferences of the user who is currently logged into the Console. The pages in the property sheet are:

- General
- Notification
- Schedule
- Access
- Preferred Credentials

General

The General Preferences page allows a Super Administrator to change an administrator's password and access privileges.

Figure 1–8 Administrator General Preferences

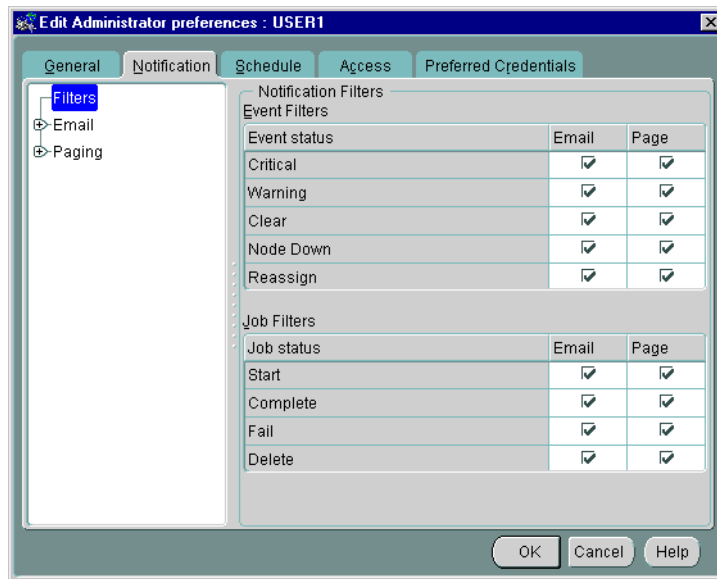


Notification

The Notification page allows you to set up paging and email notification methods for the administrator. This page consists of a hierarchical tree list and a right-side modal area that changes according to the object selected in the navigator. The tree list consists of three top-level objects:

- Filters (Selected by default)
- Email
- Paging

Figure 1–9 Administrator Notification Preferences



Filters

Notification filters allow each administrator to specify when to send an email notification as opposed to a page notification as a result of a job or event status change.

- **Event Notification Filter:** The Event Notification Filter allows you to filter email/pages sent to an administrator according to the event's level of severity. Filtering is set at the user level by checking or unchecking the Email/Page options on this property sheet page. You can select any combination of the following levels of event severity. For example, you can set filtering so that an administrator is notified via email if there is a warning and notified via the paging system if there is an alert. Selecting all levels of severity provides no filtering.
 - *Critical*
 - *Warning*
 - *Clear*
 - *Node Down*
 - *Reassign*

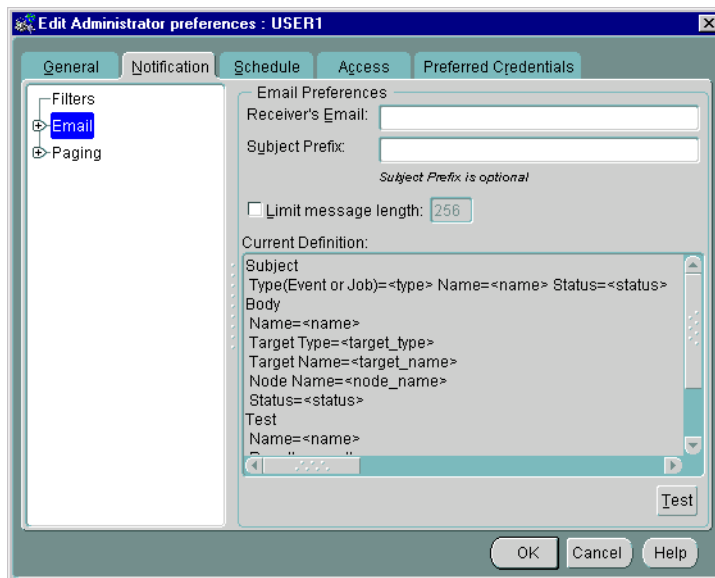
Job Notification Filter The Job Notification Filter allows you to filter email/pages sent to an administrator according to job status. As with the Event Notification Filter, filtering is set at the user level by checking or unchecking the desired option. Selecting all job statuses provides no filtering.

- *Start*
- *Complete*
- *Fail*
- *Delete (Job removed)*

Email

This page allows you to specify notification parameters for email.

Figure 1–10 Administrator Notification Preferences: Email



- **Receiver's Email:** Enter the email address of the administrator to whom the mail is being sent. You can add multiple email addresses by using commas, spaces, or tabs as delimiters between each address.

- **Subject Prefix:** Enter an optional prefix that is appended to the E-mail subject. That allows administrators to quickly identify messages from Enterprise Manager in their mail.
- **Limit message length:** Allows you to specify the maximum message body length of an E-mail notification. By default, this length is unlimited.
- **Current Definition:** Displays the current settings for format and content of E-mail notifications. To set or change the default parameters, expand the Email object in the tree list. You can select format and content options for the Subject line and Message Body. You can further expand the Message Body object to specify options at the Per Event Test level.
- **Test:** Click the Test button to check the validity of E-mail configuration. A message displays to inform you of the status of the test email.

Email Subject This page allows you to select the content and format of the email subject line. By default, Type, Name, and Status are selected. You use the left/right arrows to move items back and forth between the Available and Selected lists.

Once an item has been selected, you can control the order in which the item appears by selecting it in the list and using the up/down arrows (located to the immediate right of the Selected list) to reposition the item within the list.

Note: The following information also applies to inserting content into the Email/Paging message body.

- **Available:** Lists available content.
- **Selected:** Lists currently selected content and the order in which the content should be displayed.
- **Use Abbreviated Format:** For Type (Job or Event), Status, and Target Type, you can select the Use Abbreviated Format option. When selected abbreviations are used in the Subject line. The following abbreviations are used by the system

Table 1–1 Target Type Abbreviations

Abbreviation	Target
DB	Database

Table 1–1 Target Type Abbreviations

Abbreviation	Target
LSR	Listener
N	Node

Table 1–2 Status Type Abbreviations

Abbreviation	Status
ST	Started
C	Completed
F	Failed
D	Deleted
W	Warning
CR	Critical
CL	Cleared
NU	Node Unreachable
AC	Assignee Changed

Timestamp Format

This option is available if Timestamp is chosen from the Selected list. You use the pull-down menu to select one of the pre-defined formats.


Do not include titles

When selected, classification titles such as Status, Service Name, and Timestamp are omitted from the message.

Paging

This page allows you to specify notification parameters for paging.

Figure 1–11 Administrator Notification Preferences: Paging



- **Carrier:** After a paging server is specified by a super administrator and the name of the paging carrier service and the corresponding paging carriers have been configured, select the name of the paging carrier service from the pull-down list.
- **PIN:** Enter the PIN for your paging carrier. This entry is only required for alphanumeric pagers. Enterprise Manager does not provide PIN support for numeric pagers.
- **Limit message length:** Allows you to specify the maximum message length of paging notifications. By default, this length is set to 256 characters.
- **Current Definition:** Displays the current settings for format and content of paging notifications. To set or change these parameters, expand the Paging object in the tree list. You can select format and content options for the Message Body. You can further expand the Message Body object to specify options at the Per Event Test level.
- **Test:** Click the Test button to check the validity of the paging configuration. For alphanumeric pagers, enter the PIN number. For numeric pagers 700, is sent.

Click on the Send button to send a test page to the specified pager. A message informs you of the status of the test page.

If the test fails, check the log file. If tracing is enabled for paging, you can also view the paging trace log file in the ORACLE_HOME\sysman\log directory on the machine running the paging server.

Paging Status Codes for Numeric Pages

Numeric pages need to be interpreted as follows:

For job notifications, you will receive a 3 digit number which indicates the job status.

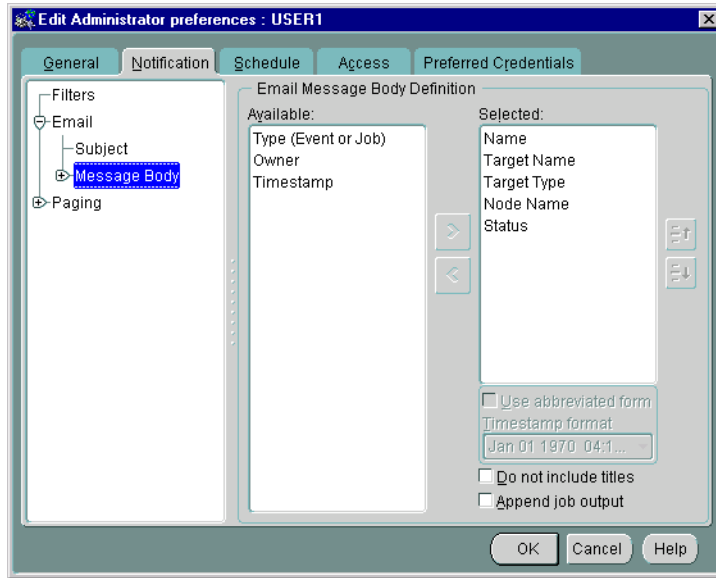
For event notifications you will receive the a 3 or 4 digit number indicating the event status.

The event status and job status codes are listed as follows:

- 100 = Job Started
- 200 = Job Completed
- 300 = Job Failed
- 400 = Job Deleted
- 500 = Event Cleared
- 600 = Event Warning
- 700 = Event Critical
- 800 = Event Node Down
- 900 = Event Unknown
- 1000 = Event Assignee Changed

Email/Paging Message Body Definition

This page allows you to select the content and format for the body of the email or page message. By default, Name, and Status are selected. You use the left/right arrows to move items back and forth between the Available and Selected lists.

Figure 1–12 Administrator Notification Preferences: Email/Paging Body

Expanding the Message Body object in the navigator and selecting per Test allows you to use a subset of the following option settings on a per test basis. See "Email Subject" on page 1-17 for more information on manipulating message content.

- **Available:** Lists available content.
- **Selected:** Lists currently selected content.
- **Use Abbreviated Format:** For Type (Job or Event), Status, and Service Type, you can select the Use Abbreviated Format option. When selected abbreviations are used in the Subject line. See Table 1–1, "Target Type Abbreviations" and Table 1–2, "Status Type Abbreviations".

Timestamp Format

This option is available if Timestamp is chosen from the Selected list. You use the pull-down menu to select one of the pre-defined formats.

Do not include titles

When selected, classification titles such as Status, Service Name, and Timestamp are omitted from the message.

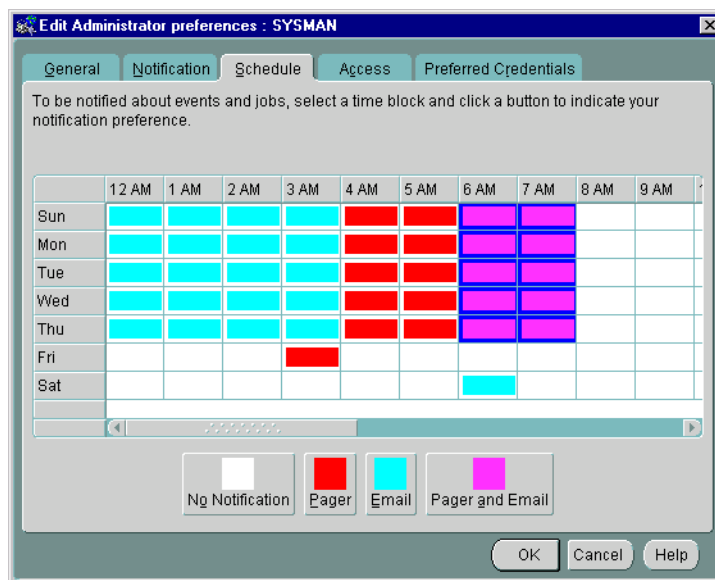
Append job output

When selected, job output is appended to the end of the notification. For instances where job output is large, you may want to specify a limit to the message length. If the job output pushes the notification length past the specified limit, then the job output will be truncated, not the message itself.

Schedule

Use this property sheet page to indicate when you want to be paged and/or e-mailed. Determine the day, hour, and the method of notification. You will receive notifications only on the objects you have permission to access.

Figure 1-13 Administrator Notification Schedule



An administrator can be notified by email, paging, or both email and paging. Paging is recommended for urgent jobs, events, or critical systems.

To determine notifications:

1. Select in the appropriate day/hour to set the notification method for that time period.

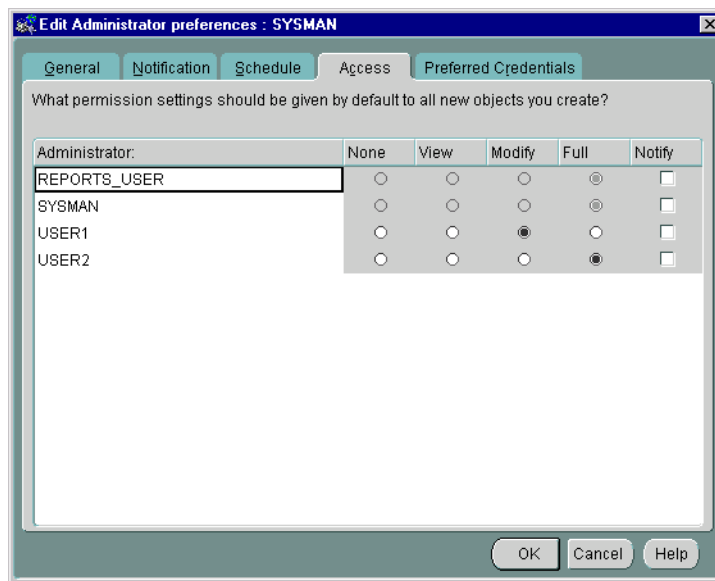
2. Click No Notification, Pager, Email, or Pager and Email at the bottom of the dialog to set the notification type for
3. Repeat steps 1 and 2 for each notification method.

If you want to set the same notification an entire day or an hour of the day, click on the heading for that row or column. For example, if you want to set up paging notification for Saturday, click on the Sat heading after you have selected Pager at the bottom of the screen. You can also drag the cursor to “paint” across multiple cells.

Access

Use this property sheet page to determine the default permissions you want to assign to other administrators for the objects you create. This allows other Enterprise Manager administrators to share objects, such as events and jobs, that you have created. If you work on a team, this page lets you assign access privileges to all the members of the team at one time. When an object is created by an administrator, that administrator is the owner and automatically has full permissions. The owner's permissions cannot be modified.

Figure 1–14 Administrator Default Permissions



Access-control permissions apply to event, group, and job objects. When these objects are created, the default permissions are assigned to other Enterprise Manager administrators according to the selections in this page. These initial permissions can be overridden with the Permissions page of the object's property sheet.

For more information, see the following sections:

- "Group Access Page" on page 4-7
- "Job Access Page" on page 5-20
- "Event Access Page" on page 6-51

Note: Changing your default access level does not retroactively change access levels on existing objects.

The levels of access that you can assign to an Enterprise Manager administrator are shown in Table 1–3, "Administrator Access Levels".

Table 1–3 Administrator Access Levels

Permission Level	Description
None	This access does not allow the administrator to view this object anywhere.
View	This access allows the administrator to view the object, inspect object properties, view job/event status and outputs, and receive notifications if the object is an event.
Modify	This access allows the administrator to edit the object's properties except those reserved for Full permission.
Full	This access allows the administrator to delete the object, modify permissions for other administrators, and change the ownership of the object.
Notify	This access allows the administrator to receive event notifications on the objects. Notify permission cannot be assigned if the administrator's permission level is set to None.

Preferred Credentials

The Administrator Preferred Credentials property sheet page displays a list of targets in the network, along with the target type and the administrator name for accessing the target. The property sheet page is accessed with the Administrator Preferences option of the Configuration menu. Each row in the list of the property sheet includes:

- Target Name
- Target Type (such as database, listener, or node)
- Credentials (Check mark indicating you have already specified connection information. The check mark is either grey (credential set in a previous session) or green (credential set in the current session).

You can click on a column heading to sort on that column. See Figure 1–15, "Administrator Preferred Credentials" for an illustration of the property sheet.

Note: Individual instances of an Oracle 9i Real Application Cluster are listed. Oracle recommends that all instances of a Oracle 9i Real Application Cluster use the same preferred credentials as the Oracle 9i Real Application Cluster.

Also, the Intelligent Agent authenticates the user name and password for all jobs and many events that it runs. These credentials are used when you access a network service in the Navigator or Group, and when running jobs and registering events. If you do not set the preferred credentials correctly, jobs and events may fail. See "Job Credentials" on page 5-4 and "Registering Events" on page 6-4 for more information.

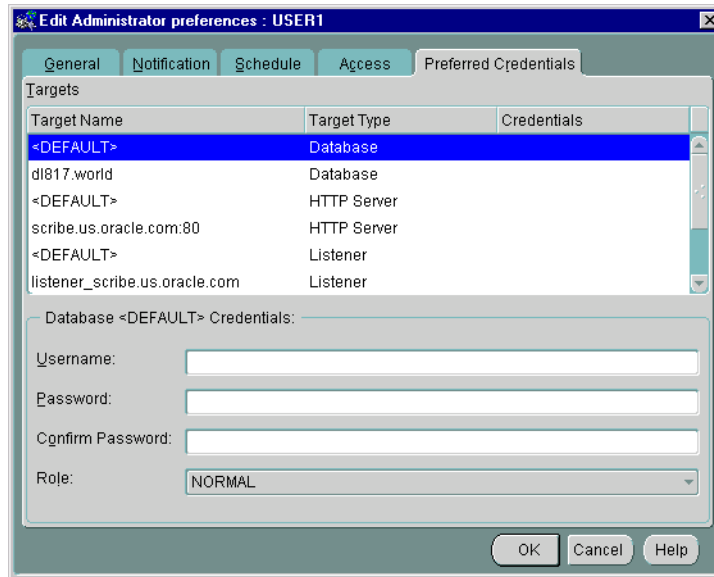
Important: Changes to the Preferred Credentials will not be automatically propagated to previously registered jobs and events. In order to update these jobs and events with the new preferred credentials, you must de-register and subsequently re-register the jobs and events.

Oracle recommends that if some services of a particular type require the same credentials, use the <DEFAULT> credentials selection for that service type to have those credentials used by all services of that type.

For example, if Administrator_1 has a username and password of foo/bar on all the databases that he manages, the administrator needs to enter foo/bar as his credentials in the <DEFAULT> database entry. All database administration tasks will use these <DEFAULT> credentials.

Furthermore, if the administrator has just one database which requires different credentials from foo/bar, he can enter different credentials for that database by selecting that specific database.

For NT users, you must set the preferred credentials for the node (where the NT Intelligent Agent resides) to be the same as the user that is set up to "Logon as a batch job."

Figure 1–15 Administrator Preferred Credentials

Setting Administrator Credentials

Select any row to update the credential fields for the object identified in the row.

Username

Enter the username. This field is required if a password needs to be entered.

Password

Enter the password associated with the username.

Confirm Password

Confirm the password.

Role (Database only)

Select the role from the pull-down list. You need to login with the SYSDBA or SYSOPER role to start up or shut down a database.

Changes to the credentials are recorded when you move to another line. When you are finished, save all your updates by clicking the OK button at the bottom of the property sheet.

Saving Logon Information as a Preferred Credential

Another convenient way to specify Preferred Credentials for discovered services is to select the Save as Preferred Credential option when manually connecting to a service from the Console or integrated application. For example, you attempt to access a newly discovered database from the Console Navigator. Enterprise Manager displays the Connect Information dialog so that you can log in to the database. You type in the requisite information and check the Save as Preferred Credential option. Click OK to log on to the database and save the logon information as a Preferred Credential for this service.

Saving the Contents of a List

Save list allows you to extract information from a multi-column list (such as the Job or Event History panes) and save it in a variety of formats (text, HTML, or Comma-Separated-Values). For example, you can select one or more entries in the Event History pane. By selecting the Save List from the Object menu, you can save the list information (Event, Target, Target Type, Severity, Date/Time, Assigned To, Owner) to an HTML file. The dialog consists of the following elements:

File Name

The name and directory to which you want to save the file information.

Browse

Displays the File Save to dialog. Choose the desired location and file name.

Format

Select HTML (table), Text row, or comma-separated-values (used by spreadsheets or other applications).

View

Active only when HTML is selected as the file format, allows you to view the generated HTML table through the default browser for your system.

Rows

- All (total number) rows – Export all rows in the multi-column list.
- Selected (number of rows selected) rows – Export only selected rows in the multi-column list.

Configuring Enhanced Notifications (Paging/Email)

Enhanced notifications allows Enterprise Manager administrators to be notified either via email or pager. Email configuration allows you to specify the mail service information for your system. Paging configuration allows you to define paging servers, specify pager numbers, and specific paging parameters. For explicit information on setup and configuration, see the *Oracle Enterprise Manager Configuration Guide*.

Management Regions

A management region is a subset of the managed nodes and a subset of management servers. The same Enterprise Manager repository can contain multiple management regions, but a particular node is a member of one and only one management region. You can only assign the Management Server to be in one and only one management region.

The management regions feature allows you to partition the nodes of a particular repository and assign them to a subset of the available Management Servers.

Once this partitioning is performed, the failover and load-balancing operations guarantee that cross-regional traffic is eliminated. The Management Servers that belong in the same partition distribute the management load amongst themselves.

For example, if you have a large, global deployment of Enterprise Manager, or have Enterprise Manager deployed across a mixture of LANs and WANs, using management regions can greatly improve performance over slow networks by allowing you to assign a subset of Management Servers and a subset of discovered nodes to a management region, thus eliminating cross-regional/cross-network communication. Additionally, the ability to segregate a large network makes the use of management regions ideal for mapping firewall boundaries.

If no Management Server is available to service a particular management region, even though Management Servers exist and are serving other regions, you can log on using a Management Server from a different region. The Oracle Enterprise Manager Console will not experience any disruption from an unmonitored management region. You can connect to a Management Server in any Management Region and you will be able to see all nodes in all Management Regions.

Defining a New Management Region

There are two ways to define a Management Region:

- Using the Console
- Using Enterprise Manager Configuration Assistant

By default, Enterprise Manager Configuration Assistant creates the initial “default” Management Region. Under this configuration, all Management Servers that use the existing Repository, as well as all discovered nodes within the Repository, are placed within this DEFAULT Management Region. This single, default management region is sufficient for most situations. However, in order to take advantage of the management region functionality, you must create additional management regions and specify a subset of discovered nodes with a subset of Management Servers within each management region.

To create a new Management Region using the Enterprise Manager Console:

1. Logon as a super administrator to any of the Management Servers working in DEFAULT region.
2. Select Define Management Regions from the Console Configuration menu to access the Management Regions property sheet.
3. Select the Regions tab from the Management Regions property sheet. The Regions page appears.
4. Click the Add Regions button on the Regions page. The Add Management Regions dialog appears.
5. Create a new region by typing the its name into the Management Region Name field and clicking the OK button.

Alternatively, you can use Enterprise Manager Configuration Assistant to create new Management Regions by editing an Oracle Management Server configuration and pointing it to a new repository.

Adding a Management Server to a Region

In this example, we are assuming that there are two already existing Management Servers (OMS1 and OMS2) using the same Repository. In this example, both OMS1 and OMS2 are running.

1. Logon as a super administrator to any of the Management Servers working in DEFAULT region.
2. Select Configure Management Regions from the Console Configuration menu to access the Management Regions property sheet.
3. Select the Regions tab from the Management Regions property sheet. The Regions page appears.

4. Click the Add Regions button on the Regions page. The Add Management Regions dialog appears.
5. Create a new region by typing its name (R1 for this example) into the Management Region Name field and clicking the OK button.
6. Select the Assign Management Servers tab on the Management Regions property sheet. The Assign Management Servers page appears. Assign OMS1 to R1.
7. Select the Assign Nodes tab from the Management Regions property sheet. The Assign Nodes page appears. Assign the required nodes to R1. Important: To assign nodes to R1, the Management Server for R1 must be running. Click OK.

Adding Discovered Nodes to a Management Region

1. Logon as a super administrator.
2. Select Define Management Regions from the Console Configuration menu to access the Management Regions property sheet.
3. Select the Regions tab from the Management Regions property sheet. The Regions page appears.
4. Click the Add Region button on the Regions page. The Add Management Regions dialog appears.
5. Create a new region. For example, R3, by typing the name into the Management Region Name field and clicking the OK button.
6. Install Management Server or Management Servers.

Install the Management Server in a new ORACLE_HOME.

If you are installing from the database CD, choose the install type: Oracle9i Management and Integration>Oracle Management Server.

If you are installing from the separately licensable packs CD, choose the install type: Oracle Enterprise Manager Packs and Management Infrastructure>Custom>Oracle Management Server.

Refer to the installation guide provided with the database release or the installation guide provided with the separately licensable packs.

7. Configure this Management Server (or Management Servers) to be in the new region, R3, using the Enterprise Management Configuration Assistant.

8. Bring up this Management Server (or Management Servers). At this point there will be two regions available: DEFAULT and R3.
9. Log on as a super administrator using one of the new Management Servers in R3.

Discover all the nodes you need to add. They will automatically be placed in the R3 region.

Removing a Management Region

A Management Region must be empty before attempting to remove it. First, ensure that there are no nodes in the region. Second, ensure there are no Management Servers in the region.

1. Logon as a super administrator.
2. Select Define Management Regions from the Console Configuration menu to access the Management Regions property sheet.
3. Select the Regions tab from the Management Regions property sheet. The Regions page appears.
4. Select the region you wish to remove.
5. Click the Delete Region button. You are asked if you wish to continue.
6. Click Yes.

Managing HTTP Servers

HTTP server management allows you to discover, start up, shut down, and monitor the server's operational status via the Event system. Additional functionality (extra job tasks and event tests) are available when the e-Business Management Tools (part of the Diagnostics Pack) are installed.

Discovering HTTP Servers

You discover a web server as you would any other managed target from the Enterprise Manager Console. See "Discovering Targets" on page 3-6. The discovery process does not require the web server to be running. Once discovered, you set the server access authorization through the Enterprise Manager Preferred Credentials (Node). See "Preferred Credentials" on page 1-25.

Important: When running on UNIX, if the web server is listening on a port number less than 1024, the administrator must have root privileges in order to submit startup or shutdown jobs to the server.

Starting Up or Shutting Down a HTTP Server

Once an HTTP server appears as a managed object within the Console Navigator, you can select Startup from the context-sensitive menu. HTTP server management utilizes the Enterprise Manager Job system to perform the actual startup operation. For this reason, you can also start the HTTP server by selecting Create Job from the Console Job menu, selecting Web Server as the Target Type, and selecting the Startup job task. After the startup job task is submitted, a dialog displays indicating the task has been submitted to the Job system. Status messages pertaining to this task appear in the Console's Job Pane.

To stop an HTTP server, select the desired HTTP server from the Console Navigator and select Shutdown from the context-sensitive menu. As with web server startup, you can also shut down the server from the Edit Web Server General page, or use the Enterprise Manager Job system directly.

Alternatively, you can select Edit to display the Edit Web Server General Page. The current web server state is selected (Shut Down or Started). To start the web server, select Started and click OK.

Determining the Status of HTTP Servers

Select a discovered HTTP server from the Console Navigator. HTTP server status and configuration information displays in the detail view. Information displayed pertaining to the selected web server includes:

- **Server Name:** Server's Internet host name. Example: xyz-machine.your_company.com. If this machine does not have a registered DNS name, an IP address will appear.
- **Port:** The network port to which the web server listens.
- **Version:** Apache server version number.
- **Server Root Directory:** The top-level directory where all web server-related files (configuration, error, and log files) are stored.
- **Configuration File Location:** Location of the web server configuration file.

- **Status URL:** The location of the Apache-generated HTML page that provides the current server statistics in an easily readable form.

Paging/Email Blackout

Paging/Email Blackout allows an administrator with super administrator privileges to suspend paging and email notifications for specified targets and/or services that have been previously discovered in the Navigator. Paging/email blackouts deactivate enhanced notification (email/paging), thus preventing Enterprise Manager administrators from being flooded with emails and pages if a managed target/service is brought down. For example, if a target is brought down on a regular basis for scheduled maintenance, a super administrator can schedule a paging/email blackout for that target to prevent enhanced notification during the maintenance period.

In addition to paging/email blackouts, the Enterprise Manager administrative framework also allows you to specify target-level blackouts to suspend all management and data collection activity for specific targets in your enterprise. See "Target-level Blackouts" on page 1-37.

Specifying Total Paging/Email Blackout

Total paging/email blackout specifies that a blackout be started immediately with an indefinite duration. The super administrator must manually turn off Total Paging/Email before any new blackout schedules can take effect or existing blackout schedules resume. To specify a total paging/email blackout:

1. Choose Set Paging/Email Blackout from the Console's Configuration menu.
2. Check Total Paging/Email Blackout
3. Click OK.

Note: You can also access Total Paging/Email Blackout menu item from a target's context-sensitive menu in the Enterprise Manager Navigator.

Defining a Paging/Email Blackout

To define a paging/email blackout for a specific node:

1. Select a target from the Navigator and then choose Set Paging/Email Blackout from the Console's Configuration menu. Optionally, you can select an object in the Navigator and access Set Paging/Email Blackout through the context-sensitive menu.
2. Click Create. The requisite naming and scheduling UI are displayed. Important: Be sure that Total Paging/Email Blackout is not checked for either the specific service or its parent target.
3. Specify a Blackout Name. All Blackout names for this node/service must be unique. Note: the Target name reflects the node/service selected in the Navigator and is not editable from this dialog.
4. Select an Occurrence

Select the frequency that you want the paging/email blackout to occur. The choices are Once, On Interval, On Day(s) of Week, and On Date(s) of Month.

- Once—schedules the blackout only one time beginning and ending on the dates and times you choose.
 - On Interval—allows you to schedule a specific time interval between paging/email blackout periods. The interval can be a combination of hours and minutes, or number of days. Select the value you want to change and click on the scroll buttons. You can also type in a new value. Select an Effective Period in which your blackout schedule will be valid.
 - On Day(s) of Week—allows you to schedule the blackout on one or multiple days (Sunday, Monday, etc.) of the week. Click on the days of the week to select the days you want the blackout period scheduled and set a start time or duration. Select an Effective Period in which your blackout schedule will be valid.
 - On Day(s) of Month—allows you to schedule the blackout on one or multiple days (1 - 31) of the month. Click on the dates of the month to select the dates you want the blackout period scheduled and set a start time or duration. Select an Effective Period in which your blackout schedule will be valid.
5. Click OK.

Important: When selecting *On Day(s) of Month* or *On Day(s) of Week*, be aware that each day terminates at 12:00 A.M. This is important when scheduling paging/email blackout periods that span the day change (12:00 A.M.).

For example, if you wish to schedule a blackout that begins on Tuesday at 10:00 P.M and lasts for six hours, you must define two blackout periods:

- Tuesday from 10:00 P.M. to 11:59 P.M.
 - Wednesday from 12:00 A.M. to 4:00 A.M.
-
-

Making a Copy of an Existing Blackout Schedule

To schedule a paging/email blackout schedule for a node that is identical to that of another managed node:

1. Choose Paging/Email Blackout from the Console's Configuration menu. Optionally, you can select an object in the Navigator and access Paging/Email Blackout through the context-sensitive menu.
2. Select a Blackout Schedule entry from the Blackout Name list.
3. Click Create Like. A new entry is created in the list prefixed by "Copy of."
4. Modify the blackout period parameters (such as Blackout Name) as desired and click OK.

Turning Blackout Schedules On and Off

Once you have defined one or more paging/email blackout periods for a node/service, they persist in the Blackout Name list associated with the node/service until deleted. You can turn paging/email blackout on or off by checking the box located to the left of the individual blackout name. You can edit the Occurrence parameters to reflect the present date and time. Paging/email blackouts are set ON by default.

Deleting Paging/Email Blackout Periods

1. Choose Paging/Email Blackout from the Console's Configuration menu. Optionally, you can select an object in the Navigator and access Paging/Email Blackout through the context-sensitive menu.
2. Select the desired Blackout Name.
3. Click Delete.

Viewing Blackout Periods

Once you have defined blackout periods on multiple nodes, you can view them by selecting the desired target in the Console Navigator and choosing Paging/Email Blackout from the Navigator menu. If you select a high-level folder from the Navigator, such as Databases, Groups, or Nodes, rather than an individual target, all blackout periods defined in the container type are displayed.

Target-level Blackouts

Target-level Blackouts allow Enterprise Manager users to suspend any or all management and/or data collection activity on one or more managed targets. This capability permits maintenance or emergency operations to be performed. Target-level Blackouts must be set on the target running a 9i version of the Intelligent Agent and cannot be set from the Console. See the Intelligent Agent User's Guide for more information.

Specifically, blackouts can suspend:

- **Events:** All events registered on a target will not be evaluated or triggered for the duration of the blackout.
- **Jobs:** All jobs submitted to a target will not be scheduled or run for the duration of the blackout.
- **Data Collections:** All current historical data collection activities for a target are stopped. However, loading of data collected for a target **prior** to the blackout will continue as long as the database is up. New collections can be submitted but will not proceed unless the blackout ends.

The Standalone Console

Beginning with Release 9.0 when you launch the Enterprise Manager Console or various other Enterprise Manager applications, you are prompted to choose between launching the product standalone (i.e. not connecting to the middle tier Management Server) or logging into a Management Server.

Launching the Console standalone allows a single administrator to perform simple database schema, instance, storage, security, and other database tasks by connecting directly to the target database(s).

Launching standalone does not require a middle tier Management Server or Intelligent Agents on target machines.

This chapter discusses the topics listed below:

- Choosing to Launch the Console Standalone
- Starting the Standalone Console
- Adding Databases to the Tree in the Standalone Console
- Connecting to a Database in the Standalone Console
- Connecting to the Database As a Different User
- Viewing If You Are Connected As SYSDBA
- Removing a Database from the Tree
- Changing from Using the Console to the Standalone Console
- Editing Local Preferred Credentials in the Standalone Console
- Changing from Using the Standalone Console to the Console

Choosing to Launch the Console Standalone

When you launch the Enterprise Manager Console, you are prompted to choose between launching the product standalone or logging into a Management Server.

Choose to launch the Console standalone when you want to connect directly to your managed target(s) to perform administration tasks. With Enterprise Manager Release 9.0 the standalone Console only supports connecting directly to database targets, no other target types are currently supported.

Launching standalone does not require a Management Server or Intelligent Agents on managed targets. Consequently, when you launch the Console standalone, you do not have access to functionality typically available through the Management Server and Intelligent Agent, such as:

- Management of several different target types (e.g. database, web server, application server, applications, etc.)
- Sharing of administrative data among multiple administrators.
- Proactive notification of potential problems.
- Automation of repetitive administrative tasks.
- Backup and data management tools
- Customization, scheduling, and publishing of reports
- Running the client from within a web browser.

Starting the Standalone Console

On Windows-based platforms, you start the Console from the Windows Start Menu.

You can also start the standalone Console from the command line using the command:

```
C:\> oemapp console
```

On UNIX platforms, you start the Console from the command line using the command:

```
$ oemapp console
```

Figure 2–1 Enterprise Manager Login



When the dialog appears, choose "Launch standalone" and press OK.

Note: The login choice is remembered for the next time you log in whether the last login was Launch standalone or Login to the Oracle Management Server. If you had selected Login to the Oracle Management Server, the Management Server is remembered.

To bypass the Console login, you can enter the following command at any supported operating system command line:

```
oemapp console oem.loginmode=standalone
```

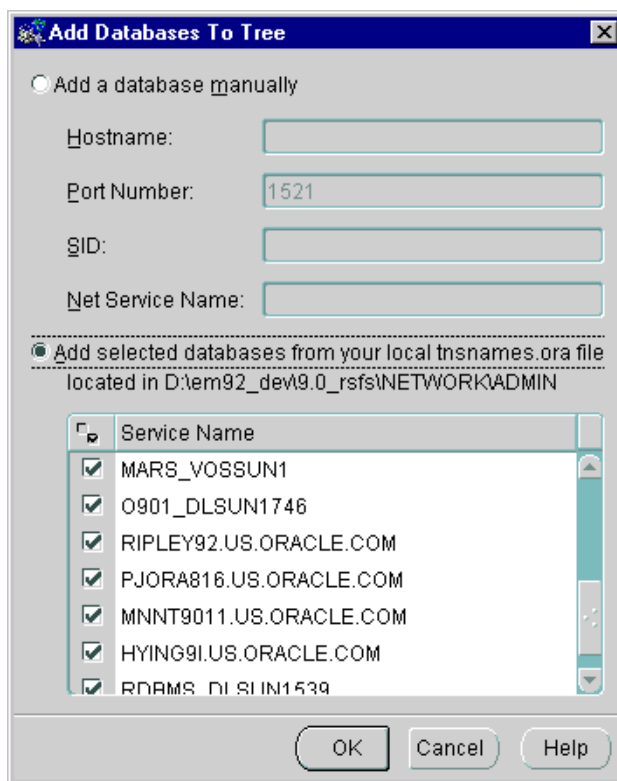
By entering the command, you will immediately see the standalone Console.

If you are starting the standalone Console for the first time, the left panel of standalone Console is empty because you have not yet added the databases you want to manage. The Add Databases To Tree dialog appears automatically so that you can add them to the navigator tree.

Adding Databases to the Tree in the Standalone Console

The Add Databases To Tree dialog appears automatically when you start the standalone Console for the first time; it is also available from the Navigator menu.

Figure 2–2 Add Databases to Tree



The Add Databases To Tree dialog allows you to manually enter the Net service names or add them from the local tnsnames.ora file.

Add a database manually

You can add databases to the standalone Console navigator tree by manually filling in the following fields:

- **SID:** the database system identifier, usually the instance name, such as ORCL

- **Hostname:** the machine or node name where the database is located
- **Port Number:** the database listener port address, usually 1521 or 1526
- **Net Service Name:** A name which uniquely identifies a database when connecting to a machine. It is usually the global database name.

For example: ORCL.world.

Note: Adding a database manually automatically updates the local tnsnames.ora file located in your <Oracle_Enterprise_Manager_Home>/network/admin directory.

Add selected databases from your local tnsnames.ora file

You can populate the standalone Console navigator tree by reading the database service names from the local tnsnames.ora file located in your Oracle Enterprise Manager home. The Add Databases To Tree dialog displays a list of databases identified in your tnsnames.ora file from which you can select or deselect. Click the column header to the left of Service Name to either select or deselect all the databases. If you have deselected all the databases, you can choose specific databases by selecting their checkboxes.

Note: Currently only TCP/IP service names can be added manually for the standalone Console. If other network protocols are required, add them by entering them in the tnsnames.ora file using the Oracle Net Configuration Assistant. All protocols are supported when you import selected services from your tnsnames.ora file.

Connecting to a Database in the Standalone Console

There are three ways to connect to a database for the standalone Console:

- Click the plus symbol next to the database icon in the Console navigator tree. The preferred credentials are used if the connection information has not been set previously.
- Double-click the database icon in the standalone Console navigator tree. The preferred credentials are used if the connection information has not been set previously.
- Select the database and then select the Connect item from the Navigator menu.

If no preferred credentials are set in the Oracle Enterprise Manager Console, the Database Connection Information dialog box appears. If preferred credentials are already set, you will connect to the database using this login information.

In the Database Connect Information dialog, enter the following information to connect to the database.

Username

The Oracle username for the database to which you are connecting.

For example: system

Password

The password for the username to which you are connecting.

For example: manager

Connect As

You can select from a pull-down list whether you want to connect to the database with NORMAL, SYSOPER, or SYSDBA privileges. Select NORMAL to connect to the database as an ordinary user. Select SYSOPER to connect to the database with special operator privileges, such as capabilities to shut down and start up the database. Select SYSDBA to connect to the database as a user with full database privileges such as the capability to grant any privileges to any user.

To use SYSOPER and SYSDBA privileges, a password file or OS group authentication must be created and set up for your database.

Note: In 9i, the SYS account may only be used with SYSDBA or SYSOPER (not NORMAL). You also do not set up a password file as with previous database versions.

Save As Local Preferred Credentials

Saving preferred credentials is an option which enables you to store login information in a local file such as username, password, and role (NORMAL, SYSOPER, or SYSDBA). Passwords are always stored in encrypted format.

This login information is used when a connection is established for the database instead of having to type a username and password each time.

You can set preferred credentials by selecting the Save As Local Preferred Credentials checkbox or you can save credentials later by using the Edit Local Preferred Credentials dialog available from the Configuration menu.

Connecting to the Database As a Different User

In the standalone Console, you do not have to disconnect as a user before reconnecting as a different user.

If you are already connected to a database and you want to reconnect as a different user

- Right-click the database icon and select Connect from the context-sensitive menu.
- Select the Connect item from the Navigator menu.

When the Database Connect Information dialog appears, enter a different username and press OK.

You will automatically be disconnected and reconnected.

Viewing If You Are Connected As SYSDBA

When a connection is made to a database, a connection icon is displayed on top of the database icon in the standalone Console tree.

Next to the database name, you will see the username and role, if applicable, that you are connect as. For example, if you have connected as user SYS with the SYSDBA role, you should see "sys as SYSDBA."

Removing a Database from the Tree

To remove a database from the standalone Console navigator tree, follow the steps below:

1. Highlight the database you want to remove.

2. From the Navigator menu, choose Remove Database from Tree item.
3. A message appears, saying, "Remove the <name> database from tree. Are you sure?" Click the Yes button.

Note: Removing the database removes the entry from the standalone Console Navigator tree. It does not remove the entry from the tnsnames.ora file or physically remove the database. If desired, you have the option of adding the database to the navigator tree again.

Changing from Using the Console to the Standalone Console

When the Console is connected to an Oracle Management Server, the Console navigator tree is populated with the discovered databases.

The following steps describe how to copy these databases so they can be used in standalone mode.

The databases displayed in the navigator tree can be added to the tnsnames.ora file and their associated preferred credentials saved to a local file:

1. Use the Add Services to tnsnames.ora from the Configuration menu to copy the services to the local tnsnames.ora file.

Note: The passwords are encrypted in the local file to prevent the file from being used on other machines.

To make these databases available in standalone mode, you can perform the following steps:

2. Start the standalone Console.
3. From the Navigator menu, select the Add Databases To Tree item.
4. In the Add Databases to Tree dialog, select the Net service names to add from the local tnsnames.ora file.
5. To save the local preferred credentials of a service, select the Edit Local Preferred Credentials item from the Configuration menu.

Editing Local Preferred Credentials in the Standalone Console

Note: The passwords are encrypted in the local file, dbastudio-<os_username>.crd, so that they cannot be copied to another machine and used by a different user.

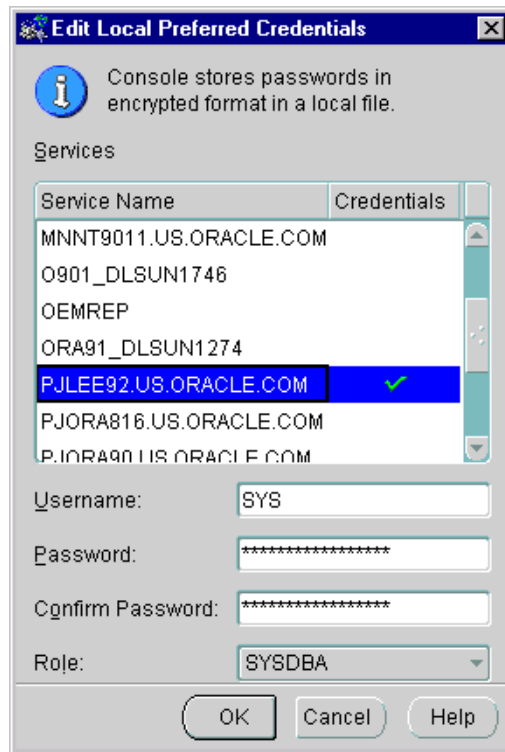
The standalone Console stores a list of databases that are displayed in the standalone Console navigator tree. If preferred credentials are specified for any of these databases, the username, encrypted password, and role are added to the local file. The local record, dbastudio-<os_username>.crd, is located in the <Oracle_Enterprise_Manager_Home>/sysman/config/pref directory.

The login information is used when a connection is established to the database instead of having to type a username and password each time.

To save or edit the local preferred credentials of a service, you can perform the following steps:

1. From the Configuration menu, select the Edit Local Preferred Credentials item. The Edit Local Preferred Credentials dialog displays a list of services.

Figure 2–3 Edit Local Preferred Credentials



2. In the Edit Local Preferred Credentials dialog, select the service name of the database and enter or update the preferences for connecting to that database.
 - Username: Enter the username. This field is required if a password has been entered.
 - Password: Enter the password. You can leave this blank if you want to be prompted for a password when you connect to the database.
 - Confirm Password: Confirm the password. The contents of this field must be the same as the contents of the Password field.
 - Role: Select the role from the pull-down list. You must login with the SYSDBA or SYSOPER role to start up or shut down a database.
3. Click the OK button at the bottom of the Edit Local Preferred Credentials dialog to save your updates.

Note: You can save or overwrite the preferred credentials by selecting the "Save As Local Preferred Credentials" checkbox in the Database Connection dialog.

Changing from Using the Standalone Console to the Console

After you have started the standalone Console, and you want to connect to a management server (switching therefore from standalone into distributed mode), you must close and restart the Console.

If you have been using standalone Console, the list of databases displayed in the navigator tree is retrieved from a local file.

At some stage in the future, you may decide to administer jobs, events, and groups; run the database applications through a web browser; or perform backup and data management tasks. These tasks require you to run the Console connected to an Oracle Management Server.

If you now decide to run the Console connected to an Oracle Management Server, the databases displayed in the navigator tree may be different than those in standalone mode, because the list of databases is retrieved from the repository and not from a local file.

To add the databases that were available from the standalone Console, you will need to discover the services from the Oracle Enterprise Manager Console if the databases are not already available from the Console connected to the Oracle Management Server.

Navigator

The Navigator graphically displays network objects and allows you to administer the objects. The Navigator tree displays a direct view of the network's nodes and services, the objects they contain, and the relationships among objects. The topics discussed in this chapter include:

- Navigator Pane
- Navigator Menu
- Discovering Targets
- Manipulating Objects in the Navigator
- Removing a Node from the Navigator
- Node Removal Failure

Navigator Pane

The Navigator pane provides:

- Identification of the objects or services on nodes in the managed environment.
- Views of the objects in a network environment and the relationships among the objects. By expanding an object, you can display any objects it contains.
- Methods of accessing and launching administration tools on the objects.
- A source of objects with which to populate groups.
- A source of objects upon which to launch DBA tools and other integrated applications.
- A source of objects for copying with the drag and drop method.

Because the Enterprise Manager Console uses a master/detail-type user interface, objects selected in the "master" Navigator tree control what is displayed in the "detail" pane to the right. This simple, yet effective interaction paradigm is consistent for all Enterprise Manager applications.

Populating the Navigator Tree

The Navigator tree is populated with objects by selecting Discover Nodes from the Navigator menu. The types of services in the network may include, but are not limited to:

- Databases
- Groups
- Listeners
- Nodes
- HTTP Servers

The actual number and types of objects that appear in the Navigator differ according to the types of targets being monitored and any options you have installed in your enterprise environment.

Expanding Objects in the Navigator

Each object type in the Navigator tree is identified by an icon and name. If there is a '+' or '-' to the left of an object's icon and name, the object is a container that can be expanded to display other objects. A container that is represented by a folder icon is

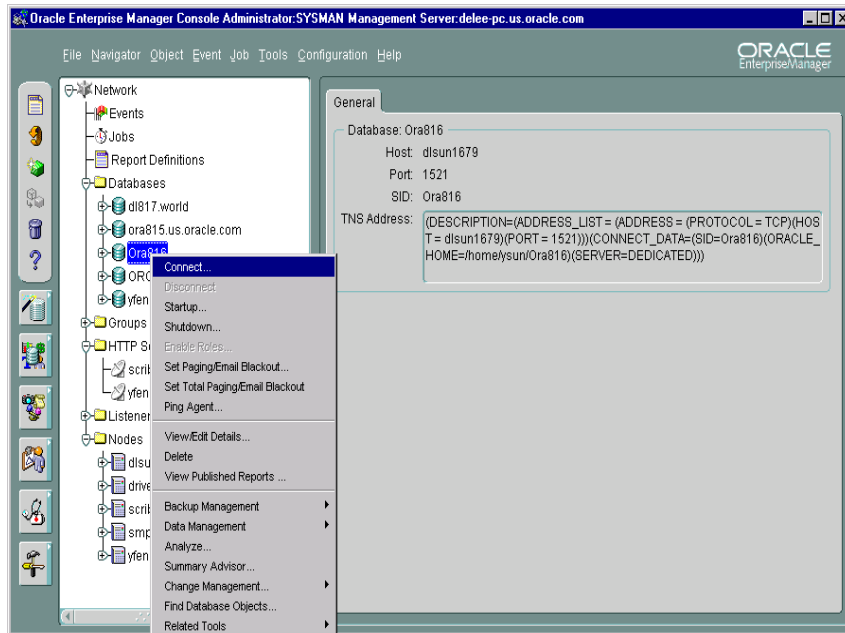
a logical grouping, or collection, of one specific type of object, such as databases. Other containers are objects that hold multiple types of objects. See Figure 3-1, "Navigator Pane and Context-sensitive Menu" for an illustration of a Navigator pane.

The top-most object in the Navigator tree is the network container. The network folder can contain:

- Events
- Jobs
- Report Definitions
- Databases folder
- Groups folder
- HTTP Servers
- Listeners folder
- Nodes folder

In addition to the folders listed above, there may also be folders for additional targets that have been discovered in your enterprise.

You can expand tree containers to view the objects and relationships in the environment. For example, you can expand a node to view the databases and listeners on the node. If you expand a database, you can access common administration areas such as schema, instance, and security. When a connection is made to a database, an icon displays on the database in the tree. See "Database Administration" on page 10-1 for more information on integrated database administration tools that are accessible from specific Navigator objects.

Figure 3–1 Navigator Pane and Context-sensitive Menu

Launching Tools

To launch a database tool in the context of a database or database object, select the object in the Navigator tree that you want to access. You can then execute a tool from the Tools menu or with the Related Tools option of the right-mouse menu.

You are connected to the database according to the preferred credentials that have been set up for the database. If connection to the database fails for any reason, the Database Connect Information dialog box displays. At that point, you can enter the requisite information and optionally save that information as a preferred credential. See "Preferred Credentials" on page 1-25 for information.

Navigator Menu

The Navigator menu allows you to manage objects in the Navigator pane. The menu options are enabled according to the object selected in the Navigator tree. Usually the Create, Create Like, Edit, and Remove menu options are available when an object is selected. See Figure 3–1, "Navigator Pane and Context-sensitive Menu" for an illustration of the Navigator menu.

Clicking on an object in the Navigator with the right mouse button, displays a *context-sensitive menu* of all the options you can use to manipulate the object, as well as related tools.

For information on operations for an object type, see the chapter in this guide that discusses the specific application that manipulates the object type. See Chapter 10, "Database Administration" for an overview of the database application tools.

Refresh

Refreshes the current Navigator view.

Find

Locates discovered databases within the Navigator Tree. It is available when you select a Database or Node icon on the tree.

Discover Nodes

Discovers services on a node in the network. You can also display the status of network services. See "Discovering Targets" on page 3-6 for more information.

Refresh All Nodes

Refreshes services for all discovered nodes.

Ping Agent

Attempts to contact the Intelligent Agent on a specified node. If it fails, either the node is down or the Intelligent Agent isn't running. If either of these is true, no Job or Event notification will come through.

Connect

Displays the Connect Information dialog allowing you to specify logon information for a selected target (such as a database) in the Navigator.

Disconnect

Log out of the selected target.

View By Schema

When viewing a database schema in the Navigator, all objects are grouped by schema.

View By Object

When viewing a database schema in the Navigator, all objects are grouped by object type.

Enable Roles

Displays the Enable Roles dialog allowing you to enable roles for an administrator.


Application SQL History

Displays SQL statements produced by Enterprise Manager's database administration functions. This option allows you to view the last 100 SQL statements (maximum) executed by the application you are using (i.e., Instance Management, Schema Management, and others) against the selected database.

Discovering Targets

The Navigator provides service discovery functions for identifying network targets populating the Navigator tree through the Discovery Wizard. The Discovery Wizard is activated each time you select Discover Nodes from the Console's Navigator menu.

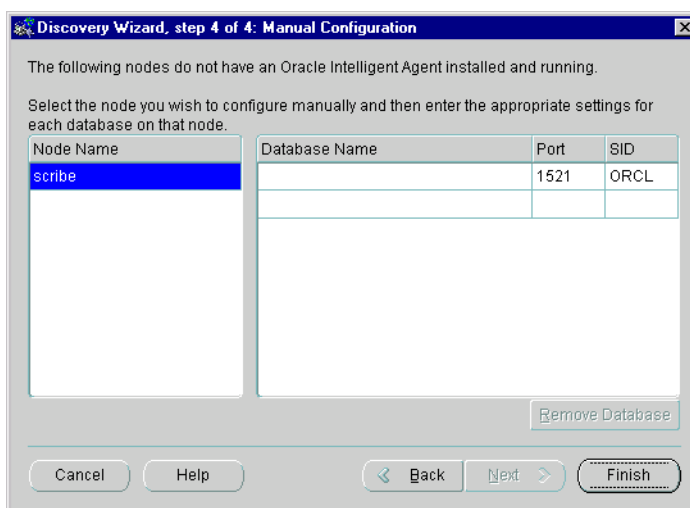
Figure 3–2 Discovery Wizard



The Discovery Wizard searches the network for the targets you specify. If that target has an Oracle Intelligent Agent running, it will be added to the Enterprise Manager Console Navigator for management along with any services running on that target. You will be able use these services as targets in jobs, events, and groups.

If there is no Oracle Intelligent Agent running, the Discovery Wizard gives you the option of performing a manual configuration. This allows you to define an Oracle database on that target so that it appears in the Console Navigator. While manual configuration permits database access from the Navigator, its use is not recommended since Enterprise Manager services, such as jobs and events, will not be available. For more information on the Intelligent Agent and the discovery process, see the *Oracle Intelligent Agent User's Guide*.

Figure 3–3 *Discovery Wizard: Manual Configuration*



When refreshing an existing target, the Discovery Wizard verifies and updates the list of services on a target that had been previously discovered. If the target was previously Manually Configured, the Discovery Wizard will check if that target now has an Oracle Intelligent Agent and if so, provide that information with the option to discover using that Intelligent Agent, or to update the previous configuration.

Note: Services on the machine where the Management Server is running are automatically discovered when you start the Console and connect to the Management Server.

Adding a Service

To discover new services, choose **Discover Nodes** from the Console Navigator menu to call up the Discovery Wizard. Any services running on that target will, if successfully discovered, appear in the Navigator. The Discovery Wizard guides you through the service discovery process. For more information on service discovery and refresh, refer to Oracle Enterprise Manager online help.

Refreshing a Target

Refreshing a target verifies and updates the list of services on a target that has been previously discovered. You can refresh the discovery of a target by selecting a target from the Navigator, clicking the right mouse button, and choosing **Refresh Node** from the context-sensitive menu.

To refresh an existing service, choose **Refresh Nodes** from the Console Navigator menu to call up the Refresh Wizard. The Refresh Wizard performs a "rediscovery" of any target selected in the Console Navigator. This menu option is greyed out until a specific target is selected from the Navigator Nodes folder.

Removing a Target

Select a target from the Console Navigator. Using the right-mouse button, choose **Delete** from the context-sensitive menu.

Pinging the Intelligent Agent

Pinging the Intelligent Agent on a monitored node allows you to check whether the Management Server can communicate with the Intelligent Agent on a monitored node. Proper Intelligent Agent operation is required for successful service discovery. If discovery is successful, the discovered node and services appear in the Navigator tree. If discovery has failed, an error message displays.

The following are common problems if the service discovery has failed:

- The node must have an Oracle Intelligent Agent started and running.
- TCP/IP network protocol must be used.
- If a database has not been discovered, make sure an entry for the database is in the `tnsnames.ora` file on the node where the Intelligent Agent is running. On a UNIX platform, check the `oratab` file.
- If a database has not been discovered, you may need to stop and restart the Intelligent Agent if the database was not configured or installed when the Intelligent Agent was last started. When the Intelligent Agent starts, database

entries are written to the `services.ora` file in the `ORACLE_HOME/network/agent` directory (Unix example).

Determining Node Properties

Since Enterprise Manager allows you to manage a heterogeneous environment, there will be situations where you may need to determine operational properties of that node (network name, operating system and version, and Intelligent Agent version). For example, you may want to find out which version of the Intelligent Agent is running on each node within your enterprise to determine which nodes require Intelligent Agent updates. To determine node properties, select the desired node from the Nodes folder in the Navigator. Pertinent node information is displayed in the detail pane of the Console.

Manipulating Objects in the Navigator

The Navigator interface provides easy manipulation of the services and objects in your managed network from the Enterprise Manager Console. From the Navigator, you can apply some or all functions available from the Console and any integrated applications to selected objects.

Administering Objects

To administer an object, select the object in the Navigator tree and choose the administration task from the Navigator menu, or use the menu options available when you click the right mouse button on the object. The menu options available vary according to the object selected. See "Navigator Menu" on page 3-4 for more information. When you create or edit an object, the property sheet for that object displays. For information on the property sheets, see the chapter on the DBA tool that administers the object.

Copying Navigator Objects

You can drag and drop some objects in the navigator to make copies of the object in different locations. For example, you can drag and drop a database user or role from one database to another to add that user or role to another user. However, if you drag and drop a group that resides within another group to a different group, it will move the group instead of copying it.

Removing a Node from the Navigator

Before removing a discovered node from the Enterprise Manager Navigator, you must remove all jobs and events submitted against that node. If the Intelligent Agent is down when you attempt to remove the node, and jobs and events have not been cleared previously, Enterprise Manager will allow you to remove the node. However, when the Intelligent Agent is restarted and the node re-discovered, the Intelligent Agent will be out of sync with the Oracle Management Server. The following example illustrates how this situation can occur.

1. You register an event with the Intelligent Agent on a managed node.
2. The Intelligent Agent goes down.
3. You remove the node from the Navigator without removing the event.
4. The Intelligent Agent is started and the node is rediscovered.
5. The Console Event window is clear.
6. You attempt to register the same event with the managed node.
7. Enterprise Manager displays a message stating the event is already registered even though it is not.

The Intelligent Agent remembers the current status of jobs and events through a series of generated files. Oracle Intelligent Agents 8.0.6 and higher and Intelligent Agents 8.1.6 and higher automatically synchronize with the Oracle Management Server by removing old Intelligent Agent overhead files. For older versions of the Intelligent Agent, you must remove these files manually.

To manually remove the Intelligent Agent overhead files:

1. Stop the Intelligent Agent (if it is currently running).
2. Go to the Intelligent Agent directory (`$ORACLE_HOME/network/agent`).
3. Delete files with the following extensions: `.q`, `.jou`, `.inp`
4. Restart the Intelligent Agent.

Node Removal Failure

Typically, node removal from the Navigator fails because of problems encountered by the Management Server. The following are the most probable causes.

- The Management Server encountered an error condition while accessing the repository. Check the Management Server log for detailed error messages.

(ORACLE_HOME/sysman/log/oms.log). If problem resolution is not obvious from the error messages, contact Oracle Support.

- The Management Server could not open a database session to the repository. Ensure that the repository is up and running. You may have to shut down other Enterprise Manager applications that may be connected to the repository. Check the Management Server log for more detailed error messages.
- Different DNS servers may resolve node names in different ways. For example, if a Management Server is running on a machine whose DNS server resolves a particular node as "A", the Management Server discovers the node as "A". Next, the Management Server is run on another machine whose DNS server resolves the same node as "B". If you switch back to running the Management Server on the first machine (originally resolved the node as "A"), attempting to remove the node from the Console Navigator will now fail since the repository entries will be out of sync.

The solution is to run the Management Server on the machine from which you last discovered the node. In the example above, you would need to run the Management Server on the second machine (resolved the node as "B") and then remove the node from the Navigator. Always ensure that the DNS setup is consistent for all nodes where the Management Server is run.

The Group system lets you organize the objects you manage into logical categories for more efficient management and administration. You can organize groups based on any criteria you determine, such as function, department, geographical location, or number of administrators. You can then place the groups on a map or graphical view of the network making it easy to locate, diagnose, and act on encountered conditions. The system also lets you assign jobs to groups and monitor their status. The Groups system is especially useful for managing environments with many databases and services.

The topics discussed in this chapter include:


- Group View Pane
- Managing Groups
- Viewing Reports for Targets within a Group

Group View Pane

The Group view pane provides a graphical custom view of your managed environment and is displayed when any group is selected in the Console Navigator. As with any object in the Navigator, the appropriate group view is displayed in the detail area of the Console. See Figure 4–1, "Navigator and Group Panes" for an illustration of a Group view.

Note that groups are also displayed in the Navigator. To display the members of a group in this pane, expand the group by clicking the "+" symbol next to the group name. By default, the Group folder is empty.

Figure 4–1 Navigator and Group Panes



Managing Groups

A group is a collection of targets, such as databases, listeners, nodes or HTTP servers, that share a common location or function. You can create, modify, and remove groups to further organize your network view. Different types of objects, like databases and listeners, can be grouped together.

Groups appear in two places: the Navigator and the Group View pane, once a group is created.

A group is represented in the Navigator by an icon and a name. You can double-click on the group's icon to expand the group, and view and update the group's members and sub-groups.

The Group View pane gives you a graphical view of the objects in the group and their individual status. You can also add background maps to the Group pane to better visualize locations of particular objects.

Grouping nodes or services can simplify tasks that are applied to all members of the group. For example, in order to execute a SQL script on all the databases in the CHI_MAIL group, you can use the Job Scheduling services to schedule a job on the group. The job that executes the SQL script is scheduled on all the databases in the group. The job will not be applied to any other objects in the group.

Note: Jobs and events applied to a target before it joins a group will not be applied to the group retroactively. You must re-apply any jobs or events you want to apply to the entire group.

Important: If you submit a job or event against a group that contains a mixture of manually configured and discovered targets, the job or event will only be submitted against the discovered targets. If all the targets in a group are manually configured, then the job/event submitted against that group will fail.

Automatic Group Refresh

The Group View page can be refreshed automatically or manually.

Manual Refresh

The Group View page can be refreshed manually by clicking the Refresh button located at the bottom of the Group View page. When you refresh manually, only the currently selected group is refreshed.

Automatic Refresh

When refreshed automatically, all group views currently open for viewing by the administrator are refreshed.

Automatic refresh, along with the refresh frequency, can be set from the Group View page or from the General page of the Edit Administrator Preferences property sheet. Administrators can set a refresh interval greater than or equal to the minimum refresh interval set for all administrators by the Super Administrator. The minimum refresh interval can be changed from the General page of the Edit Administrator Preferences property sheet. Changing the minimum refresh interval can only be performed by Super Administrators. By default, the minimum refresh interval is five minutes.

In general, it is best to refresh group views manually. However, for situations where a Group View page needs to be displayed for extended monitoring periods, automatic refresh can be used.

Important: Setting an automatic refresh frequency too high (low interval value) may impact network and system performance.

Creating a Group

To create a group:

1. Right-click on the Groups folder to display the context-sensitive menu and choose Create. The Create Group property sheet appears.
2. On the General page of the Create Group property sheet, fill in the requisite identification information.
3. From the Available Targets list, select the specific target(s) you want to appear in the group. By holding down the Shift (in sequence) or Alt (random) key, you can select multiple objects.
4. Click Add to move the items to the Selected Targets list.
5. Click on the Access tab to set administrator access privileges.

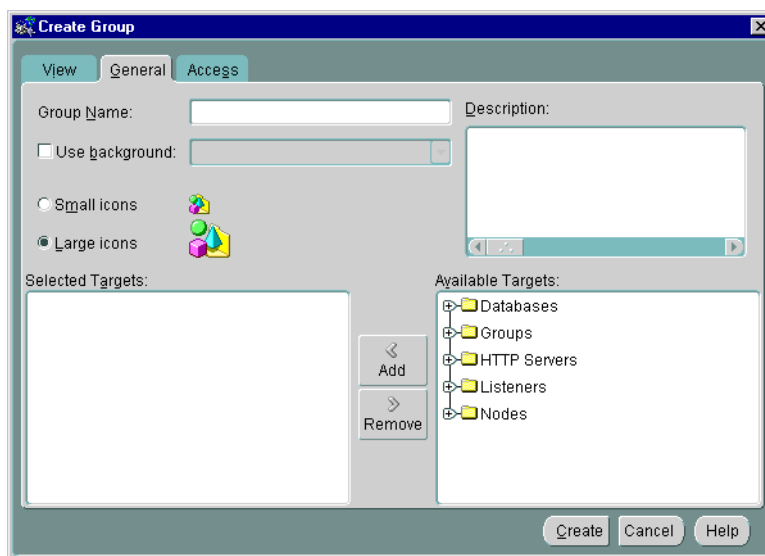
6. Click on the View tab to see a graphical representation of the group. You can re-arrange the icons on this page according to your requirements.
7. Click Create to save the group.

Once the group has been defined, it appears in the Navigator. Selecting the group in the Navigator displays the View, General, and Access pages for that specific group in the detail view. You can modify the group as needed.

To add an object to a group, click on the General tab. Select the desired object(s) from the Available Targets list, and then click Add. Alternatively, you can drag and drop objects within the Navigator to any group.

To add an existing group to another existing group, you can drag and drop one group to another within the Navigator. To add an existing group to a new group, select the desired group from the Available Targets list on the General page and click Add.

Note: When you create a group, it appears at the root level of the Groups folder. This also applies when you create subgroups; even though the subgroup appears at the appropriate level in the group hierarchy, it also appears at the root of the Groups folder.

Figure 4–2 Create Group General Page

Group General Page

This page allows you to define the group and its content. The General page consists of the following:

Group Name

Enter a unique name for the new group in the Group Name field of the dialog box.

Use background

Check this option if you want to display an image in the background of the group's View pane. When checked, you can select one of the standard images from the drop-down list. These standard files are provided in the `oracle_home\classes\oracle\sysman\resources\images` directory located on both the machine running the client Console and the machine running the Management Server. The Java-based Console reads the images from the local images directory. The browser-based Console reads the images from the machine running the Management Server.

The sample images included are:

- `asia.gif`
- `europe.gif`

- japan.gif
- usa.gif
- world.gif

You do not need to specify an explicit path for graphics located in this directory. Specify only the file name, such as `usa.gif`. You can use `.GIF` or `.JPG` graphic files. When a new image is added to a group, the file name is stored in the repository.

Note: If you want to use your own custom images, and make them available to all Consoles across your enterprise, you will need to install the images on all Management Server machines as well as all machines running the Java-based Console. Any new images added to the image directory will not appear in the drop-down list. Administrators wishing to use the new graphic must specify the exact name of the file.

Icons

Determine whether you want to display large or small icons in the group.

Selected Targets

List of targets that have been added to the group.

Available Targets

Tree list displaying all targets and existing groups that can be added to the current group.

Add

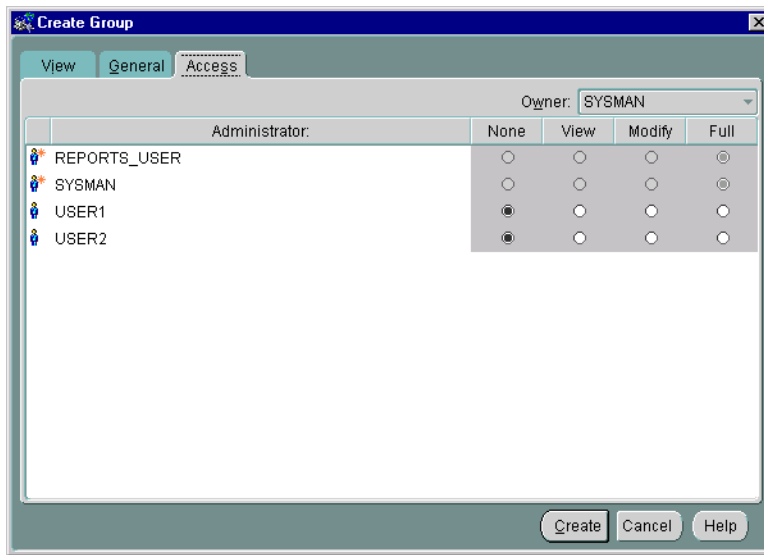
Moves objects selected in the Available Targets list to the Selected Targets list.

Remove

Moves objects from the Selected Targets list back to the Available Targets list.

Group Access Page

Determine the permissions that you want to assign to the group with the Access page. This allows other users to view or modify the group. Any permissions assigned on this page over-write any user default permissions. See "Access" on page 1-23 for an explanation of the permissions that can be assigned.

Figure 4–3 Group Access Page

Manipulating Group Views

Groups are populated by adding targets from the specific group's Available Targets list on the General page. Groups can also be populated by dragging and dropping databases, groups, nodes, listeners or any other discovered targets in the Navigator into the appropriate group in the Navigator Groups folder. You can create, update, and save any number of groups. In addition, you can add a graphic as a background to the group for visual identification or to graphically show the location of nodes. For example, you can use a group drawing of a city or state for the background of your group. Once a group is created, it can be manipulated like any other object in the Navigator.

Monitoring Status

Probably the most important aspect of the Enterprise Manager Group system is that it provides an efficient way to monitor the alerts reported by the event system. If the object or group has events registered against it, a flag showing the state of the event condition is displayed. If an object in a group has more than one event registered against it, the flag will represent the most severe alert condition.

- If the flag is green, there are no problems on the monitored objects.

- If the flag is yellow, there is a condition detected that should be checked.
- If the flag is red, there is a severe problem detected by an event and the object requires immediate attention.
- If the flag is gray, the status of the object is unknown. For example, the UpDown event is registered to a database but the node is unavailable. The gray flag will only display if an up/down test is explicitly registered for the object.
- If there is no flag, then the object is not being monitored. Note that groups within groups do not show a status monitor when the event is clear (green).
- If there is a yellow hexagon, there is an error state (yellow hexagon). An error state indicates there is a problem with the evaluation of the event condition, as opposed to a threshold being met. Examples of error states are: registering an Archive Full event against a database in non-archivelog mode, registering an event that monitors segments but specifying a filter that excludes all available segments.

Groups inherit the worst state of any of the members. If one target in a group is down (status of the target is unknown), the group displays a gray flag.

Note: For status monitors to display on a group or object, an Intelligent Agent must be running on the node where the object is located.

See Chapter 6, "Events" for more information on events.

Expanding Objects

Some group objects, such as database objects and defined group objects, can be expanded in the Group detail view by double-clicking on the object's icon. You can double-click on some objects to open property sheets. The property sheets allow you to both view and alter the properties of the objects. You can double-click on databases to display property sheets.

Some group objects, such as databases, can also be expanded within the Console Navigator. All groups appear in the Navigator's Group container.

See Figure 4-1, "Navigator and Group Panes" for an illustration of an expanded node in the Group pane.

Groups

When you double-click on a group icon, the Edit Group property sheet displays for that group. You can add or remove objects from the group via the Available/Selected Targets lists on the General page. In addition, you can delete objects from the group by selecting them and pressing the Delete key, or clicking the right mouse button and selecting Remove Item from Group.

Databases and Other Discovered Targets

When you double-click on a database or other discovered target within a group, you connect to that service. The instance property sheet displays if that service is a database. If the connection to the database fails for any reason, the Login Information dialog box displays. See "Preferred Credentials" on page 1-25 for information on preferred credentials.

Launching Applications from a Group

You can launch a database application using an object in the group. Select a database icon in the Group View pane, then select a tool from the Console Tools menu or Tool drawers. You are connected to the database according to the user credentials that have been set up for the system. See "Preferred Credentials" on page 1-25 for more information on user credentials.

Adding Objects to a Group

To add objects to a group, drag and drop the object from its location on the Navigator to the group.

Note: Only one copy of an item can exist in the group.

When you add objects in a group, the updates are reflected in every occurrence of the group. Any updates to a group are automatically saved as the updates are made.

Deleting Objects from the Group

To delete objects in a group view pane, you can select the object and press the Delete key, choose Remove Item from Group from the context-sensitive menu, or remove items from the Selected Targets list on the General page. Objects in the

Navigator can be removed via context-sensitive menu option or by selecting the object and pressing the Delete key.

Modifying a Group

You can modify the properties of a group.

1. Select a group in the Navigator.
2. Click the General tab. The name of the group is automatically entered in the name field. The name cannot be modified.
3. Modify the properties for the group.
4. Click the Apply button when you have finished.

Removing Groups

There are two types of group removal:

- Remove the group from within another group. This does not affect the existence of the group in the system or in any other groups.
- Remove the group completely (delete). This removes the group from every group where it has been copied.

After you select a group, there are multiple ways to remove it:

- Delete the group completely by choosing a group (that is not inside another group) and selecting the Delete menu item from the group's context menu.
- Delete the group completely by choosing a group (that is not inside another group) and pressing the <DELETE> key.
- Remove a sub-group from within another group by expanding the desired group in the Console Navigator and choosing Remove Item from group from the context-sensitive menu. This will remove the sub-group from the selected group, but will not affect the sub-group's existence in other groups.
- Remove a sub-group from within another group by selecting the Delete menu item from the group's context menu. This will remove the sub-group from the selected group, but will not affect the sub-group's existence in other groups.
- Remove a sub-group from within another group by pressing the <DELETE> key. This will remove the sub-group from the selected group, but will not affect the sub-group's existence in other groups.

Viewing Reports for Targets within a Group

If you have already configured the Enterprise Manager Reporting website, you can view a variety of generated reports for objects within the group. See Chapter 8, "Enterprise Manager Reporting" for more information.

The Job system allows you to automate standard and repetitive tasks, such as executing a SQL script or executing an operating system command. With the Job system, you can create and manage jobs, share jobs with other administrators, schedule execution of jobs, and view information about the jobs. Jobs can be scheduled on a single node or multiple nodes in the network, provided that the node has an Intelligent Agent running on it. The topics discussed in this chapter include:

- Job Process
- Job Detail View
- Job Menu
- Creating, Modifying, or Viewing a Job
- Oracle Job Tasks

Job Process

The Job process includes:

1. Creating a job. This involves:
 - a. Determining the type and destination of the job.
 - b. Determining the tasks.
 - c. Determining task dependencies of the job.
 - d. Determining the parameters for each task.
 - e. Scheduling the times that the job executes.
 - f. Assigning permissions to other administrators for notification purposes.
2. Submitting a job to the selected destinations in the network system.
3. Viewing the job history to review the results of the job.

Job Tasks

The Job system provides a variety of predefined job tasks, or you can submit your own tasks by executing a SQL*Plus script or running an operating system program. Job tasks are implemented in the Tool Command Language (Tcl) scripts with Oracle extensions (OraTcl) to include database-specific commands. You can write your own Tcl script and submit it with the Run Tcl job task. For more information on custom job scripts, see the *Oracle Intelligent Agent User's Guide*.

The tasks are grouped by the target type of the task:

- Database
- Node
- Listener
- HTTP Server

Note: All target types include Node tasks.

The tasks allow you to perform such operations as:

- Execute operating system commands or shell scripts.

- Execute SQL and DBA commands.
- Perform database administration tasks such as starting up and shutting down Oracle databases.
- Start up and shut down Listeners.

See the online help for Oracle job tasks and "Oracle Job Tasks" on page 5-29 for information on Oracle predefined job tasks and their parameters.

You can combine two or more tasks into one job, called a composite job. Composite jobs consist of separate tasks that are constructed such that some tasks may or may not execute upon completion of another task. For example, if a composite job consists of two tasks, starting up a database and then running a SQL script, you can specify that the script be run only if the database was successfully started. Here, you specify a dependency between the two tasks that determine whether the next task is executed. The Job system allows you to specify one of three dependencies for any task: Always (default), Only on Success, or Only on Failure.

You can create jobs that can be used as *fixit* jobs that are to be run when a condition is signalled by the Event system. Fixit jobs are not scheduled. See Chapter 6, "Events" for information on monitoring events in the system.

Note: You need to set up a password file to perform administration tasks, such as start up or shut down, on a remote database. See the "Configuring a Remote Database for Backup or SYSDBA Administration" in the Oracle Enterprise Manager online help for more information.

Writing SQL*Plus Scripts

Creating your own job tasks using the SQL*Plus command language allows you to automate any number of complex database operations via the Job system and Intelligent Agent(s). The success or failure of a given SQL*Plus script is determined by the exit condition returned when the script completes. Because a SQL*Plus script can return an error condition that may not be recognized by the Intelligent Agent, such as a non-ORA error message, your job task can appear to complete successfully even though a command issued from within your script has failed. To avoid this situation, you must specify an operational clause when using a SQL*Plus EXIT command in your script.

For example:

```
> EXIT sql.sqlcode  
> EXIT warning  
> EXIT 666
```

Using the EXIT command without specifying an operational clause (default) commits and exits with a value of SUCCESS. By explicitly defining the conditions under which your script terminates, the Intelligent Agent can ascertain whether your job has succeeded or failed. For more information on SQL*Plus and the SQL*Plus command language, see the *SQL*Plus User's Guide and Reference*.

Job Credentials

Jobs are normally run with the preferences of the administrator who submitted the job. This ensures that jobs cannot be used to perform functions the administrator could not perform if logged into the machine directly. For example, to write a job output file to the `ORACLE_HOME/network/agent/` directory, the administrator must have permissions to write to that directory on that node.

Because jobs are categorized by the type of target they act on, the job system knows what credentials to pass to the Intelligent Agent. The Job system uses Enterprise Manager System Preferences (Preferred Credentials) to determine what preference information needs to be passed. When a job runs on a node, the job system passes the administrator preferences for the managed node. In addition, the Job system gives you the option of overriding the node preferred credentials when defining a job. This allows the creator of the job to run the job on the node using a different username and password. When a job runs on a service, such as a database, the job system also passes the administrator preferences for the service. See "Preferred Credentials" on page 1-25 for information on administrator preferences.

Important: You must set up valid user credentials for the nodes on which you want to run jobs. Node credentials are required for all jobs. If credentials are not set up correctly for an Windows NT node, you may get the "Failed to Authenticate User" error message. Make sure the Windows NT node account specified in your preferred credentials possesses the "log in as batch jobs" user right. See the *Enterprise Manager Configuration Guide* or your Windows NT documentation for instructions on creating NT user accounts.

Administrator preferences for nodes and/or services are given the following prioritization:

1. Preferred Credential Overrides (Highest Priority)
2. Preferred Credentials for the selected node or service.
3. Preferred Credentials for the Default target.
4. Null Credentials (Lowest Priority) **Note:** Jobs and events will fail.

Submitting Jobs

The Job system is simple to use because the task of scheduling and managing jobs is centralized in the Enterprise Manager Console. The administrator only needs to submit a job once, regardless of the number of targets on which the job will run.

When you submit a job, the Management Server sends the job information to the appropriate Intelligent Agents on the targets you selected. The Intelligent Agents are responsible for executing the job on the specified schedule and returning job status messages to the Console through the Management Server. Once submitted, jobs will run regardless of whether you are logged in or not.

Note: There is usually a slight delay between submitting the job and the notification by the Intelligent Agent.

To schedule a job, you do not have to connect at the time of job creation to the node on which the job will be run. You only need to submit the job from the Console and specify the targets on which it should run. The target can include nodes, databases, listeners, user-defined groups that have been created with the Group system, or any other discovered services.

The Job system of Enterprise Manager allows you to efficiently run jobs on multiple remote nodes by transferring job information to the Intelligent Agents servicing the nodes. When a job is executed, it is run by the Intelligent Agent on that node, thus minimizing network traffic between the remote node and the Console and Management Server. In addition, jobs can be run on multiple nodes simultaneously because there is an Intelligent Agent residing on each node. Jobs can only be run on nodes where an Intelligent Agent is running. If you send a job to a group, the job is only scheduled on the nodes in the group where the Intelligent Agent is running.

When you submit a job to one or more remote sites, it is possible that any one of those site may be down. If a site or its Intelligent Agent is down, the Management Server queues any job requests that could not be delivered to the site. Once the site can be contacted, the Management Server submits the queued job to the Intelligent Agent, which in turn executes the job on the node.

Important: If the Agent goes down for any period of time while the job is running, the Agent will return a job status message "Agent was down." This does not affect job operation, but merely indicates that the repeating job did not run as scheduled during the period the Agent/node was down.

Cancelling Jobs

You can cancel a job by selecting the desired job from the Job Active window and choosing Remove Job from the Job menu. Care should be taken when cancelling a job. Cancelling a job that is currently running will interrupt the job process and terminate the job. Problems can arise when the job is composed of multiple job tasks with sequential dependencies.

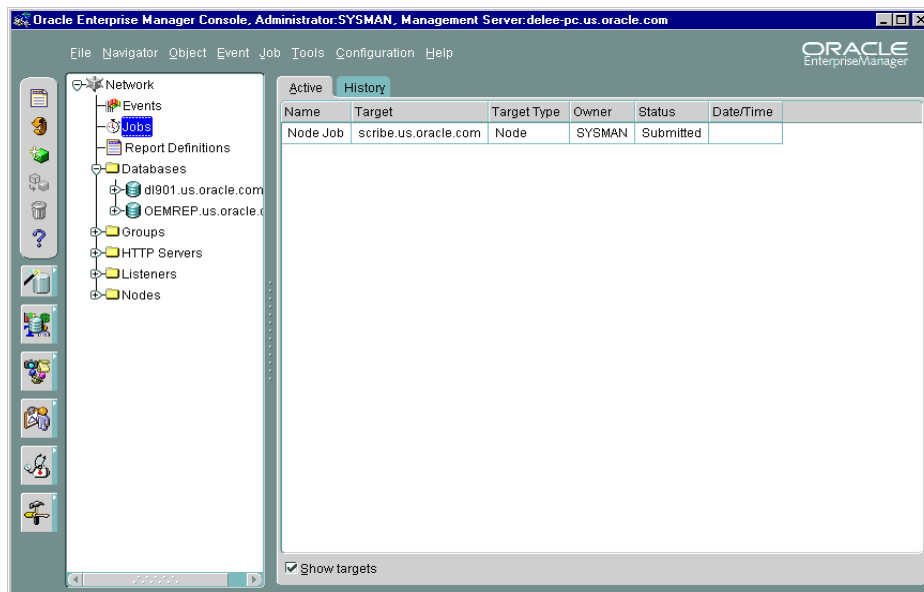
Job Detail View

You can view the different pages of job information by selecting the page tabs in the Job detail view. There are two pages in the detail view:

- Active
- History

You can switch between the pages by clicking the tab of each page. The rows in both pages can be sorted on any column by clicking the column heading.

Figure 5–1 Job Detail View in the Console



Active Job Page

The Active page contains a summary of the active jobs on the network. These are jobs that you have submitted to the job system and are not yet completed. Each row is an execution of a particular job scheduled on a specific destination. While a job may execute multiple times, the job listed in the Active page is the one that is currently scheduled or running. You can use the Edit Job menu option to display the details of the selected job. When selecting a job from the Active pane, the only job attribute you can modify is the job permission. See "Creating, Modifying, or Viewing a Job" on page 5-12.

You can double-click on a job listed in the Active page to view the job details.

Name

Name of the job.

Target

Destination of the job. (Displayed only when Show Targets is selected.)

Target Type

database, listener, node, HTTP Server, or other managed targets.

Owner

Administrator submitting the job.

Status

Status of job (Displayed only when Show Targets is checked.) is one of the following:

- *Submitted*: The job has been submitted to the job target running an Intelligent Agent.
- *Scheduled*: The job has been successfully delivered to the Intelligent Agent and is scheduled for execution.

Note: Under certain circumstances, a job will remain in a Submitted state. The most likely cause would be the Intelligent Agent on the node to which you are trying to submit the job is down, or the node is not connected to the network.

- *Started*: The job execution has started. After the job executes, the job execution is displayed in the Job History page. If this is the last scheduled execution of the job, the job is removed from the Active Jobs page. Otherwise, the job remains in the Active Jobs page and has the status of Scheduled. Unless you view the Active Jobs page at the exact time that the job is executing, you would not see the Running status.
- *Pending Deletion*: The job has been selected for deletion. When the deletion is successful, the job is removed from the Active Jobs page and added to the Job History page.
- *Fixit*: The fixit job has been submitted.
- *Fixing*: The fixit job is executing. A fixit job remains in the Active Jobs page until it is deleted.

Date/Time

This is the time the Intelligent Agent returns after the job has been scheduled by the Intelligent Agent.

Context-sensitive Menu options

The following options are available by selecting a job and clicking the right mouse button.

View Published Reports

Displays the Job status report. The Enterprise Manager Reporting environment must be configured before using this menu option. See Chapter 8, "Enterprise Manager Reporting" for more information.

Edit/View Job

Displays the property sheet for the selected job. Only the permissions, notifications, and target list can be changed.

Create Like

Displays a copy of the selected job. You can edit the property sheet and submit this job.

Copy to Library

Copies the selected job to the Job Library if it does not already exist there.

Remove Job

Delete the selected job. If the job has not been saved to the Job Library, a warning dialog appears indicating that you are about to perform a remove operation. You may wish to copy the job to the job library before removing it.

History Page

Job History contains a list of previous job executions. These are jobs that have been submitted to an Intelligent Agent and have executed successfully or unsuccessfully. You cannot modify these jobs.

- Job Name is the name of the job.
- Target of the job.
- Owner of the job.
- Target type
- Status of job is one of the following:
 - Completed: The job has executed successfully.
 - Failed: The job execution has failed.
 - Deleted: The job has been deleted.

- For the other status categories, see "Active Job Page" on page 5-7.
- Finish Time is the time when the job finished, failed, or was deleted.

Refreshing the History Page

You can refresh the job history list at any time by clicking on the Refresh icon in the Console toolbar or choosing Refresh Job History from the Console's Job menu. Job history is automatically refreshed each time you move from the Active tab to the History tab.

Clearing the History Page

You can clear the History page by choosing Clear Job History from the Console's Job menu, or from the context-sensitive menu.

Displaying Job Output

You can double-click on a job listed in the Job History page to display the Job property sheet and view the Job Output dialog box, if output exists for the job. If no output is produced by a job, a message displays that states that there is no output for the job. If the output includes only blank spaces, the dialog box is blank.

Context Menu options

The following options are available by selecting a job and clicking the right mouse button.

View Job

Displays the property sheet for the selected job.

Create Like

Displays a copy of the selected job. You can edit the property sheet and submit this job.

Copy to Library

Copies the selected job to the Job Library if it does not already exist there.

Remove Job

Delete the selected job. If the job has not been saved to the Job Library, a warning dialog appears indicating that you are about to perform a remove operation. You may wish to copy the job to the job library before removing it.

You can save the jobs from the History page to a file, then clear the jobs from the History page. This prevents the History page from being overloaded with obsolete

jobs that occurred in previous days. Choose the Report Job History option from the Job menu to save the history to a report. This report is a file on the machine where the Enterprise Manager Console is running.

You can create a new job similar to a job in the History page with the Create Like option. See "Creating, Modifying, or Viewing a Job" on page 5-12 for more information.

Job Menu

The Job menu allows you to create, modify, save, submit, and manage jobs. The menu options are enabled depending on the items selected in the Job pane. See Figure 5-1, "Job Detail View in the Console" for an illustration of the Job menu.

Create Job

Allows you to create a new job. See "Creating, Modifying, or Viewing a Job" on page 5-12 for more information.

Create Job Like

Allows you to create a new job like the selected job in the Job pane. See "Creating, Modifying, or Viewing a Job" on page 5-12 for more information.

Edit Job

Allows you to modify the job selected in the Job pane. The property sheet is the same as the property sheet for creating a new job, however, you can only modify job permissions. See "Creating, Modifying, or Viewing a Job" on page 5-12 for more information.

Copy Job to Library

Copies the selected job to the Job Library if it does not already exist there.

View Job

Displays the property sheet for the selected job in the Job pane. The property sheet is in read-only format. Active jobs can be removed but not modified. See "Creating, Modifying, or Viewing a Job" on page 5-12 for more information.

Remove Job

Removes the selected job from the Active or History page of the Job pane. When you delete a job, there is usually a slight delay while the request is processed. See "Cancelling Jobs" on page 5-6 for specifics on job cancellation.

Job Library

Displays the Job Library dialog. See "Job Library" on page 5-12 for more information.

Clear Job History

Clears the jobs listed in the Job History page.

Refresh Job History

Refreshes the job history list. Job history is refreshed each time you move from the Active tab to the History tab. However, to refresh the job history list while currently viewing the History pane, you must select this menu item.

Remove Job

Clears the selected job listed in the Job History page.

Job Library

The Job Library dialog contains a list of the jobs that you have created and saved. In the dialog, you can view summary information about a job:

- **Job Name** is the name of the job.
- **Description** is the user-supplied description of the job.
- **Owner** is the administrator that is assigned as the owner of the job.

These jobs can be submitted to the job system with the Submit button. You can use the Edit button to modify a job selected in this page. You can also double-click on a job listed in the Job Library page to edit the job. You can create a new job based on an existing job with the Create Like button.

Creating, Modifying, or Viewing a Job

When you create, modify, or view details of a job, similar property sheets display. See Figure 5-3, "Job Tasks Page" for an illustration of a Job property sheet. The property sheets contains:

- General Page
- Task Page
- Parameters Page
- Schedule Page
- Access Page

- Progress Page (this page only appears when a job is selected from the Active or History pane)

Attention: When submitting a job that consists of multiple tasks, an error may occur if the string of arguments that was sent is longer than the internal buffer. If that error occurs when submitting a job, divide the tasks among multiple jobs and resubmit the jobs.

Creating a New Job

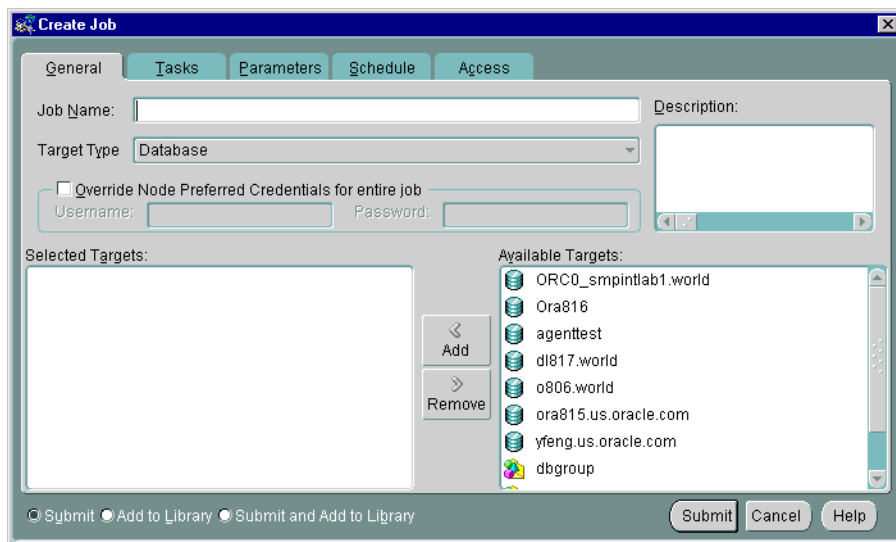
1. Select Create Job from the Job menu to display the Create Job property sheet.
2. Complete the pages of the Create Job property sheet.
3. Determine whether the job is ready to submit.
 - a. Select the Submit option and click Submit to send the job to the Intelligent Agents at the selected destinations. The job appears in the Active page.
- OR -
 - b. Select the Add to Library option and click Add. The job appears in the Job Library. You can modify or submit a saved job at a later time.
- OR -
 - c. Select the Submit and Add to Library option and click Submit and Add to send the job to the Intelligent Agents at the selected destinations and save the job to the job library. The job appears in the Active page and the Job Library. You can modify or submit a saved job at a later time.

Note: There is usually a slight delay between the time you submit the job and Intelligent Agent notification.

General Page

The General page allows you to specify the primary attributes of a job, such as, job name, description, target type, targets, etc.

Figure 5–2 Job General Page



Job Name

Enter the name of the new job.

Description

Enter the optional description of the job.

Target Type

Select the target type from the pull-down list: Database, Listener, Node, Group, HTTP server, or other service that is integrated into the Console.

Override Node Preferred Credentials for Entire Job

Allows the submitted job to bypass the Console’s current preferred credential settings and use a specified Username and Password.

Selected Targets/Available Targets

Select the targets of the job in the Available Targets list and click the Add button to move the target to the Selected Targets list. The targets are determined by the Job Type. The targets include databases, listeners, nodes, HTTP Servers, and groups of these objects.

- For an operating system task, a list of nodes and groups containing nodes displays.
- For a database task, a list of databases and groups containing databases displays.
- For a listener task, a list of listener and groups containing listeners displays.
- For an HTTP Server task, a list of HTTP Servers and groups containing HTTP Server displays.

Note: Only network objects that have been discovered correctly and have an Intelligent Agent running are included in the list of available targets. See "Discovering Targets" on page 3-6 for more information.

Job Task Page

The Task page allows you to choose the tasks that you want the job to perform.

Available Tasks

Tree list of available job tasks. Tasks vary depending on the Target Type selected on the General page. Select a task and click on the < (Add) button to include the task in the job. You can add multiple tasks to the job from the Available Tasks scrolling list. See the online help for Oracle job tasks and "Oracle Job Tasks" on page 5-29 for information on Oracle predefined job tasks and their parameters.

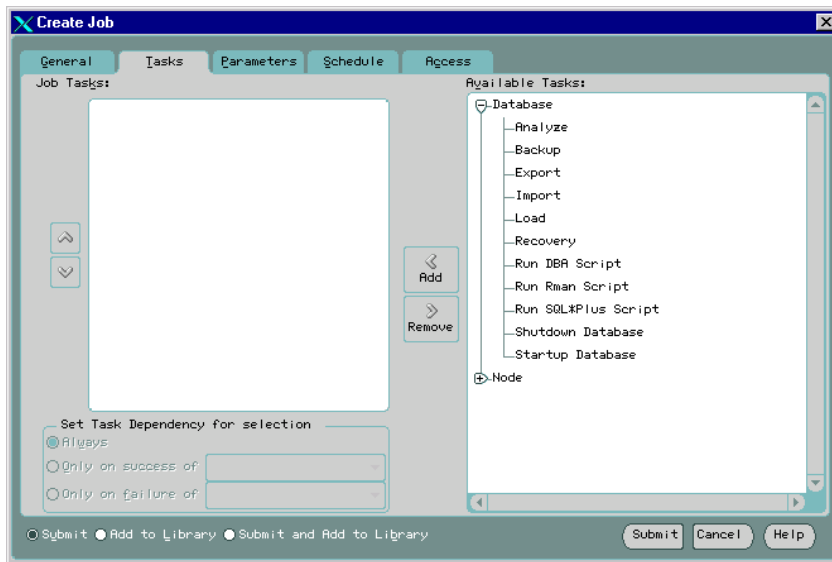
Job Tasks

Tasks currently selected for the current job. You can remove tasks from this list by selecting the task and clicking on the >> (Remove) button.

Up/Down Arrows

Use the arrow buttons to change the order of the tasks or to make a task conditional on a previous task. Select a task in the Job Tasks list and click on the up or down arrow button to position the task.

Figure 5–3 Job Tasks Page



Conditions for running a task

Select a task in the Job Tasks window and select from the following:

- Always - The task is always executed regardless of the failure or success of other tasks.
- Only on success of - The task is only executed if the task selected in the list is successful.
- Only on failure of - The task is only executed if the task selected in the list has failed.

Note: If any task does not execute, control moves to the next task on the same level as the task that did not execute.

Important: If your targets are running with pre-9i Intelligent Agents, you can add up to five job tasks per job. Adding more than five tasks will generate an error message stating the maximum number of input files have been exceeded. This limitation does not apply to targets running 9i Intelligent Agents.

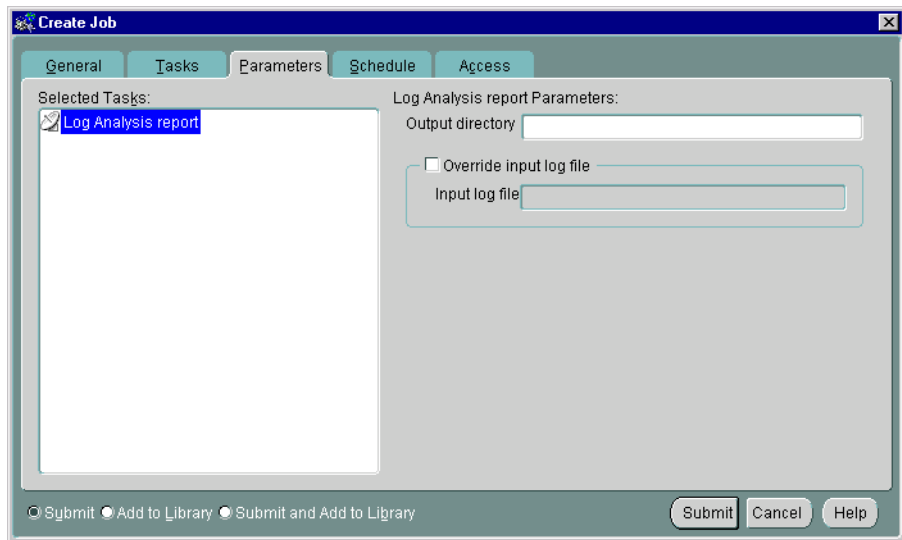
Stopping Jobs

Normally, a job stops when all tasks have been executed. In some cases, you may want the job to stop before completing the task sequence, which might be the case when one task in the sequence fails. To handle situations like this, you can include the Halt Job task as part of a composite job to stop execution at any point within the task sequence. You specify the Halt Job task dependency as you would any other task.

Job Parameters Page

On the Parameters Page, you specify parameter settings for the selected job tasks. To set the parameters for a task, select the task in the Selected Tasks list. The parameters for the selected task are displayed on the right side of the Parameters Page.

Figure 5–4 Job Parameters Page



Selected Tasks

Select the task for which you want to set parameters.

Task Parameters

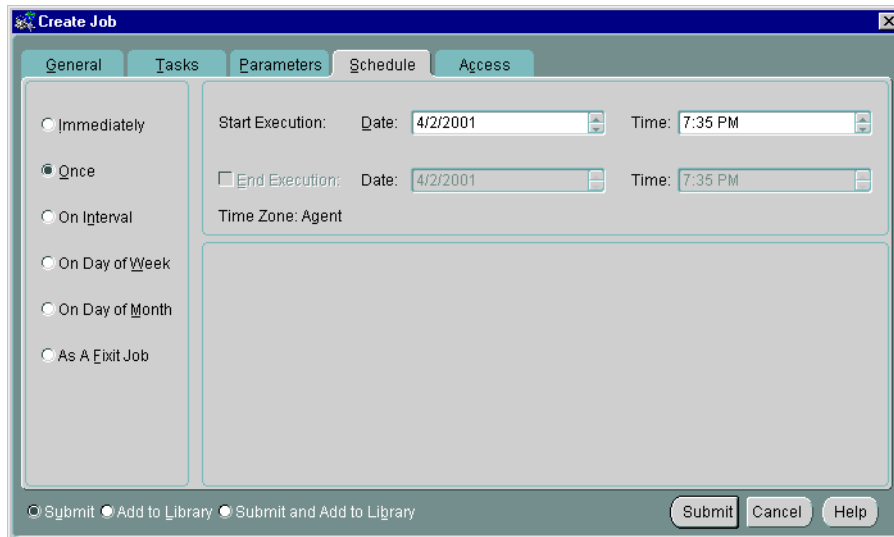
Specify the parameters for the selected task. You can enter values in the entry boxes or select values from the pull-down lists. The parameters vary according to the job task. See the online help for Oracle job tasks and "Oracle Job Tasks" on page 5-29 for information on Oracle predefined job tasks and their parameters.

For some jobs, you can override the preferred credentials for connecting to the service. This allows you to enter a username and password. See "Job Credentials" on page 5-4 and "Preferred Credentials" on page 1-25 for information on administrator preferences.

Job Schedule Page

The Schedule page allows you to schedule the execution of the job task. Select the frequency that you want the task executed. The choices are Immediately, Once, On Interval, On Day of Week, On Date of Month, or As a Fixit Job.

Figure 5-5 Job Schedule Page



Immediately

Submits the task as soon as you finish the setup process. The task executes only one time.

Once

Schedules the task only one time at the date and time you choose.

On Interval

Allows you to schedule a specific time interval between task executions. The interval can be a combination of hours and minutes, or number of days. Select the value you want to change and click on the scroll buttons. You can also type in a new value.

On Day of Week

Allows you to schedule the task on one or multiple days (Sunday, Monday, etc.) of the week. Click on the days of the week to select the days you want the task scheduled.

On Date of Month

Allows you to schedule the task on one or multiple days (1 - 31) of the month. Click on the dates of the month to select the dates you want the task scheduled.

Note: If you choose a day, such as 31, that is not in a month, the job will not be run in that month.

As a Fixit Job:

Check this box if you want to use this job as fixit job to correct an event condition. The fixit job must be submitted to the target where the event is being monitored. A fixit job cannot be scheduled.

The job can be selected from the Fixit Job list in the Event property sheet's Fixit Jobs page after it has been successfully submitted to an Intelligent Agent. For information on fixit jobs with events, see "Event Fixit Jobs Page" on page 6-54.

Start Execution

Choose the first date and time that you want the task executed. This is the starting time for any task scheduled on an interval.

Select the month, day, or year in the Date field and click on the scroll buttons to change the value. You can also type in new values.

Select the hour, minute, or AM/PM in the Time field and click on the scroll buttons to change the value. You can also type in new values.

End Execution

Choose the last date and time that you want the task executed. This option does not apply if you chose the Immediately or Once execution options.

- Select the month, day, or year in the Date field and click on the scroll buttons to change the value. You can also type in new values.
- Select the hour, minute, or AM/PM in the Time field and click on the scroll buttons to change the value. You can also type in new values.

Time Zone

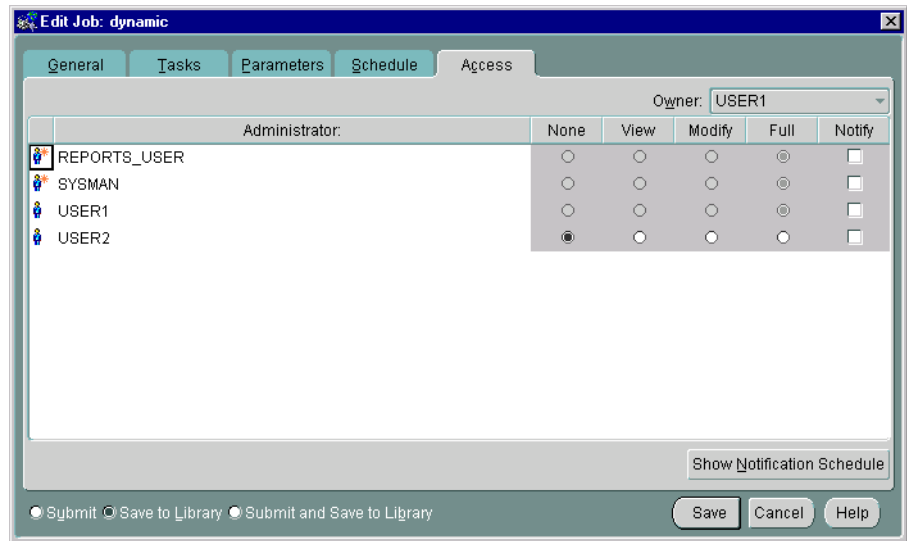
Static text displaying the time zone to be used for job execution.

Note: Only the Agent time zone is available with this release. Here, the Intelligent Agent schedules the task execution at each target based on the actual system time of each Intelligent Agent. Tasks are not necessarily run simultaneously.

Job Access Page

Determine the permissions that you want to assign to administrators with the Access page. This allows other administrators to view or modify the job. Notifications can be assigned with this page. Any permissions assigned on this page supersede any user default permissions that have been assigned with User Preferences.

Figure 5-6 Job Access Page

**None**

Does not allow the administrator to view this job anywhere.

View

Allows the administrator to view the job and inspect job properties.

Modify

Allows the administrator to suspend and resume the job. The administrator will also be able to enable notifications for active jobs.

Full

Allows the administrator to delete the job from the Active or History pages of the Job window, modify permissions for other administrators, and enable notification.

Notify

Allows the administrator to receive enhanced notifications (paging and email) for the job. Notify permission cannot be assigned if the administrator's permission level is set to None.

Note: All administrators possessing at least the View permission will receive notifications through the Console when a job is executed.

Any permissions assigned on this page supersede any administrator default permissions. See "Access" on page 1-23 for more information.

Show Notification Schedule

Show Notification Schedule displays when administrators are scheduled to receive notifications.

To remove an administrator from receiving an enhanced notification, display the context menu (press the right mouse button) on any time block. The context menu provides options for adding and removing recipients of the notifications. Administrators that appear via the Add Recipient and Remove Recipient menu options are those whose schedules match the selected time slot.

Skipped Job Notification

When using a 9i version of the Intelligent Agent, notifications are also sent to administrators whenever a job is skipped, as might be the case when the Intelligent Agent is unavailable, when the target is blacked out (see the *Intelligent Agent User's Guide* for more information), or when the length of previous job executions run beyond the time the next execution is scheduled to begin.

Job Progress Page

The Job Progress page is available for jobs in the Active and History pages of the Job window. The Progress page provides a log of the job's activities.

The Progress page contains all status changes that have been received for a specific job. Each row in the page summarizes a status change of the job. If you display the Progress page for an execution in the History page, the page typically displays Submitted, Scheduled, Started, and Completed or Failed statuses for that execution. If you display an execution from the Active Jobs page, the Progress page displays only those status changes that have been received.

When you display the Progress page, the page displays the statuses only for the target and execution time of the job occurrence selected. To view the status changes associated with other targets or execution times, select other targets or execution

times from the Target or Execution pull-down lists. You can also select <All> in either list to view all statuses. If the job has been Deleted on a destination, the Deleted status always displays at the top of the Progress Page.

The following options are available on the Progress page:

Target

Select the target of the job occurrences you want to view from the pull-down list. Select <All> for all destinations. The list of job occurrences changes according to the selection.

Execution

Select the execution time of the job occurrences you want to view from the pull-down list. Select <All> for all executions. The list of job occurrences changes according to the selection.

Save List

Select the Save List button to save the list of job occurrences as a local file using the standard Windows file dialog box.

Notification Details

This option displays the notifications for this job.

Show Output

If output exists for the selected occurrence of a job, you can display the output in the Output dialog box. You can also double-click on a selected job to display output.

The columns in the Progress page contain the following information:

Status

See "Active Job Page" on page 5-7 and "History Page" on page 5-9 for information on the status of a job.

Target

This is the target of the specific occurrence of the job.

Date/Time

This is the time the Management Server was notified by the Intelligent Agent that a status change had occurred.

Job Output Dialog Box

The Output dialog box is displayed by double-clicking on an entry in the Progress page, or selecting an entry and clicking on the Show Output button, if enabled. This

dialog displays any output, including error messages, as a result of the execution of an occurrence of a job. If no output is produced by a job, a message displays stating there is no output for the job. If the output includes only blank spaces, the dialog box is blank.

With the Output dialog box displayed, the following options are available:

Close

Select the Close button to exit the dialog box after viewing it.

Save As

Select the Save As button to save the contents of the dialog box to a text file.

Changing the Maximum Limit for Job Output Because the output of some jobs may be quite large, Oracle Enterprise Manager provides you with the option of specifying the maximum size for any job output returned by the Intelligent Agent. If the job output exceeds the maximum limit, any output generated beyond this maximum is truncated. By default, the maximum job output size is set to 128K (smallest allowable value). The largest permissible value is two megabytes.

You can set this parameter (`oms.vdg.joboutput.maxsize`) in the following file:

```
$ORACLE_HOME/sysman/config/omsconfig.properties
```

Example omsconfig.properties entry:

```
oms.vdg.job_output_maxsize=128K
```

Units must be specified in kilobytes (K) or megabytes (M).

Modifying Active Jobs

Certain job properties can be dynamically modified for submitted jobs in the Active Jobs page. These job properties are: job targets, job permissions, and enabling notification. Dynamic modification of these properties means you can select the active job, change the properties and apply the changes. The changes will be dynamically applied to all targets of the job. The other job properties (tasks, task parameters, and job schedule) are currently not dynamically modifiable. See the next section *Modifying Active Jobs* for more information.

The owner of the job, i.e., the Enterprise Manager administrator that submitted the job, can change any of the dynamically modifiable properties: targets, permissions, notification. Administrators who have "Full" permission for the job can change the

job permissions and enable notifications for any administrator. Administrators who have "Modify" permissions can enable notifications for any administrator.

Alternative Method of Modifying an Active Job

If you want to change other job attributes besides targets, permissions, and notifications, you need to first delete the job from the Active page, then re-submit the job with the necessary changes. It is more convenient to save the job definition first in the Job Library, and then make the necessary changes. You must be the job owner in order to modify it.

1. If you have not already done so, copy the job to the Job Library using the Copy to Library menu option.
2. Select Job Library from the Job menu. The Job Library dialog appears.
3. Select the job from the library.
4. Click Edit. The Job property sheet appears.
5. Update the pages of the Job property sheet and determine whether the job is ready to submit.
 - a. Select the Submit option and click OK to submit the job to the Intelligent Agents at the selected destinations. The job appears in the Active window.
- OR -
 - b. Select the Save to Library option and click OK. The job appears in the Job Library. You can modify or submit a saved job at a later time.
- OR -
 - c. Select the Submit and Save to Library option and click OK to submit the job to the Intelligent Agents at the selected destinations and save the job to the job library. The job appears in the Active window and the Job Library. You can modify or submit a saved job at a later time.

Viewing Job Details

To view the details of a particular job, double-click on a job in the Active or History Job pane. The Job property sheet for that job appears.

Example: Creating a Job

This example illustrates how to complete the General, Tasks, Parameters, and Schedule Pages when creating a job. It also describes how to save and submit a job.

1. Select Create from the Job menu to create a new job. The Create Job property sheet displays.
2. Enter a name for the new job in the Job Name field of the General Page. You may also enter a description for the job in the Description field.
3. Select Database from the Target Type pull-down list.
4. Select a target from the Available Targets list, then click on the << (Add) button to add the target to the Selected Targets list. These are targets where an Intelligent Agent is running. The Job Credentials must be set up correctly for these targets. See "Job Credentials" on page 5-4 for more information.
5. Repeat the previous step for another target. You are choosing the targets where the job will be run.
6. Click on the Task Page tab of the Create Job property sheet.
7. Select Run SQL*Plus Script from the Available Tasks list, then click the << (Add) button to add the task to the Selected Tasks list. For this example, add only the Run SQL*Plus task. You can specify multiple tasks for a job and make a task conditional on a previous task.
8. Make sure the task is set to run Always. If you specified multiple tasks for this job, you could make a task conditional on a previous task.
9. Click on the Parameters Page tab of the Create Job property sheet to set the parameters:
 - Enter "SELECT * FROM dba_users;" in the Script Text box.
 - Do not check the Override Preferred Credentials box. If you do not override the credentials, the information that is set up with the administrator Preferred Credentials property sheet is used. See "Job Credentials" on page 5-4 for permissions needed to run job tasks. This particular job task requires database credentials that permit access to the data dictionary views. Make sure the Preferred Credentials for the chosen target have this.
10. Click on the Schedule Page tab of the Create Job property sheet to schedule the execution of the job.
11. Select the On Interval under Run Job to execute the job on a specific interval.

12. Set the Start Execution Date of the job to 9/1/2001. Select the month value in the Start Execution Date field and enter 7. You can also click the up and down arrows to change the value when a number is selected.
13. Repeat the process in the previous step for the day and year values of the Start Execution Date.
14. Change the Start Execution Time to 12:00 AM. Use the same procedure that you used to set the date.
15. Check the End Execution box to set a final execution date for this job.
16. Set the End Execution Date of the job to 9/3/2001. Set the End Execution Time of the job to 12:00 AM.
17. The Time Zone should be set to Intelligent Agent (currently static text). This job will execute at the local time zone where the Intelligent Agent is located. This is the time zone of the destination.
18. Select the Every ... Days button to set the job frequency for day interval. Click on the up arrow to change the value to 3 days.
19. Click on the Access Page tab of the Create Job property sheet to assign permissions on this job for other administrators.
20. Allow all other Enterprise Manager administrators to view and receive notifications for this job.
21. Click the Submit and Add to Library button to save the new job in the Job Library and submit the job to the selected destinations.
22. If you want to modify this job later, select the job in the Job Library dialog and choose the Edit option. You can also double-click on the job.

A submitted job is sent to the Intelligent Agents at the selected destinations. The Intelligent Agent for a target begins processing the job, the job appears in the Active Jobs page in the Job window. If the job is processed successfully, the job will start executing on 9/1/2001 at 12:00 AM. After an execution of a job, it is moved to the Job History page of the Job window. You can view the progress of the job and any output in the Job Progress Page.

Note: If you have a domain user set up, you must set the domain password to be the same as the local password in order for scheduled jobs to run when they are submitted using the domain user account. (Windows NT only)

Required Administrator Permissions

The following table summarizes the permissions required to perform various activities against jobs in the Enterprise Manager Console. The owner of a job refers to the Enterprise Manager administrator that submitted the job.

Table 5–1 Administrator Permissions for Jobs

Action	None	View	Modify	Full	Owner	Super User	Comments
JOBS - Dynamic modification							
Add/remove targets	No	No	No	No	Yes	No	
Change permissions, including your own permissions	No	No	No	Yes	Yes	Yes	
Set Notification checkbox for any administrator	No	No	Yes	Yes	Yes	Yes	
Change all other job properties: tasks, parameters, fixit job, schedule	No	No	No	No	No	No	Not supported in 9i
Change owner	No	No	No	No	No*	No	* New behavior. When administrator is deleted, all jobs are reassigned to the new owner.
JOBS - In the Library							
Change owner (library)	No	No	No	Yes	Yes	Yes	
Add/remove targets	No	No	Yes	Yes	Yes	Yes	
Change description, tasks, parameters, schedule, fixit job	No	No	Yes	Yes	Yes	Yes	
Change permissions; enable/disable Notify preferences	No	No	Yes	Yes	Yes	Yes	
Delete job	No	No	No	Yes	Yes	Yes	
Submit job from the library	No	Yes	Yes	Yes	Yes	Yes	
JOBS - In the Console							
Delete active job	No	No	No	No	Yes	Yes	
Clear history	No	No	No	Yes	Yes	Yes	

Oracle Job Tasks

This section lists the Oracle predefined job tasks and parameters for:

- Oracle databases
- Operating systems or hosts (Nodes)
- Listeners
- HTTP Servers

This information is entered in the Job Task Page and Job Parameters Page of the Create Job property sheet. The name and the parameters are listed for each task.

Oracle Database Tasks

These are the tasks that can be run on databases and database groups. In addition, you can run operating system or host job tasks.

- Run SQL*Plus Script
- Run DBA Script
- Shutdown Database
- Startup Database

Note: You need to set up a password file to perform administration tasks on a remote database. See the *Oracle Enterprise Manager Configuration Guide* for more information.

Enterprise Manager Wizard Database Tasks

Specific database job tasks are used by the Enterprise Manager data and backup management wizards that are accessed through the Console Navigator popup and main menus. The job tasks are:

- Analyze
- Import
- Export
- Load
- Backup
- Recovery

Since these job tasks are used only by their respective wizards, the tasks will not have directly editable parameters. To modify parameters for any Wizard database task, click on the Load Wizard button located on the Parameters page of the Job property sheet to activate the associated Wizard. For more information about these wizards and associated job tasks, see the Enterprise Manager online help.

Run SQL*Plus Script

This job executes a SQL*Plus script, allowing any legal SQL or PL/SQL scripts to be run, including all SQL*Plus formatting commands. You can copy and paste the text of the script into the Script Text box of the Parameters page (Create Job property sheet), or simply type SQL commands in the Script Text box.

Parameters:

1. SQL Parameters. Enter one or more arguments that you want the script to use.
2. Override Preferred Credentials. You can use the preferred credentials that have been set up for the database, or you can enter a username and password. If you check the box to override the credentials, then you need to enter:
 - a. User Name. Username for accessing the database.
 - b. Password. Password for the username.

Note: See "Preferred Credentials" on page 1-25 for more information.

3. Script Text. You can copy and paste your script into the Script Text box.

Hint: If you need to determine whether a SQL error has occurred during the running of a SQL script, include "WHENEVER SQLERROR EXIT SQL.SQLCODE" at the beginning of the script. If a SQL error occurs, the job status is set to failed.

Run DBA Script

This job executes a Server Manager line mode script that contains DBA commands.

Parameters:

1. Override Preferred Credentials. You can use the preferred credentials that have been set up for the database, or you can enter a username and password. If you check the box to override the credentials, then you need to enter:

- a. User Name. Username for accessing the database.
- b. Password. Password for the username.
- c. Connect As. Select the role you want to connect as from the pull-down list.

Note: See "Preferred Credentials" on page 1-25 for more information.

2. Script Text. You can copy and paste your script into Script Text box.

Shutdown Database

This job task shuts down an Oracle database instance.

Parameters:

1. Mode:
 - Immediate
 - Abort
2. Connect As:
 - SYSDBA
 - SYSOPER
3. Override Preferred Credentials. Check the box if you want to override the preferred credentials that have been set up for the database. If you check the box to override the credentials, then you need to enter:
 - a. User Name. Enter the username for accessing the database.
 - b. Password. Enter the password for the username.

Note: See "Preferred Credentials" on page 1-25 for more information.

Startup Database

This job task starts up an Oracle database instance.

Parameters:

1. Startup State. Select the start up state from the pull-down list:
 - Startup instance, mount, and open database

- Startup instance and mount database
 - Startup instance only
2. **Parameter File.** Enter the initialization parameter filename you want to use for the database. This file is located on the node where the Intelligent Agent and database reside. For example with a database on a Unix platform:

```
/private/oracle/admin/ora8db/myinit.ora
```

If you do not enter a filename, the default platform-specific initialization file is used.

3. **Override Preferred Credentials.** Check the box if you want to override the preferred credentials that have been set up for the database. If you check the box to override the credentials, then you need to enter:
- a. **User Name.** Enter the username for accessing the database.
 - b. **Password.** Enter the password for the username.

Note: See "Preferred Credentials" on page 1-25 for more information.

4. **Mount Mode.** Select the mount option from the pull-down list:
- Exclusive
 - Normal
 - Parallel
5. **Connect As.** Select the connecting role from the pull-down list:
- SYSDBA
 - SYSOPER
6. **Restrict Connections.** Check this box if you want to start the database in Restricted mode.
7. **Force Startup.** Check this box if you want to start up the database with the Force option. Refer to the Oracle 8i Administrator's Guide for more information on startup options.

Halt Job

Normally, a job stops when all tasks have been executed. In some cases, you may want the job to stop before completing the task sequence, which might be the case when one task in the sequence fails. To handle situations like this, you can include the Halt Job task as part of a composite job to stop execution at any point within the

task sequence. You specify the Halt Job task dependency as you would any other task.

Operating System or Node Tasks

These are the tasks that can be run on the host's operating system.

- Broadcast Message
- Run OS Command
- Run Tcl Script
- Halt Job

Broadcast Message

This job allows you to submit a message to the selected target using the platform-specific mechanism. To send the message to a target, you may need to have permissions on specific directories. For example, you may need permissions on `/dev/console` (system console device) to send a message to a Unix destination.

Note: On a Windows platform, this task sends the message to ALL users on the network. To send a message to specific users, use the Run OS Command task to execute the `net` command with the `send` option. See the Windows online help for information on `net` command line arguments. You can also enter

```
net send /help
```

at the MSDOS command prompt.

Parameters:

Message Text. Enter the message text that you want sent to the selected destinations.

Run OS Command

This is a generic method of running any program or script that is executable on that host, provided your credentials allow you to do that.

Parameters:

1. OS command or shell script name. The command or script must be accessible from the node where the Intelligent Agent and database reside. You may have to include the path for the Intelligent Agent to locate and execute the command or script. For example: `ls`

2. One or more arguments to the command. For example: `-1 /export/oracle`

Run Tcl Script

This job executes a Tcl script. This is a generic method of running any Tcl script that is executable on that host, provided the preferred credentials allow that. See "Preferred Credentials" on page 1-25 for more information.

Parameters:

1. Parameters. One or more command-line arguments that you want the script to use. The arguments must be delimited by quotes.

Note: Multiple parameters, such as "one two three", are treated as only one parameter. To ensure that the parameters entered in the field are treated as separate arguments and to ensure that the Tcl script functions in future releases, include the following at the beginning of the Tcl script:

```
set argc [llength $argv]
if { $argc == 1 } { set argv [lindex $argv 0]}
```

2. Script Text. Type or copy the Tcl script into the Script Text box. To use an existing script file, click Import to display the Import File dialog.

Tcl Script Examples

For information on writing Tcl job tasks, see the *Intelligent Agent User's Guide*. For information on Tcl, see "Tcl and the Tk Toolkit," by John K. Outsterhout, published by Addison-Wesley Publishing Company, 1994. For examples of Tcl job scripts, review the scripts located in `ORACLE_HOME\network\agent\jobs\oracle` subdirectories on the machine where an Intelligent Agent has been installed. Do not edit these Tcl scripts.

The following is an example of a Tcl script (Unix platform) that logs on to a database and runs a SQL statement:

```
set argc [llength $argv]
if { $argc == 1 } {set argv [lindex $argv 0]}
set connect_str [lindex $argv 0]
set sql_statement [lindex $argv 1]
set lda [oralogon $connect_str]
set curl [oraopen $lda]
```

```

orasql $curl $sql_statement
set result_row [orafetch $curl]
while {$oramsgrc == 0} {
    puts $result_row
    set result_row [orafetch $curl]
}
oraclose $curl
oralogoff $lda

```

When the script is executed with the Run Tcl Script task, the following are examples of command line arguments that should be entered in the Parameters field:

```
"scott/tiger@or817.world" "select * from emp"
```

The following is an example of a Tcl script (Unix platform) that displays the contents of a file if it exists and triggers a third-party event if it does not exist:

```

set argc [llength $argv]
if {$argc == 1} {set argv [lindex $argv 0]}
set myfile [lindex $argv 0]
append mymessage "File not found:" $myfile
if {[file exists $myfile]} {
    catfile $myfile
} else {
    puts $mymessage
    orareporevent /user/host/file/alert $oramsgrc(nodename) 1 $mymessage
}

```

When the script is executed with the Run Tcl Script task, the following is an example of a command line argument that should be entered in the Parameters field:

```
"/export/oracle/network/agent/dbsnmp.ver"
```

Note: When `orareporevent` is used to trigger a third-party event with a job script, you need to create and register an event that has the "Unsolicited events" box checked. See "Event General Page" on page 6-46 for more information.

Listener Tasks

These are the tasks that can be run on Listeners. In addition, you can run operating system or host job tasks.

- Shutdown Listener

- Startup Listener

Shutdown Listener

This stops the Listener. The preferred credentials for the node must have a user that has system administration privileges. See "Preferred Credentials" on page 1-25 for information on user preferences.

Parameters:

Password. Enter a password for the listener if you choose to override the default password.

Startup Listener

This can be invoked to start the Listener. The preferred credentials for the node must have a user that has system administration privileges. See "Preferred Credentials" on page 1-25 for information on user preferences.

Parameters:

None

Important: See "Numeric Pager Job/Event Ids" on page 6-57 for specific ids.

HTTP Servers

These job tasks provide limited control your web servers. Although you can start or stop a managed HTTP server directly from the Console Navigator, you can also use the Job system directly to schedule web server shutdowns or startups.

- Shut Down HTTP Server
- Start Up HTTP Server

No parameters are required for either job task.

Job Tasks Run through Wizards

Some Enterprise Manager wizards and applications use the job system to perform specific operations. Job tasks used by these wizards/applications appear in the job tasks list, but cannot be used directly as with regular job tasks discussed in the previous section. If one of these job tasks is selected, a button appears on the Job

property sheet Parameters page allowing you to start the related application/wizard. The following job tasks are used in conjunction with specific wizard/applications:

- Backup
- Recovery
- Export
- Import
- Load
- Analyze
- Run RMan script

The Event system allows you to monitor your network for specific conditions, such as loss of service or lack of storage, that may occur in your managed environment. You select tests to run on managed targets (databases, nodes, listeners, or other services), then set the threshold parameters for which you want to be notified. You can share events with other administrators, in addition to being able to notify specific administrators when an event condition occurs. For some event tests, you can also choose to execute a *fixit* job that automatically corrects the problem.

The following topics are discussed in this chapter:

- Event System Overview
- Event Detail View
- Event Menu
- Event Viewer
- Event General Page
- Oracle Event Tests
- Event System Features and Requirements

Event System Overview

The Event system allows you to efficiently monitor large enterprise. Using the Event system and Intelligent Agents, you can effectively monitor any number of databases, nodes, or other services 24 hours a day, and be alerted when a problem or specific condition is detected. You can also pinpoint only the services you wish to monitor. The Event system can be extended to include other third-party applications that detect events independent of the Intelligent Agents.

In the Event system, event settings are stored based on the administrator registering the event. This allows administrators of large systems to customize their event systems to their preferences and tasks. Administrators receive messages for events for which they have been selected to receive notifications by other administrators.

The Event system includes the following processes:

1. Creating an event by completing the Event property sheet pages. This involves:
 - a. Determining the monitored targets.
 - b. Selecting the event tests that you want to run.
 - c. Determining the threshold parameters for the event tests.
 - d. Determining how often the event condition is to be checked.
 - e. Specifying a fixit job to be run when an event triggers. (Optional)
 - f. Assigning permissions to allow other administrators to share the event or be notified if the event condition is met.
2. Saving and modifying an event.
3. Registering, or submitting, an event to the Intelligent Agents on the monitored targets.
4. Interpreting and correcting an event occurrence.
 - a. Logging information pertinent to your interpretation of the event to the Event log.
 - b. Assigning the Event to a different administrator if appropriate.

Using Events

You need to create and register events, which are simply a group of event tests that you want to run on your managed systems. Oracle Enterprise Manager includes a

variety of predefined event tests that you can use when creating events. The event tests are grouped by target type, for instance:

- Database
- Listener
- HTTP Server
- Concurrent Manager
- Node

Creating Events

You can create events using the predefined event tests that have been installed with Oracle Enterprise Manager. See "Event Categories and Types" on page 6-8 for more information.

The events are created with information entered in the Event property sheet. You determine parameters such as the target that is monitored, the specific tests to perform, the frequency that the event test is executed, and whether other administrators can share the events and which administrators should be notified if the event condition is met. See "Access" on page 1-23 for more information. Some event tests have parameters with threshold values that you can customize for your system. See "Event Parameters Page" on page 6-49 for more information. To use the Event system, an administrator must have sufficient privileges to access database objects from the Console. Under most circumstances, full DBA privileges are not required, nor would be appropriate to assign full DBA privileges to every administrator. For this reason, the OEM_MONITOR role was created.

Enterprise Manager Monitor Role Beginning with Oracle 8.0.6 databases and higher, the OEM_MONITOR role is created by the Oracle database creation scripts. This role permits access to database functionality within Enterprise Manager. For example, running events against a database (tablespace full, buffer cache hit ratio) or browsing through the objects in a database via the Console Navigator tree. These types of functionality require database credentials on which to perform these operations. Rather than granting the powerful DBA role to the database credentials, many administrators prefer to provide only the necessary privileges required to do these operations. Granting the OEM_MONITOR role to the database credentials, ensures that the user has the minimum sufficient privileges required for these operations.

Note: You need to create the OEM_MONITOR role using the SYS account.

If you need to create the OEM_MONITOR role manually, here are the steps you need to perform:

1. Create a role called OEM_MONITOR

```
drop role OEM_MONITOR;  
create role OEM_MONITOR;
```

2. Grant the 'connect' role to OEM_MONITOR

```
grant connect to OEM_MONITOR;
```

3. Grant the system privileges 'Analyze any' & 'Create table' to OEM_MONITOR

```
grant analyze any to OEM_MONITOR;  
grant create table to OEM_MONITOR;
```

4. Create the SELECT_CATALOG_ROLE role as defined in sc_role.sql .

5. Grant the SELECT_CATALOG_ROLE to the OEM_MONITOR role

```
grant select_catalog_role to OEM_MONITOR;
```

You are now ready to grant the OEM_MONITOR role to the database user that will be used as “database preferred credentials” in Enterprise Manager. In addition to granting the OEM_MONITOR role to a user, you must also ensure that the QUOTA for the user account is set to UNLIMITED.

The “Continued Row” event test needs to analyze results into a table so it needs both the "analyze any" and "create table" privileges.

Note: The "analyze any" privilege is used by the "index rebuild" event to compute statistics.

Registering Events

Events are registered, or submitted, to specific targets, such as nodes, listeners, or databases. The status of a registered event is viewed in the Registered page of the

Event pane. **Note:** The "Show Targets" checkbox is checked at the bottom of the Registered page.

The event scripts are executed on nodes with the permissions of the Intelligent Agent. However, some of the database event tests, such as Continued Rows, require access to system tables and require additional permissions. You need to set up preferred credentials for the monitored database with an administrator that has system privileges. See "Preferred Credentials" on page 1-25 for more information.

The Intelligent Agent is responsible for detecting when a specific event condition has occurred. The Intelligent Agent first notifies a Management Server, which in turn notifies interested administrators either through the Oracle Enterprise Manger Console, or by external means such as Email or Paging.

The Management Server is responsible for registering event information with the appropriate Intelligent Agents on nodes in the network. You determine the frequency that an Intelligent Agent checks an event. See "Event Schedule Page" on page 6-50 for details on setting the frequency interval for an event. An exception to this is the Up/Down (node) event test, which is checked at an interval set by the system itself. See "Fault Management Event Tests" on page 6-9 for more information on this event test.

Event Occurrences

When an alert condition occurs, the Intelligent Agent is responsible for notifying the Management Server. Each event is logged in the repository and can be viewed and acknowledged in the Alerts page of the Console. See Figure 6-6, "Event Menu and Detail View" for an illustration of the Event pane.

Event Notifications

Events can consist of multiple event tests. If any one of these tests identify a specified condition, the event is triggered and a notification is sent to the Console. If enhanced notification is configured for your system, paging and/or email notifications are sent.

Event notification occurs as follows:

- A notification is sent when the threshold of an event test exceeds the level specified by parameter values. If the event does not have parameters, a notification is sent when the event occurs.
- If the event test condition remains above the threshold specified, a new notification is not sent. If the condition does not exist when the next test is run,

the event clears. Notifications are also sent (email/paging) when an event clears.

- If an event test condition changes from warning to critical or critical to warning, a new notification is sent to the Event pane or via E-mail or paging.
- If you acknowledge and move an alert to history, a new notification is not sent to the Alerts page unless a moved warning changes to a critical alert.

Notifying Administrators

Enterprise Manager administrators can be notified in various ways, such as electronic mail or paging, depending on the administrator's setup and permissions. You need to set up the notification services and determine the administrators that need to be notified for the events. See "Event Access Page" on page 6-51 and "Access" on page 1-23 to determine the administrators that receive notifications. See "Notification" on page 1-14 to determine how and when an administrator is notified.

If you plan to notify administrators with email or paging, you need to make sure the following is set up properly:

- The settings for the system modem.
- The notification schedules for the administrators.
- The mail and paging services that are used to contact administrators.
- The mail address and paging numbers for each administrator.

Interpreting Events

An event is composed of one or more event tests. While an individual event test may result in a different status (For example, some clear, some are in alert), there is a general status for the Event. To determine the general severity for the event, the following rules apply in succession:

- a. If the event includes an UpDown event test, and this test triggers, then the general status of the Event is "Unknown" (gray flag)
- b. Otherwise, if the event includes a test that reaches an alert state, then the general status of the Event is "Critical" (red flag)
- c. Otherwise, if the event includes a test that reaches a warning state, then the general status of the Event is "Warning" (yellow flag)
- d. Otherwise, if the event includes a test that is in error, then the general status of the Event is "Error" (yellow hexagon)

- e. Otherwise, all tests should be clear, so the general status of the event is "Clear"

You can still see the individual status of each event test in the Event Viewer.

Event Colors and Icons All events return values and some events produce output messages. The events return different status icons depending on the severity of the event. These severity levels are determined by parameter thresholds you set for the event tests during event creation. The colors are displayed on the event severity icon that is located:

- Next to the event name listed in the Alerts page of the Events detail view.
- On the object in the Group detail view if it is part of a group. See Chapter 4, "Groups" for information.

The colors of the event severity icons are:

- Error State (yellow hexagon with an exclamation point)

An error state indicates there is a problem with the evaluation of the event condition, as opposed to a threshold being met. Examples of error states are: registering an Archive Full event against a database in non-archivelog mode, registering an event that monitors segments but specifying a filter that excludes all available segments.

- Event cleared (green flag)
- Warning (yellow flag)
- Critical (red flag)
- Unknown (gray flag)

A gray flag represents an "unknown" state where it is not possible for Enterprise Manager to ascertain the event status because the node is unreachable or the Intelligent Agent is not available. The gray flag will appear on the group pane and as the flag for the event in the Alert tab if your event includes at least one up/down event test (any target: node/database/listener). When the gray flag occurs, it will be set for the Event. When you see the event in the Event viewer, the flag for the up/down event test will be gray, and the flags for the other event tests will remain the color of their original state.

If you have an event that does not include any up/down event test, then even if the target node or Intelligent Agent become unavailable, the unknown state will not trigger for that event; the current severity for the event will be unchanged.

Note: Some events, such as Probe and User Blocks events, do not return a warning value because the warning threshold parameter is not used. The event has either occurred or not occurred.

Correcting Problems

When an event occurs, you need to correct the problem. In some cases, you can create a fixit job that responds to specific event conditions. See "Event Fixit Jobs Page" on page 6-54 for more information. These situations are noted in the online help for Oracle events.

In other cases, the solution may require the attention of a system administrator. For example, space management event conditions may require an administrator to increase space requirements and resource management conditions may require an administrator to adjust initialization parameters. The online help for Oracle Event Tests has recommendations on how to resolve many of the common event condition.

If the Diagnostics Pack is installed, advice and/or related tools and charts are available to help administrators diagnose the problem. For additional information on Oracle database problems, refer to the Oracle Server Administrator's, Tuning, and Reference Guides. For network problems, refer to the Oracle networking guides for your system.

Event Categories and Types

The Oracle event tests for the database, listener, and node destination types are grouped into categories:

- Fault Management event tests
- Space Management event tests
- Resource Management event tests
- Performance Management event tests
- Unsolicited event tests
- User-Defined Monitoring

Only the UpDown event tests are included with Oracle Enterprise Manager. These fall under the 'Fault' category for the selected target type. Additional advanced

events for all categories are available with the optional Oracle Diagnostics Pack. Beginning with the Oracle Diagnostics Pack for Enterprise Manager version 2.2, operating system-specific tests are also available for NT and various UNIX platforms.

See the online help for Oracle predefined event tests, "Oracle Event Tests" on page 6-56, and the Diagnostic Pack documentation for information on events and their parameters. You can also refer to the *Enterprise Manager Event Test Reference Manual* for a comprehensive look at all available event tests. All the Node events are supported on Unix and Windows NT platforms. For other platforms, see your platform-specific documentation.

Fault Management Event Tests

This category of event tests monitors for catastrophic conditions on the system, such as a database, node, or listener is down. Immediate action must be taken by the administrator. Examples of event tests available in this category include:

- Alert
- UpDown

Most of the fault management event tests do not require any threshold values because the event test only checks whether the service is up or down or if the event condition occurred. For the Alert event test, the event test checks whether error messages are written into the database alert log file.

The UpDown event tests are provided with the Enterprise Manager base product. These event tests check whether a database, listener, or node is available. With the UpDown event test for databases or listeners, you can use the Startup Database or Startup Listener task as a fixit job to re-start the database or listener. To avoid executing that job when the database or listener is brought down intentionally, you need to remove the event registration or blackout the target.

Space Management Event Tests

This category of event tests track possible space problems, such as running out of space on a disk or archive device. Examples of space management event tests in this category include:

- Disk Full
- Archive Full

To check for space management events, set a threshold on the free space left. For example, set an alert if the free space on a disk falls below a specific number of bytes. In order to properly choose the threshold value, you need to know the characteristics of the tablespaces. For example, you would want to know whether the tablespaces contain online transaction processing (OLTP) tables or decision support tables. The former usually has a very fast growth rate, while the latter almost never grows.

Resource Management Event Tests

This category of event tests track possible resource problems, such as exceeding datafile or lock limits. Examples of resource management event tests in this category include:

- Datafile Limit
- Lock Limit

To check for resource management events, set a threshold on the percentage of a resource used. For example, you can set an alert if the percentage of the datafile resource used is greater than a specified value.

Performance Management Event Tests

This category of event test monitors the system for performance problems, such as excessive disk input/output or library cache miss rate. Examples of events in this category include:

- Disk I/O
- Library Cache

To check for performance management events, set a threshold on a system value. For example, you can set an alert if the library cache miss rate is greater than a specific value. The set of threshold values is system specific, depending on the hardware platform, number of users, and other factors.

Unsolicited Event Tests

Unsolicited event tests are events that have been initiated outside the Enterprise Manager Event system. An event is considered unsolicited if it is raised by a process other than the Oracle Intelligent Agent, but is running on the same node as the Intelligent Agent. These events are usually checked and provided by third-party software. Creating an unsolicited event allows you to integrate and monitor

third-party events. Essentially, there are two phases to setting up an unsolicited event:

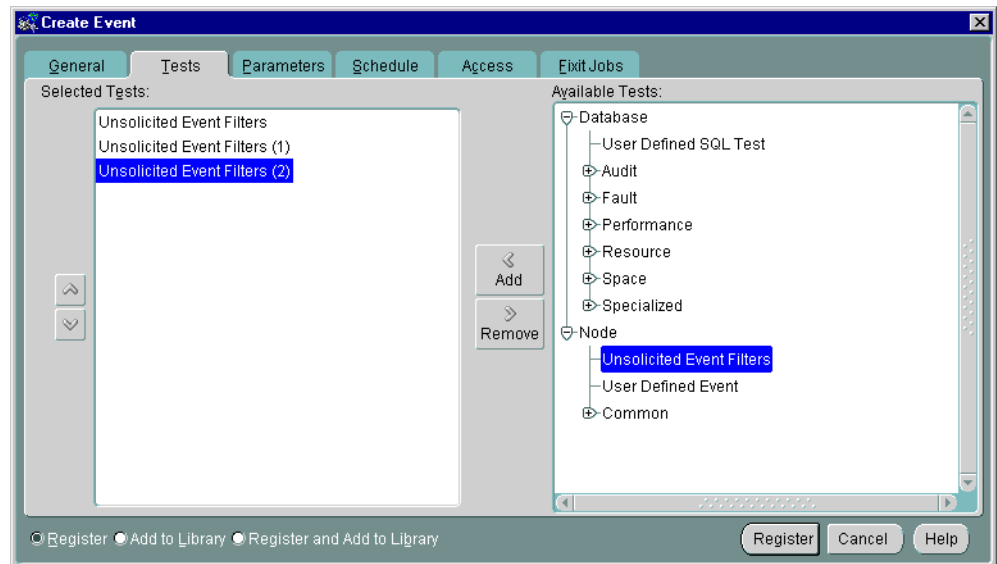
- Registering interest in an unsolicited event
- Raising the unsolicited event

Registering Interest in an Unsolicited Event

In order to receive unsolicited events, you must create and register an event that is expecting to receive unsolicited events. The event should have the event test "Unsolicited Event Filters." This event test also allows you to filter on only those unsolicited events you are interested in.

To register interest in an unsolicited event, choose the Unsolicited Event Filters event test in the Test page and complete the Parameters pages.

Figure 6–1 Test Page: Unsolicited Event



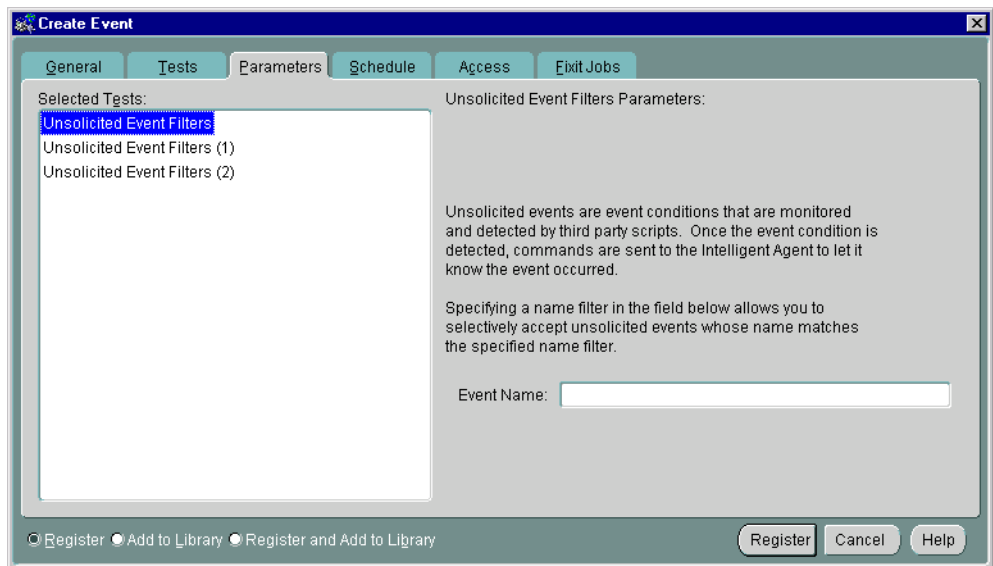
You can have more than one unsolicited event test per event. Information on how to fill out the Parameters page is discussed in the next section. After completing the

unsolicited event, you can save and submit the event. See "Event General Page" on page 6-46 for more information.

Setting the Parameters Property Sheet for Unsolicited Events

Because unsolicited events originate outside the Event system, you may wish to screen only for specific external events. The Parameters page for the Unsolicited Event Test allows you to filter unsolicited events based upon the event name.

Figure 6–2 Parameters Page: Unsolicited Event



Event Name

This is the four-part name of the event of the form:

`/vendor/product/category/name`

You can enter any character string but all four parts and the forward slashes (/) are required. The `eventname` is assumed to be in 7-bit ASCII, so that it never changes regardless of platform or language. The name of the event that fires must match the value specified in this parameter field in order for the unsolicited event to fire.

You can enter a wildcard "*" to specify no filtering. In this case, all unsolicited events that are raised by third-party applications will trigger this unsolicited event.

Raising Unsolicited Events

When third-party applications detect their own external events, they can notify the Intelligent Agent by raising the event. To raise unsolicited events, users have a choice of a command-line interface (`oemevent` executable) or an OraTcl verb (`orareporevent`). The related syntax is as follows:

```
oemevent [event_name] [object_name] [severity] [message]
orareporevent [event_name] [object_name] [severity] [message]
```

where `event_name` is the name of the event that triggered the unsolicited event. `object_name` (valid discovered target name) is the name of the object that the event is monitoring, `severity` is the level of severity for the event, and `message` is the text string to be displayed in the Enterprise Manager Console. For additional details, please refer to the *Intelligent Agent User's Guide*.

Note that the severity is specified as a character string in `oemevent` and as an integer in `orareporevent`. Also, note that the `event_name` must be a four-part string of the form `/a/b/c/d`, where the different elements may be used to organize the event test within a hierarchy of event tests. For example, `/myevents/node/files/filefound` may be an event test you developed. It relates to nodes, more specifically space on nodes, and it monitors for the existence of a particular file.

For information on OraTcl and event scripts, see the *Oracle Intelligent Agent User's Guide*.

Raising Unsolicited Events through the Enterprise Manager Job System

Typically, unsolicited events are evaluated and raised by third-party software. Enterprise Manager allows you to implement monitoring of unsolicited events through the Job system and Tcl. You create a Tcl job and submit it as a periodic job. The Tcl contains logic to evaluate the underlying test and decide whether it needs to raise the event and at what severity level. Since the job is submitted as a periodic job, the underlying test is evaluated periodically like all regular Enterprise Manager event tests. Techniques such as those in the following examples, allow users of Enterprise Manager to implement and customize event monitoring specific to their environments.

Example 1: Simple Job Raising an Unsolicited Event It is possible to submit a job with an imbedded OS command task that executes the `oemevent` program and passes the program all necessary arguments. All users that have registered for the unsolicited event raised by `oemevent` will receive the event notification. The event has to be known to the administrator submitting the job that raises it.

The job that raises the event may contain enough logic to evaluate the underlying test and decide whether it needs to raise the event and at what severity level. Such a job may be submitted as a periodic job so that the underlying test is evaluated periodically, which is similar to regular Enterprise Manager event tests.

Unsolicited events are evaluated in their own process and within the proper OS security protocols and do not pose security or robustness threats to the system. Introduced here is a procedure where the user must submit a job in order to monitor for an external event.

The following example illustrates how to implement an event test that triggers when a particular file is found. Let's call this event
`/myevents/node/files/filefound.`

The following Tcl script needs to be submitted as the job:

```
# event name
set event_name /myevents/node/files/filefound
# filename to look for comes at the first (and only) argument
set file_name [lindex $argv 0]
# check for the file, and if it's found trigger the event as critical
if { [file exists $file_name] } {
    orareporevent $event_name $oramsg(oraobject) 2 "$file_name found"
}
```

In order to receive this event, a user needs to register an event with the Unsolicited Event Filters test selected and configured to filter an event name of the format:

`/myevents/node/files/filefound.`

This event should be registered against a node and will trigger against it. The message associated with the event occurrence will contain the values of all parameters passed into `orareporevent`.

Although this event is fairly straightforward, there are two problems:

1. The event never clears – The event should eventually clear after the file disappears.
2. The event will trigger every time the above script is evaluated, even if it has triggered previously – You will receive multiple copies of the same alert. An event should not trigger unless there is a severity change.

Any other scripting language or executable program can also be used to implement the logic of an unsolicited event test. However, Tcl is preferred because it allows platform-independent implementation and the fact that the code may be sent from

the Enterprise Manager Console 'on-demand' without requiring anything to be installed on the Intelligent Agent side.

Example 2: Unsolicited Events with the Proper Lifecycle As with regular Enterprise Manager events, unsolicited events could be triggered only once per condition detected and could clear automatically if the condition that triggers the event is no longer met. Events adhering to this operational pattern are said to have a "proper lifecycle."

Typically, scripts that implement unsolicited events are composed of two basic parts:

1. The part that evaluates the event and sets its associated severity
2. The part that handles the event reporting and avoids multiple notifications if the severity does not change

The following Tcl script illustrates this two-part script implementation, as well as a technique that allows proper event lifecycle.

```
#-----
#
# Tcl Procedure
#   orareporevent1
#
# Purpose:
#   Trigger an unsolicited event only previous state is different
#
# Arguments:
#   - event_name: event test to trigger
#   - severity: new severity
#   - message: message to be attached to the event report
#
#-----
proc orareporevent1 {event_name severity message} {
    # define a 'lock' that its contents define the previous event status
    # and figure out the event state during the previous execution
    global orams
    append event_lock [tempdir] "/" $orams(jobid) ".el"
    if { [file exists $event_lock] } {
        set f [open $event_lock r]
        gets $f previous_severity
        close $f
    } else {
        set previous_severity -1
    }
}
```

```
# if event test state has changed, trigger the event at new severity
if { $previous_severity != $severity } {
    orareporevent $event_name $oramsg(oraobject) $severity $message
    if { $severity == -1 } {
        mfile $event_lock
    } else {
        set f [open $event_lock w]
        puts $f $severity
        close $f
    }
}
}

#-----
#
# Event Test Name:
#     /myevents/node/files/filefound
#
# Purpose:
#     Monitor for the existence of a particular file
#     The test triggers at warning level if the file exists, but
#     at critical level if the file is larger than the specified
#     value
#
# Arguments:
#     - filename to look for
#     - critical file size
#
#-----
set event_name /myevents/node/files/filefound
set file_name [lindex $argv 0]
set critical_filesize [lindex $argv 1]

if { [file exists $file_name] } {
    # if the file exists calculate its size in Kilobytes
    set file_size [expr [file size $file_name] / 1024]
    if { $file_size > $critical_filesize } {
        # if file is larger than the critical value, trigger as critical
        orareporevent1 $event_name 2 "Size: $file_size Kb"
    } else {
        # if file is smaller than the critical value, trigger as warning
        orareporevent1 $event_name 1 "Filesize: $file_size Kb"
    }
} else {
    # if file is no longer there, clear the event
}
```

```

    orareporevent1 $event_name -1 "File does not exist"
}

```

Example 3: An Unsolicited Event Script that Accesses the Oracle Database This example of an unsolicited event test illustrates a situation where the test evaluation involves connecting to an Oracle instance and executing some SQL against it.

This example checks the size of a particular table in the database and triggers the event when a set threshold is crossed. There is a warning value and a critical value. The size of the table is measured by counting the number of its rows.

```

#-----
#
# Tcl Procedure
#     orareporevent1
#
# Purpose:
#     Trigger an unsolicited event only previous state is different
#
# Arguments:
#     - event_name: event test to trigger
#     - severity: new severity
#     - message: message to be attached to the event report
#
#-----
proc orareporevent1 {event_name severity message} {
    # define a 'lock' that its contents define the previous event status
    # and figure out the event state during the previous execution
    global orams
    append event_lock [tempdir] "/" $orams(jobid) ".el"
    if { [file exists $event_lock] } {
        set f [open $event_lock r]
        gets $f previous_severity
        close $f
    } else {
        set previous_severity -1
    }
    # if event test state has changed, trigger the event at the new severity
    if { $previous_severity != $severity } {
        orareporevent $event_name $orams(oraobject) $severity $message
        if { $severity == -1 } {
            rmfile $event_lock
        } else {
            set f [open $event_lock w]

```

```

    puts $f $severity
    close $f
  }
}

#-----
#
# Event Test Name:
#       /myevents/database/space/tablesize
#
# Purpose:
#       Monitor the size of a particular database table
#       The test triggers at warning level when the warning threshold
#       is crossed and at critical level when the critical threshold
#       is crossed
#
# Arguments:
#       - table name
#       - critical threshold
#       - warning threshold
#       - username/password for connecting to target (optional)
#
#-----
set event_name /myevents/database/space/tablesize
set table_name [lindex $argv 0]
set critical_threshold [lindex $argv 1]
set warning_threshold [lindex $argv 2]

if { $argc == 4 } {
    set connect [format "%s@s" [lindex $argv 3] $oramsg(oraobject)]
} else {
    set connect [format "%s/%s@s" $SMP_USER $SMP_PASSWORD $oramsg(oraobject)]
}

if {[catch {oralogon $connect} lda]} {
    append msg "Cannot connect to target." "\n" $oramsg(errortxt)
    orafail $msg
}

if {[catch {oraopen $lda} cur]} {
    append msg "Cannot connect to target." "\n" $oramsg(errortxt)
    oralogoff $lda
    orafail $msg
}

```



```

set sql [format "select count(*) from %s" $table_name]
if {[catch {orasql $cur $sql}]} {
    append msg "Cannot execute SQL against the target." "\n" $oramsg(errortxt)
    oraclose $cur
    oralogoff $lda
    orafail $msg
}

if {[catch {orafetch $cur} row]} {
    append msg "Cannot execute SQL against the target." "\n" $oramsg(errortxt)
    oraclose $cur
    oralogoff $lda
    orafail $msg
}

set current_tablesize [lindex $row 0]

if { $current_tablesize > $critical_threshold } {
    orareporevent1 $event_name 2 "Table:$table_name #rows:$current_tablesize"
} elseif { $current_tablesize > $warning_threshold } {
    orareporevent1 $event_name 1 "Table:$table_name #rows:$current_tablesize"
} else {
    orareporevent1 $event_name -1 "Table:$table_name #rows:$current_tablesize"
}

```

A number of OraTcl verbs were used in this script. Refer to the *Intelligent Agent User's Guide* for details on OraTcl verbs. Note that the preferred credentials, specified in the Console, are available to the script writer via the SMP_USER and SMP_PASSWORD Tcl global variables. For jobs against a database, the values of those variables are set to the username and password specified as preferred credentials for that database. This script also allows for an optional overwrite of the preferred credentials via an optional forth input argument.

Unsolicited Event Caveats

- When raising unsolicited events against a particular target, then an unsolicited event test must be registered against that event particular target from within the Console. When raising the unsolicited event via `omevent` or `orareporevent`, the specified object name must match the target name as displayed in the unsolicited event for that target.
- Unsolicited Events cannot be used with Fixit Jobs unless you are using a 9i Intelligent Agent.

User-Defined Monitoring

The Event system provides you with two types of user-defined monitoring capability:

- User-Defined SQL Event Test
- User-Defined Event Test

The User-Defined SQL Event test is used for database-specific events. It allows you to define your own custom database events by specifying your own SQL query that will evaluate the event condition. The return value of the SQL query will be compared against thresholds you specify.

The User-Defined Event test can be used to monitor any type of event condition. It allows you to specify your own monitoring scripts that will be used to monitor the event condition. These scripts can be written in any scripting language suited to your environment.

User-Defined SQL Event Test

The User-Defined SQL event test allows you to define your own SQL script that evaluates an event test. The event tests you define should be written as queries, such as SELECT statements, that return condition values for which you are monitoring. These values are checked against the Critical threshold and Warning threshold limits you specify, and trigger the event if the threshold limits are reached.

Example 6–1 Creating a User-Defined SQL Event Test

You have a custom application that runs against the Oracle database. Each time it finds an application error, it creates an entry into a table called "error_log". Using the "User-Defined SQL Test", you can write an event test that notifies you when it finds at least 50 errors. Specifically, you define the following SQL statement:

```
select count(*) from error_log
```

This returns the number of rows in the error_log table. Since you want a critical alert raised when it reaches at least 50, you specify the Operator ">=", a Critical Threshold value of 50, and a Warning Threshold value of 30.

Support for PL/SQL Functions

If your query for the event condition requires more complex processing than is allowed in a single SELECT statement, you can first create a pl/sql function that

contains the extra processing steps, and then use the pl/sql function with the User-Defined SQL event test. Your pl/sql function must still return a value that can be compared against the Critical and Warning thresholds.

Example 6–2 Using a SQL Event Test

You need to trigger a critical alert whenever an employee's salary is \$500 higher than the highest manager's salary. You first define a pl/sql function as follows:

```
create or replace function overpaid_emp return number is
max_mgr_sal number;
max_emp_sal number;
begin
select max(sal) into max_mgr_sal from scott.emp where job = 'MANAGER' or job =
'PRESIDENT';
select max(sal) into max_emp_sal from scott.emp where job != 'MANAGER' and job
!= 'PRESIDENT';
return (max_emp_sal - max_mgr_sal);
end;
```

This pl/sql function returns the difference between the highest employee's salary and the highest manager's salary. If the difference is a positive number, then an employee has the higher pay. If the difference is more than 500, then a critical alert needs to be triggered.

When defining this event this using the User-Defined SQL event test, you define the SQL statement as follows:

```
select overpaid_emp from dual
```

Then use the Operator ">" and Warning threshold of 100 and Critical threshold of 500.

Note that ROLES are not enabled within PL/SQL functions, so any privileges that are granted via ROLES will not work from within the function. You may need to grant the privileges directly to the database user account that is used for the event. (The database user account used for the event is either the Preferred Credentials user for the database, or the overwritten preferred credentials).

Parameters

- **SQL:** Type the SQL query you want to use. You can also cut and paste SQL from an existing script or import from an existing file containing the SQL query.
- **Operator:** Select one of the following comparison operators: == (equal); < (less than); > (greater than); <= (less than or equal to); >= (greater than or equal to);

!= (not equal). This operator will be used to compare the SQL query's return value to the critical and warning thresholds you specify.

- **Critical Threshold:** Depending on the return value (number or text string) of your SQL query, enter a scalar value that will be compared against the query's return value. If the return value crosses this threshold, an Event at Critical severity will be generated.
- **Warning Threshold:** Depending on the return value (number or text string) of your SQL query, enter a scalar value that will be compared against the query's return value. If the return value crosses this threshold, an Event at Warning severity will be generated.
- **Occurrences Preceding Notifications:** Type a numeric value indicating how many times the SQL query's return value crosses the thresholds before an alert flag is displayed in the Console and before a notification is sent.
- **Override Preferred Credential:** Check this box if you want to change the user name or password or both.

User-Defined Event Tests

The User-Defined event test (available with the Oracle Diagnostics Pack) allows you to define your own scripts that monitor conditions particular to your environment. These event tests can be written in any scripting language, as long as the node that runs the script has the appropriate runtime requirements to execute the script.

The power and flexibility of User-Defined event tests lie in the ability to integrate any, custom script into the Enterprise Manager Event System and leverage the system's multi-administrator, lights-out scheduling and notification capabilities.

User-Defined events are implemented in two phases:

1. Creating your monitoring script.
2. Registering the script as a User-Defined event in the Enterprise Manager Console

Creating Your Monitoring Script

Using a scripting language of your choice, create a script that contains logic to check for the condition being monitored. Examples of these are scripts that check for disk or memory usage. All monitoring scripts should contain these basic elements:

- Code to check the status of monitored objects
- Code to evaluate the results

- Code to send the results back to the Enterprise Manager Event system

Code to check status of monitored object

Define logic in the code that checks the condition being monitored. For example, the amount of free space on a particular filesystem, memory usage, etc.

Code to evaluate the results

After checking the monitored condition, the script should return either the value associated with the monitored object OR the event status severity.

If you choose to have the script return the value of the monitored object (e.g. actual disk usage), then it means you want the Enterprise Manager Event system to evaluate the object's current value against Warning and Critical thresholds you specify. You specify these warning and critical thresholds when you register the event.

Otherwise, if you choose to have the script itself evaluate the event status severity of the monitored object, you need perform this evaluation in such a way that it falls under one of following event status severities:

Table 6–1 Severity Levels for Event Status

Severity Level	Status
<i>Script Failure</i>	The script failed to run properly. This status is represented by numeric value -2.
<i>Clear</i>	No problems with the object monitored, hence status is clear. This status is represented by numeric value -1.
<i>Warning</i>	The value of the monitored object reached the warning threshold. This status is represented by numeric value 1.
<i>Critical</i>	The value of the monitored object reached the critical threshold. This status is represented by numeric value 2.

Code to return the results back to the Event System

After evaluating the status of the monitored object, the script needs to return this result back to the Event System. The script should return the result by sending tagged information to standard output (stdout) using the syntax that is consistent with the scripting language. The result information to be sent should be enclosed by a pair of well-known tags. The following are the tags that are recognized by the Event System as it checks the information in stdout:

<oraresult> and </oraresult>

Enclose within these tags the current value of the monitored object OR the event status severity.

Example:

```
print "<oraresult>200</oraresult>"
```

Returns 200 as the value of the monitored object

```
print "<oraresult>2</oraresult>"
```

Returns an event status of 2 (event in critical state)

<oramessage> and </oramessage>

Enclose in these tags the message to be sent with the event notification if the event triggers.

Example:

```
print "<oramessage>Disk usage is high</oramessage>"
```

<orafailure> and </orafailure>

Enclose in these tags the message to be sent if a failure occurs in the script. The occurrence of an <orafailure> in the standard output is equivalent to sending a tagged <oramessage> and an <oraresult> set to -2 (script failure).

Finally, the script itself will need to be entered in the Create Event property sheet when creating the event or be located in the monitored node. The node needs to be monitored by a 9i or higher Intelligent Agent. Make sure the node has the script's runtime requirements (e.g. perl interpreter) and that the script works independently of the event system.

Special conditions

- There must be exactly one `oraresult` in `stdout`. A missing `oraresult` results in script failure. Having more than one `oraresult` in the output also results in script failure.
- There can be any number of `oramessage` tags - the entire set is sent to the Management Server
- If a notification needs to be generated, and there is no `oramessage` tag, a default message is generated. The message is of this format:

Current result: <value in `oraresult`>

When the user-defined event is evaluated, it executes the script using the Node credentials associated with the event. These are either the default Node credentials associated with the Enterprise Manager administrator who registered the event or the overwritten Node credentials specified when the event was registered, as explained in the next section. Note that any environment associated with these Node credentials will not be available when the script is run.

Register the user-defined event in the Console

Once you have created the monitoring script, you are ready to add the script's monitoring functionality to the Enterprise Manager Event system. To create and register a user-defined event for your monitoring script:

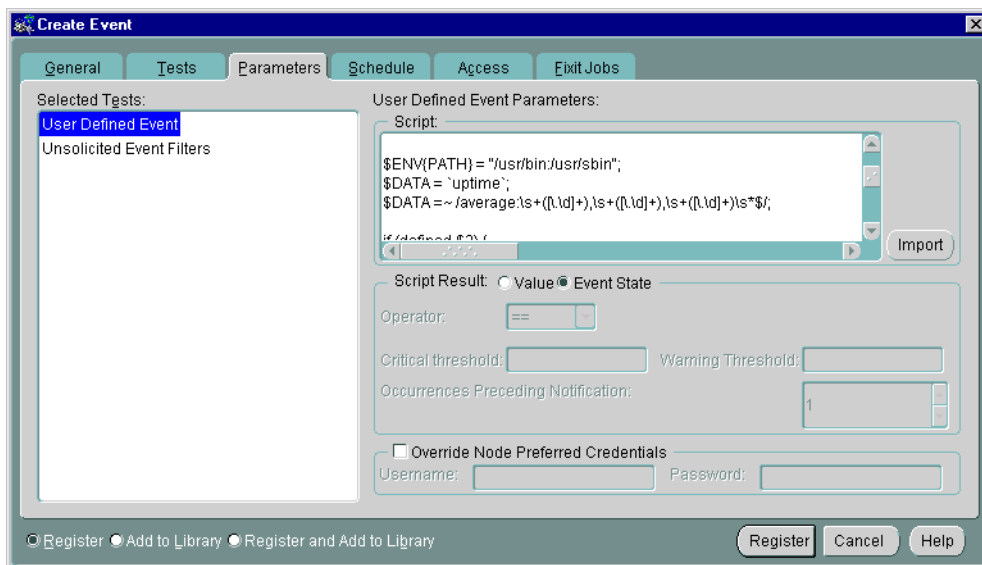
1. Choose the Create Event option from the Event menu to display the Event property sheet.
2. Complete the General page. Select the target node as the node on which the monitoring script will be run.
3. Click on the Tests tab to display available tests for the selected target.
4. If not expanded already, expand the Node object in the Available Tests tree list.
5. Select User Defined Event and click the Add button.
6. Click on the Parameters tab to display the Parameters page.
7. Click on User Defined Event in the Selected Tests list to display the user-defined event test parameters. See Figure 6-3, "User-Defined Event Parameters Page".
8. Click Import. The Load File dialog appears.
9. Select the desired script file and click Open. The contents of the script appear in the Script Text window. See Figure 6-3, "User-Defined Event Parameters Page" and "User-Defined Event Parameters" on page 6-26 for more information on other ways to specify user-defined event parameters.
10. Complete the rest of the Event property sheet.
11. Submit the Event. You have three options when submitting the Event:
 - a. Choose Submit, to register the event against the selected destinations. The new event is not saved to the Event Library.
 - b. Choose Add to the Library (or Save to Library if editing an event from the Library) to save the event to the Event Library. The event will not be submitted to the target destinations at this time. The new event appears in the Event Library dialog.

- c. Choose Submit and Add to Library (or Submit and Save to Library) to submit the event to the selected targets and save the event to the Event Library. The new event appears in the Event Library dialog.
12. Make sure the Event is registered by selecting Events from the Console Navigator and clicking the Registered tab (Detail view). Make sure "Show Targets" is checked at the bottom of the Registered page in order to see the registration status for each target.

User-Defined Event Parameters

The User-Defined Event parameters page allows you to specify the user-defined event test information required to successfully register the event in the Enterprise Manager Console.

Figure 6–3 *User-Defined Event Parameters Page*



User-Defined event test parameters consist of the following:

Script

Enter the monitoring script used for the event evaluation. You can specify this either by entering the full text of the script OR by entering the fully-qualified script name (on the monitored target).

If you choose to enter the full text of the script, and if the script is in a file locally accessible to the console, you can use the "Import" button to load the script from the file instead of manually entering the script.

If your script file resides on the monitored target, you can just specify the fully-qualified filename of the script instead of loading the script text.

Script Result

This parameter indicates the way the results of the event evaluations are returned by your script. You can specify one of two ways in which results are returned: by Value or Event State.

Value: Your script evaluates the condition and returns the value of the monitored metric. Enterprise Manager will then compare the value against specified thresholds.

The following parameters indicate how you want Enterprise Manager to evaluate the value of the monitored metric.

Operator: The operator that Enterprise Manager should use when comparing the value of the monitored metric against the specified thresholds. Select one of the following comparison operators:

- == (equal)
- < (less than)
- > (greater than)
- <= (less than or equal to)
- >= (greater than or equal to)
- != (not equal)

Critical Threshold: The value against which the monitored metric is compared using the specified operator. If it holds true, the event triggers at a Critical level.

Warning Threshold: The value against which the monitored metric is compared using the specified operator. If it holds true, the event triggers at an Warning level.

Occurrences Preceding Notifications: The number of times the event condition should hold true before a notification is sent.

Example:

You may want to create an event that monitors disk space. You can write a script that checks the amount of free disk space and returns that amount as the value to be evaluated. You may want the event to trigger at Warning level when the free disk space is below 500K, and to trigger at Critical level when the free disk space is below 200K. Hence, when defining the event, you would specify the following:

Script: Enter the script text or click Import to load an existing file. If the name of your script is "checkspace.sh" and if it is located on the monitored node, you can, for example, simply enter: `/ul/private/checkspace.sh`.

Event Parameters Page Settings

- Script Result: Choose the "Value" option
- Operator: <
- Critical Threshold: 200000
- Warning Threshold: 500000
- Occurrences Preceding Notification: 1

Event State If you choose this option, the script you write evaluates the event condition and also determines if the event has triggered at a Critical or Warning level, or has not triggered at all (e.g. the event status is Clear or the script has failed to run due to some error). In order to provide the appropriate event status to Enterprise Manager, the script should define and return the appropriate event status. For more information, see "Creating Your Monitoring Script" on page 6-22.

Override Node Preferred Credentials: When your script is executed, it runs as the operating system user specified by the Node credentials associated with the event. These credentials are either the default Node credentials of the Enterprise Manager administrator who is registering this event, or the credentials specified here. It is important to note, however, that any environment associated with the Node credentials will not be used when the script is run.

Output

If the event triggers, the value of the monitored metric is returned. The actual message to be displayed depends on the message you defined in your script via the `<oramessage>` tags. If no message is specified, the default message is: Current

result: <value of monitored metric>. If a failure occurs, then the message displayed is the message specified in the <orafailure> tag.

Bundled User-Defined Event Sample

Enterprise Manager has bundled a sample user-defined event script that 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 node.

This script, called udeload.pl, is installed in the \$ORACLE_HOME/sysman/admin directory where \$ORACLE_HOME is the Oracle directory where the Enterprise Manager is installed.

Full text of the script:

```
#!/usr/local/bin/perl

# Description: 5-min load average.
# Sample User Defined Event monitoring script.

$ENV{PATH} = "/usr/bin:/usr/sbin";

$DATA = `uptime`;
$DATA =~ /average:\s+([\.\d]+),\s+([\.\d]+),\s+([\.\d]+)\s*$/;

if (defined $3) {
    print "<oraresult>$2</oraresult>\n";
} else {
    print "<orafailure>Error collecting data</orafailure>\n";
}
```

Setting Up the Sample Script as a User-Defined Event

1. Copy the script (udeload.pl) to the monitored target. For example: /private/myhome. Make sure you have a 9i version of the Intelligent 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: `udeload.pl` You may need to set its file permissions such that it runs successfully. You should see output of this form:

```
<oraresult>2.1</oraresult>
```

4. In the Enterprise Manager Console, create a new event as follows:
 - a. In the General page, provide a name for the event, say "Test UDE". Choose "Node" as the Target Type. For targets, select the node on which you copied the script.
 - b. In the Tests page, select the "User Defined Event" test.
 - c. In the Parameters page, enter the following:

Script: `/private/myhome/udeload.pl` (... or the fully qualified path to where the script is)

Script Result: make sure the "Value" option is selected

Operator: `>=`

Critical threshold: 0.005

Warning threshold: 0.001

Occurrences Preceding Notification: 1

Override Node Credentials: Specify the credentials of an OS user that can execute the script.

In this example, we want the event to trigger at a Warning level if the 5-minute load average on the machine reaches 0.005, and trigger at a Critical level if the 5-minute load average reaches 0.001. Feel free to change these thresholds depending on your system.
 - d. In the Schedule page, set the time interval upon which you'd like this event to be evaluated. By default this is set to every 5 minutes. As a test, you can reduce this to 1 minute.
 - e. In the Access page, select the Administrators to be notified when the event triggers
5. Click Register to register the event with the Enterprise Manager Event system.

When the 5-minute load reaches at least 0.001, you should see the event trigger in the Enterprise Manager Console as well as have the selected administrators be notified of this event.

Creating and Registering an Event

Events include the target type and the event information that you want to monitor. Events can consist of multiple event tests. To create and register an event:

1. Choose the Create Event option from the Event menu to display the Event property sheet. (you can also display the Event property sheet by opening an event from the Event Library dialog.)
2. Complete the fields in the General page. On the Tests page, select the desired event tests. Complete the rest of the pages of the property sheet to create a new event.
3. When you have completed the Event property sheet:
 - a. Choose Register, to register the event against the selected destinations. The new event is not saved to the Event Library.
 - b. Choose Add to the Library (or Save to Library if editing an event from the Library) to save the event to the Event Library. The event will not be submitted to the target destinations at this time. The new event appears in the Event Library dialog.
 - c. Choose Register and Add to Library (or Register and Save to Library) to submit the event to the selected destinations and save the event to the Event Library. The new event appears in the Event Library dialog.

If you registered an event, the Intelligent Agent on the target node processes the event and the event appears in the Registered page of the Event pane. If the "show Targets" checkbox is selected, each destination target is listed separately with the event. If the "Show Targets" box is not checked, only the target name, type, and owner is shown.

Note: There is usually a slight delay between the time the event is registered and the actual notification by the Intelligent Agent.

When threshold values are exceeded for the tests in an event, the event appears in the Alerts page of the Event pane. The notification changes the color of the severity flag for the event in the Alerts page. If a destination icon is displayed in the Group pane, the flag on the icon changes color. The colors and their meaning are:

- Unknown (gray flag)

- Event cleared (green flag)
- Warning (yellow flag)
- Critical (red flag)
- Error (yellow hexagon)

Cases where an event notification is Unknown (gray flag) indicate the Intelligent Agent or node where the event is registered is unavailable or inaccessible, or the Intelligent Agent on that node is unavailable.

Warning: Do not register an UpDown event (included in the Oracle DB Fault event) against the database or node where the Repository schema is stored. If the database containing the Repository goes down, the Management Server also shuts down. Hence, the Intelligent Agent cannot inform the Management Server that the database is down.

The property sheet for creating a new event is the same as the property sheet for modifying an event, except that the event name and target type fields are always read-only. See Figure 6-8, "Event General Page" for an illustration of the Event property sheet.

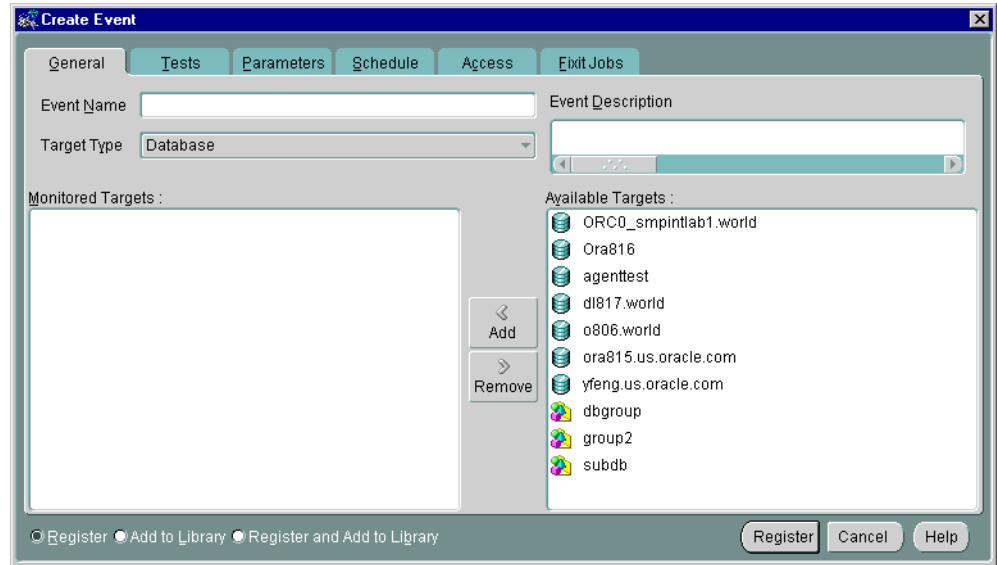
See "Event Categories and Types" on page 6-8 for more information.

Dynamic Modification of Registered Events

Dynamic event modification allows you to actively modify a registered event and have the changes automatically applied to all monitored targets of that event. For example, you can add an additional database to be monitored if you have an existing Tablespace Full event. The Intelligent Agent for the newly added database will now monitor for tablespace full conditions.

However, not all event attributes can be changed. What you are allowed to change depends on the version of the Intelligent Agent used with each monitored target. You may have older versions of the Intelligent Agent running on different targets within your enterprise and these older versions of the Intelligent Agent will only support a subset of modifications you can make using a 9i Agent.

Because pre-9i Intelligent Agents do not support dynamic event modification, if an event contains targets running pre-9i Intelligent Agents, modification will be limited. If all targets running pre-9i Agents are removed from the Monitored Targets list, then full modification of the registered event will be enabled.

Figure 6–4 Event General Page

The following are general usage guidelines for dynamic modification of events:

1. Only the Owner of the registered event can modify all parameters for the event. The owner of a registered event is the administrator who originally registered the event. The owner is shown under the "Owner" column of the Registered page of the Events pane, or via the Owner field in the Access property page of the Edit Event property sheet.
2. An event can be registered against multiple targets on different nodes, each of which is monitored by its own Intelligent Agent. The version of the Intelligent Agent determines the amount of event editing that can be performed against that particular target. You can easily determine the version of the Intelligent Agent running a particular target from the target's Node property sheet.

Figure 6-5 Node Property Sheet

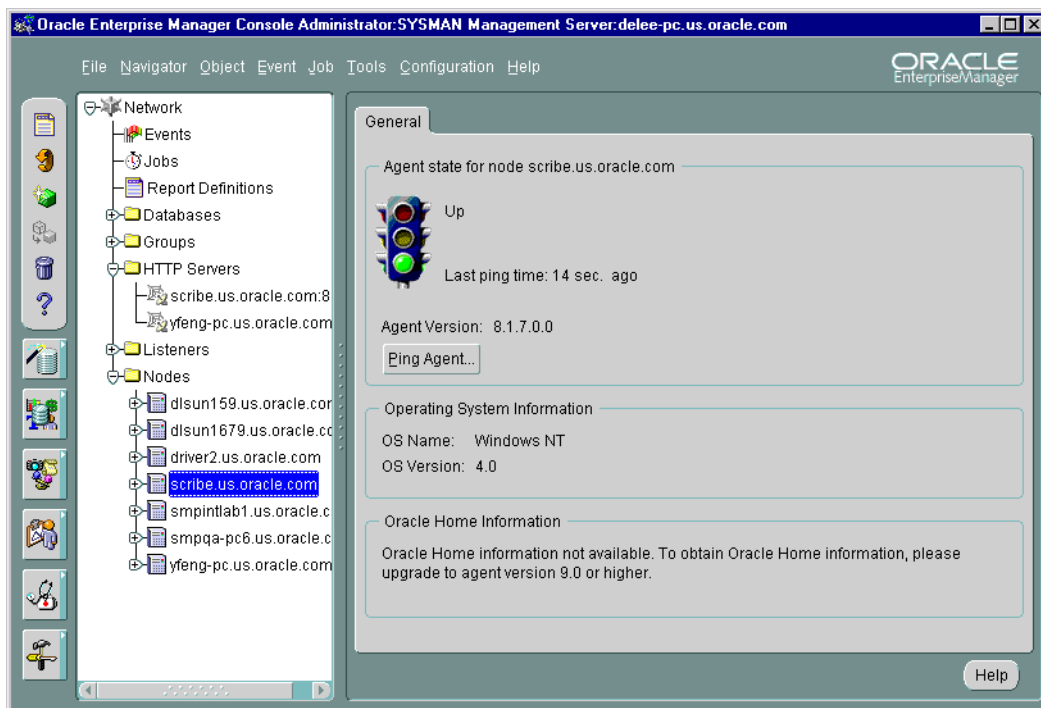


Table 6–2 Modifiable Event Attributes

Event Attribute	Located in this Property Page of the Event Dialog:	If all targets for the event are running with a PRE-9i Agent, can it be modified?	If each target for the event is running with a 9i Agent, can it be modified?	If, for the event, some targets use 9i Agents and some targets use pre-9i Agents, can the event be modified?
Event Description	General	yes	yes	yes
Monitored Targets	General	yes	yes	yes
		Adding a target creates a new event registration for that target.	Adding a target creates a new event registration for that target.	
		Deleting a target de-registers the event for that target.	Deleting a target de-registers the event for that target.	
Adding or deleting event tests	Tests	no	yes	no
Changing test parameters	Parameters	no	yes	no
Schedule - polling frequency and start time	Schedules	no	yes	no
Permissions	Access	yes	yes	yes
Enabling/Disabling SNMP traps	Access	no	yes	no
Selecting or Creating a fixit job for the event	Fixit Jobs	no	yes	no

General Behavior

When dynamically modifying events, there are general system behaviors of which you should be aware:

- When an event has selected targets with the event registered or pending running both pre-9i and 9i Intelligent Agents, then only the attributes that can be modified across all Intelligent Agent versions are supported.

- If you want to modify an attribute that cannot be modified using a pre-9i Intelligent Agent (e.g. test parameters), then you can:
 1. Modify the event by removing the targets running the pre-9i Intelligent Agent. This will enable editing of 9i targets.
 2. Modify the event attributes as necessary, i.e., test parameters.
 3. Submit the changes. The event will be modified on the 9i targets and deregistered from the pre-9i targets.
 4. Modify the event again by adding back the pre-9i target you removed. It will now contain the attribute that was just modified.
- The Event property sheet will automatically enable or disable property pages depending on the changes supported.

For example, if your event originally had a mix of targets running pre-9i and 9i Intelligent Agents, then the Tests, Parameters, Schedule, and Fixit Jobs property pages will be disabled for editing. If during the edit session, you remove the targets running the pre-9i Intelligent Agents, then the Tests, Parameters, Schedule, and Fixit Jobs property pages will now allow editing since the remaining targets support editing of those attributes. However, if any of those attributes are changed in this edit session, you will not be able to bring back the original pre-9i targets you removed. (To bring back the original pre-9i targets, you first have to submit the changes, then re-edit the event to add back the pre-9i targets).

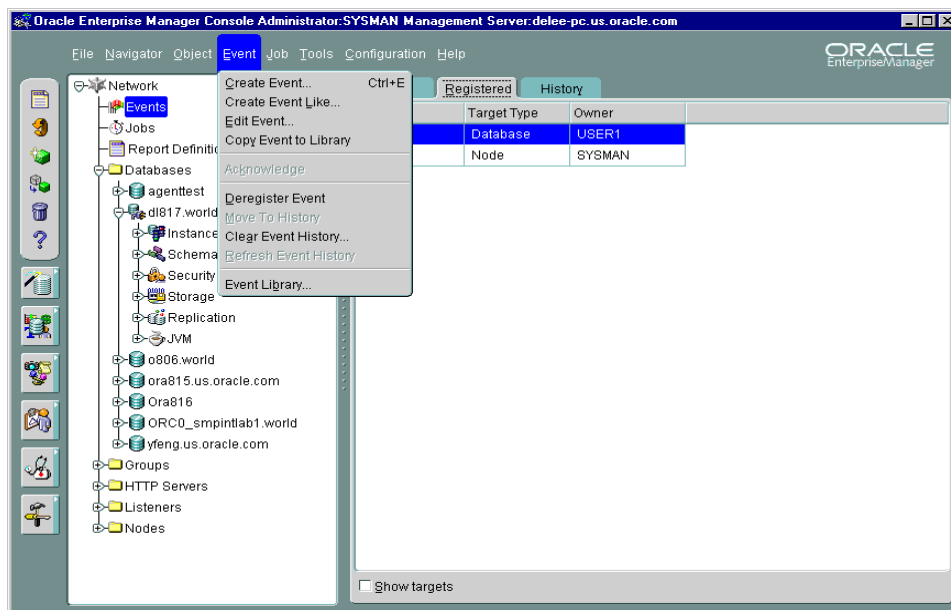
Event Detail View

The Event detail view, which displays when you select the Events object in the Console Navigator, contains the following pages:

- Alerts Page
- History Page
- Registered Page

You can switch between the pages by clicking the tab of each page. The rows in any page can be sorted on any column by clicking the column heading. See Figure 6-6, "Event Menu and Detail View" for an illustration of the Event detail view.

Figure 6–6 Event Menu and Detail View



Because the detail view changes in relation to the object selected in the Console Navigator, Enterprise Manager allows you to "undock" the Event detail view from the Console so that it can be persistently displayed. This allows you to keep an eye on the status of monitored events while you perform other tasks from the Console.

To launch the Event detail view in a floating window, select the Events object from the Navigator and choose 'Display in New Window' from the context-sensitive menu.

Alerts Page

The Alerts page displays event tests that have been triggered.

Severity

Severity of the event occurrence: critical (red flag), warning (yellow flag), clear (green flag), unknown (gray flag), or error state (yellow hexagon).

Name

Name of the event.

Target

Target where the event was triggered.

Target Type

Database, listener, node, or HTTP Server.

Date/Time

Time and date of the event occurrence.

Assigned To

Administrators assigned to work on the event occurrence.

Owner

Administrator who owns the event.

Viewing Alerts

To view details of an event that has occurred, double-click on the event in the Alerts or History page to display the Event Viewer property sheet. See "Event Viewer" on page 6-42 for more information. You can enter notes on the nature and progress of the event condition.

Note: Comments entered into the log are viewable/editable by admins with the Modify permission. After you have reviewed an event, you can move it to the History page. See "Event Viewer" on page 6-42 for more information.

History Page

The Event History page displays a history of events that have occurred and have been moved to History by an administrator or cleared by an Intelligent Agent. The Event History page displays the same columns as the Alerts page.

The History page is refreshed automatically each time you move between the History page and the Alerts or Registered page. However, to refresh the event history list while currently viewing the History pane, you must click the Refresh icon located in the Console toolbar.

To clear all entries in the History page, choose Clear Event History from the Console's Event menu. You can delete entries individually by right-clicking on a

specific event in the History page and choosing Delete Item(s) from the context-sensitive menu.

Registered Page

The Registered page displays the events that have been registered, or submitted, to monitor test conditions on network objects. The Registered page contains the following information:

Name

Name of the event.

Target

Target where the event is monitored. Displayed only when Show Targets is checked.

Target Type

Type of event destination: database, node, listener, web server, Concurrent Manager,

Status

Current registration status of the event: Registered, Registration Pending, De-Registration Pending, Modification Pending, and Registration Failed. Displayed only when Show Targets is checked. The registered event status is only updated when this page is refreshed.

Owner

Administrator who owns the event. Displayed only when Show Targets is checked.

Show Targets

When checked, the Registered page displays Target and Status information. By default, "Show Targets" is not checked.

Under certain circumstances, an event will remain in a Registration Pending state.

1. If this occurs the Intelligent Agent on the node with which you are trying to register the event is down, or the node is not connected to the network. Check the status of the Intelligent Agent by selecting the node on which the Intelligent Agent is running and viewing the Node property sheet. You can also ping the Intelligent Agent to check its availability.
2. The node with which you want to register the event was defined manually (without using the Intelligent Agent). Connections to a manually defined node will not allow you to utilize remote management functionality such as Jobs or Events. You must first de-register any jobs or events against the node, remove

the node from the Console navigator, and then rediscover the node while it is running a 7.3.4 or later Intelligent Agent.

Event Menu

The Event menu allows you to set up event and administrator information. This menu also provides options to register, track, and view specific events. Menu options are enabled or displayed according to the items selected in the Event pane. See Figure 6-6, "Event Menu and Detail View" for an illustration of the Event menu.

Note: When you register or remove an event, there is usually a slight delay while the Intelligent Agent processes the request.

Create Event

Displays the Event property sheet and allows you to create the definition of a new event. See "Event General Page" on page 6-46 for more information.

Create Event Like

Available when an existing event is selected in the Console's Event detail view, this option displays the Event property sheet with the same page and parameter settings as the selected event. You can then save the event as under another event name.

Edit Event

Displays the definition of the selected event and allows you to edit the event. This menu option appears when an event is selected in the Registered page.

Edit Event Occurrence

Displays the definition of an existing event. See "Event General Page" on page 6-46 for more information.

Acknowledge

Acknowledges the selected event in the Alerts page. When an event triggers, an entry is added to the Alerts page. In the severity column, a flag of the appropriate color is displayed along with a pair of eyeglasses. The eyeglasses also appear whenever there is a change in the status of the event (e.g. from 'warning' to 'critical') If you choose to "acknowledge" this event, then it means you are aware of this event occurrence and hence the eyeglasses will disappear. This is useful in multi-administrator environments where the presence or absence of eye glasses indicates whether or not someone has looked at the event.

Copy to Event Library

Copy the selected event in the Event pane to the Event Library.

Deregister Event

Deregisters the event. This menu option only appears when an event test is selected in the Registered page.

Move to History

Moves the selected event in the Alerts page to Event History page of the Event pane. This option is enabled when an item is selected in the Alerts page.

Refresh Event History

Updates the History pane with the most recent entries.

Clear Event History

Clears the contents of the Event History page.

Event Library

Displays the Event Library dialog. See "Event Library Dialog" on page 6-41 for more information.

Context-Sensitive Menus

If you select an item in the Event pane with the right mouse button, the context-sensitive menu for that item appears. This menu is a subset of the Event menu plus selection-specific menu options.

Event Library Dialog

The Event Library dialog displays the events that have been created and saved to the Event Library. The advantage to using the Event Library is that both events and any associated target information can be stored, copied, or modified in the library for future use. When you create an event, you have the option of submitting, saving to the Event Library, or submitting and saving to the Event Library.

This dialog contains the following information:

Event

Name of the event.

Owner

Administrator who created the event.

Editing an Event in the Event Library

Select an event and click Edit to display the property sheet for the library event. The property sheet allows you to view and modify the library event. In addition to editing, you can perform a wide variety of event-related operations such as deleting, registering, and creating new events based on an existing event in the Event Library. If an event of the same name is already actively running, you must first remove the active event from all targets before registering it again.

Refresh

Updates the library events with the current definition at any time.

Oracle Event Tests

Several predefined event tests have been installed with Oracle Enterprise Manager. These appear in the Tests page of the Event property sheet, depending on the target type selected on the General page. You can add these tests to an event. The tests include:

- Database UpDown: checks whether a database is up or down.
- Host UpDown: checks whether a node is up or down.
- Net UpDown: checks whether a listener is up or down.
- HTTP Server UpDown: checks whether a monitored webserver is up or down.

Note: Only the UpDown tests are included with Oracle Enterprise Manager. Additional advanced event tests are available with the optional Oracle Diagnostics Pack. Refer to the *Enterprise Manager Event Test Reference Manual* for a complete list of advanced event tests.

To view the specific tests assigned to an event, double-click on the event in the Event Library dialog and view the Test page of the Event property sheet. See the online help for Oracle events, "Oracle Event Tests" on page 6-56, or the Diagnostic Pack documentation for information on Oracle event tests and their parameters.

Event Viewer

The Event Viewer property sheet displays details on a selected event in the History or Alerts page. When an event triggers, you select the triggered event and bring it up in the Event Viewer. The Event Viewer contains information on why the event triggered. You can also assign the event to a particular administrator and put instructions for other administrators via the Log page.

You can enter optional comments in the Log page, which is good way to share information about an event with other administrators. Once cleared, events are automatically moved to the History page. The pages of the Event Viewer include:

- General
- Log
- Notification Details

Event Viewer: General Page

The Event Viewer General page displays statistics and author information on a selected event. To obtain information on how to respond to an event occurrence, refer to the "User Action" section of the individual event test:

The following statistics are displayed:

Target

Destination of the event.

Target Type

Database, listener, node, or HTTP server.

Last Updated

Time of last update

Owner

Administrator that created the event.

Assigned To

List of administrators to which the event can be assigned. These administrators have at least "view" access to the event.

Show Event Definition

Displays the Edit Event property sheet in view mode.

Test Name

Event test that is performed.

Severity

Severity of the event occurrence: critical, warning, clear, or unknown.

Time/Date

Time and date of the event occurrence.

Message

Message generated from the alert.

Event Viewer: Log Page

The Event Viewer Log page displays an entry whenever an event is moved to history. An event can be moved manually with the Move to History menu option or automatically when the severity of the event changes.

The Log page also allows comments to be entered on a selected event. Any administrator with permissions to modify the event can add comments in this page. Administrators can enter tips on how to resolve the problem which might be useful for other administrators. You enter comments in the text box and select the Apply or OK button to add the comment.

The information displayed in the Log page includes:

Type Entry

Text input field allowing you to add comments.

Entry

Comment that has been entered for this event.

Author

Administrator that entered the comment.

Date/Time

The date and time when the comment was entered.

Event Viewer: Notification Details Page

The Event Viewer Notification Details displays details of email and paging notifications sent for a selected event. The information displayed in Details page includes:

Severity

The severity flag associated with the event occurrence.

Administrator

Administrator that was notified.

Date/Time

The date and time of the notification.

Method

Method of notification: E-mail or page.

Notification Status

Status of the notification, indicating whether the notification was sent, is pending, or has failed.

Message

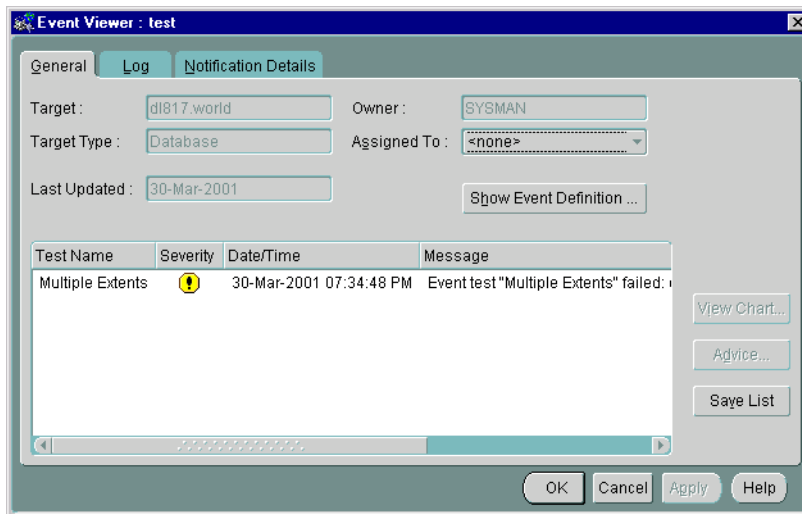
If the notification failed, this message indicates the reason for notification failure.

Responding to Event Occurrences

The online help for each event test will, in general, have a "User Action" section that provides guidelines on how to respond to that particular event tests should it trigger. See the online help Contents page for all available event tests.

Administrators can also obtain diagnostic information about the triggered event from the "View Chart" and "Advice" functionality available from the Event Viewer.


Figure 6–7 Event Viewer



The "View Chart" button allows administrators to look at real-time charts related to the event. The "Advice" button provides administrators diagnostic information to help them address the event condition appropriately.

Event General Page

On the General page, you determine the event name, target type, description, and targets to be monitored.

Figure 6–8 Event General Page**Event Name:**

Enter an event name.

Target Type:

Select the target type you want to monitor from the pull-down list. The types include Database, Listener, Node, or other service that is integrated into the Console.

If the selected Target Type is "Node", then a second pull-down list of operating systems will appear. If you choose 'All', then event tests that apply to all types of nodes, i.e. operating systems, will be available. If you choose a particular operating system, (e.g. Solaris), then additional operating-system specific event tests will be available.

The selection of the Target Type determines the list of Available Target. If you choose "Node" and a particular operating system, such as Solaris, then the list of available destinations will show all Solaris nodes that are running at least an 8.1.7 or higher Intelligent Agent. Any Solaris nodes that use older agents will not be shown.

Events can be registered against targets that have an Intelligent Agent. Targets on manually discovered nodes cannot be used as targets for an event. Hence, these

nodes will not appear on the Available Targets list. When an event is registered against a group, it will only be registered against targets that are running an Intelligent Agent. It will not be registered on any target that has been manually discovered.

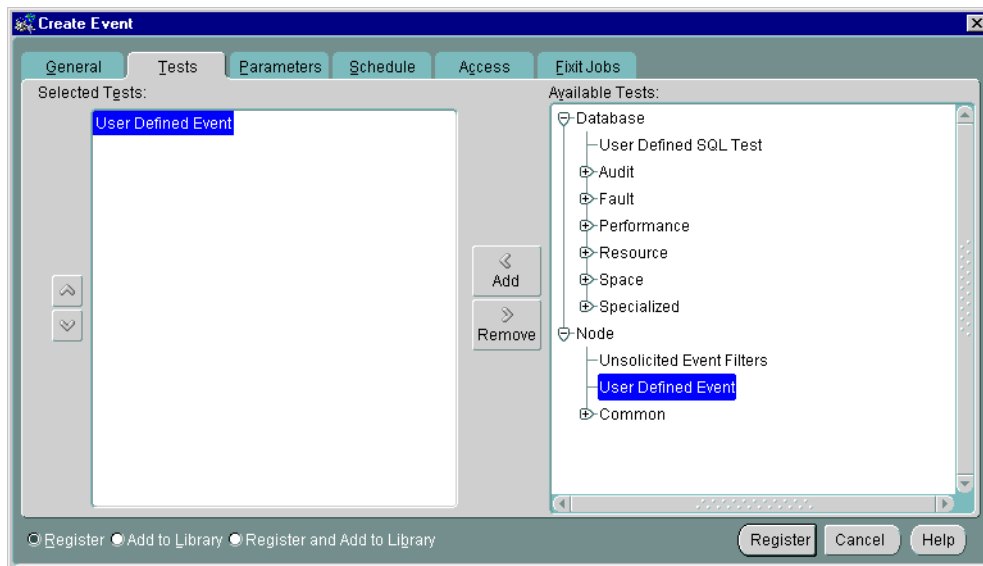
Event Description:

Enter a description or comment for the event

Event Tests Page

On the Tests page, you determine the event tests that you want to perform. Event tests are arranged hierarchically in a tree list for ease of viewing and selection. As with the Console Navigator, you can expand and compress entries in the tree list.

Figure 6–9 Event Tests Page



Available Tests:

Select the event tests in the list you want to perform in this event, then click on the << (Add) button to move the events to the Selected Events list. Double-clicking on an Available test will also move it to the Selected Tests list.

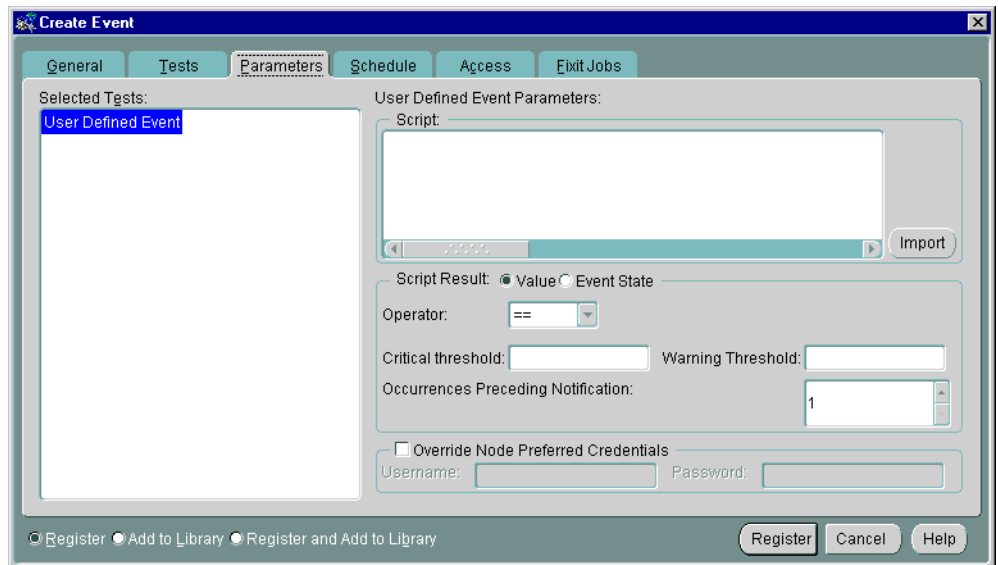
Selected Tests:

Select the event tests in the list you want to remove from this event, then click on the >> (Remove) button. Double clicking on a Selected test will also remove it from the Selected Tests list.

Event Parameters Page

The parameter settings for the selected event tests are entered in the Parameters page of the Event property sheet. The settings and types of parameters vary according to the event test selected. Some event tests do not have parameters. See the online help for Oracle events and "Oracle Event Tests" on page 6-56 for information on tests and their parameters. Further information on event tests is available in the Oracle Enterprise Manager Event Test Reference Manual.

Figure 6–10 *Event Parameters Page*

**Parameters**

The parameters for an event are displayed when the event is selected in the Selected Tests list. The parameters vary according to the event selected. Some events do not have parameters.

You can accept the default values or change the values for the parameters. To enter parameter values for an event, you can enter a value directly into a parameter field.

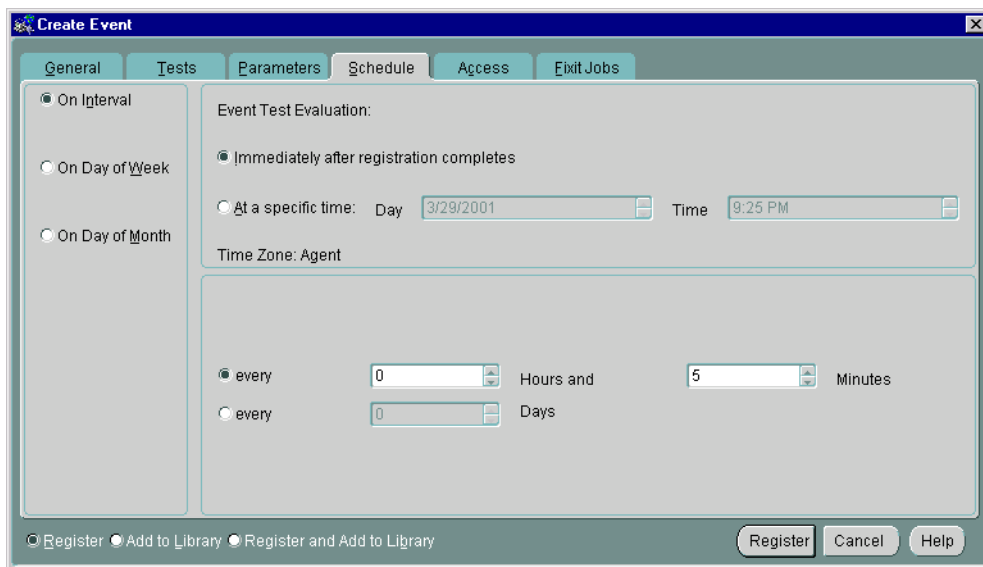
Filtering

Filtering is used in events such as **Chunk Small** and **Maximum Extents**. Examples of filters are = 'SYSTEM', LIKE '%SMP%', and IN ('SYSTEM', 'TOOLS'). Note that the quotes are single quotes. Use uppercase to match the case of the database object names. If you enter a filter value that does not select any objects or is an incorrect value, the event fails.

Event Schedule Page

The **Schedule** page allows you to schedule the evaluation of an event condition. This allows you to schedule resource-intensive events at off-peak times.

Figure 6–11 *Event Schedule Page*



You can select when you want event evaluations to occur. The choices are:

On Interval

Allows you to schedule a specific time interval at which the event monitors for a specific condition. The interval can be a combination of hours and minutes, or

number of days. Select the value you want to change and click on the scroll buttons. You can also type in a new value. This is the only schedule type allowed when there are targets running pre-9i Intelligent Agents in the "Selected Targets" list found on the General page.

On Day of Week

Allows you to schedule event monitoring on one or multiple days (Sunday, Monday, etc.) of the week. Click on the days of the week to select the days you want the event scheduled. (Available for targets running 9i versions of the Intelligent Agent)

On Day of Month

Allows you to schedule the event on one or multiple days (1 - 31) of the month. Click on the dates of the month to select the dates you want the task scheduled. (Available for targets running 9i versions of the Intelligent Agent)

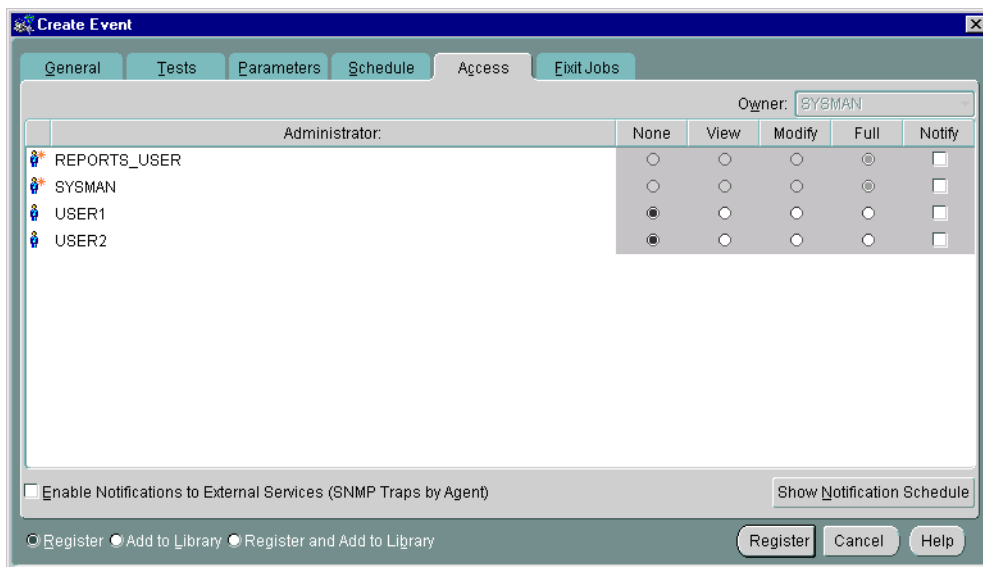
If you choose a day, such as 31, that is not in a month, the event will not be evaluated in that month.

Only the Intelligent Agent time zone is available with this release. Here, the Intelligent Agent schedules event monitoring at each destination based on the actual system time of each Intelligent Agent.

Event Access Page

Determine the administrator access permissions that you want to assign to the event with the Access Page. This allows other administrators to view or modify the event. Notifications are also assigned with this page.

Figure 6–12 Event Access Page



The levels of permission that you can assign to an Enterprise Manager administrator are:

None

Does not allow the administrator to view this event anywhere.

View

Allows the administrator to view the event, inspect event properties, and receive notifications.

Modify

Allows the administrator to modify the event’s log (See "Event Viewer" on page 6-42), enable enhanced notifications for other administrators, change event attributes in the event library, and assign triggered events to other administrators.

Full

Allows the administrator to delete the event, modify permissions for other administrators, change event attributes in the event library, clear the event history, and assign triggered events to other administrators.

Notify

Allows the administrator to receive enhanced event notifications on the objects through paging or email. Other notifications will be routed to that particular administrator's Console. Notify permission cannot be assigned if the administrator's permission level is set to None.

Any permissions assigned on this page supersede any administrator default permissions. See "Access" on page 1-23 for more information. Also, the administrator's notification schedule must be set up in order for them to receive the Email/page notification. Superusers cannot be changed from "Full" permissions.

Enable Notifications to External Services (SNMP Traps by Agent)

When checked, permits external notification (SNMP traps) to be sent from the supported SNMP service on the Intelligent Agent node. See the *SNMP Support Reference Manual* for more information.

Show Notification Schedule

Show Notification Schedule displays the notification schedule for the event. The schedule shown on this page is a combined schedule for all administrators that have been given "Notify" privileges for this event. To view administrators assigned to a particular time slot, use the right mouse button to call up the context-sensitive menu, choose the "Remove Recipient" option, and view the list of administrators. To add or remove notifications for an administrator, display the context menu (press the right mouse button) on any time block. The context menu provides options for adding and removing recipients of the notifications.

Table 6-3 summarized user permissions required to perform specific actions within Enterprise Manager.

Table 6-3 User Permissions Table

Action	None	View	Modify	Full	Owner	Super User	Comments
EVENTS - Dynamic Modification of Registered Events							
View progress/details	No	Yes	Yes	Yes	Yes	Yes	Information label appears in General page
Receive notifications (if enabled for administrator)	No	Yes	Yes	Yes	Yes	Yes	
Set permissions for any administrator including yourself	No	No	No	Yes	Yes	Yes	

Table 6–3 User Permissions Table

Action	None	View	Modify	Full	Owner	Super User	Comments
Set Notification checkbox for any administrator	No	No	Yes	Yes	Yes	Yes	
Enable SNMP traps	No	No	No	No	Yes*	No	* New behavior for 9i
Add/remove targets, change description, tests, parameters, schedule, fixit job	No	No	No	No	Yes*	No	* New behavior. Also depends on Intelligent Agent version.
Change owner	No	No	No	No	No*	No	* New behavior. When an administrator is deleted, events are reassigned to the new owner
EVENTS - In the Library							
Change owner (library)	No	No	No	Yes	Yes	Yes	
Add/remove targets	No	No	Yes	Yes	Yes	Yes	
Change description, tests, parameters, schedule, fixit job	No	No	Yes	Yes	Yes	Yes	
Change permissions; enable/disable Notify preferences; enable SNMP	No	No	Yes	Yes	Yes	Yes	
Delete event	No	No	No	Yes	Yes	Yes	
Submit event from the library	No	Yes	Yes	Yes	Yes	Yes	
EVENTS - In the Console							
Delete registered event	No	No	No	No	Yes	Yes	
Clear history	No	No	No	Yes	Yes	Yes	
Assign Event occurrences	No	No	Yes	Yes	Yes	Yes	

Event Fixit Jobs Page

A fixit job is designed to automatically correct a problem when a particular event condition is encountered. For example, you may want the Intelligent Agent to run a job to restart a database when the database instance has shut down unexpectedly. Fixit jobs are created with the Job system and must be designated as fixit jobs. The jobs must be submitted and running on the same destination that the event is set on.

The Fixit Jobs page consists of the following:

If ANY test triggers, run a fixit job:

When selected, allows a fixit job to be associated with the event. When any event test in the "selected Tests" triggers, the fixit job will run.

Fixit Job:

Drop-down list containing existing fixit jobs. If no fixit jobs currently exist, click Create to display the Create Job property sheet. Note: A newly created fixit job will not show up in the drop-down list during the current editing session. The event must be closed and then re-edited before the new fixit job will appear in the list.

Edit:

Displays the Edit Job property sheet for the fixit job selected in the Fixit Job drop-down list. The fixit job owner can edit some attributes of the fixit job.

Create:

Displays the Create Job property sheet which allows you to create a new fixit job.

Selected Tests

Displays all event tests chosen for the current event.

Note: Each event must use a unique fixit job on each destination where the event is registered. Also, when a single agent is monitoring multiple databases at a destination, create a separate event and fixit job for each database.

Event Progress Page

The Event Progress page displays when you edit an event from the Registered page of the Events pane. This page provides the current registration status for the event selected: Registered, Registration Failed, Modification Pending, or Registration Pending. In addition, the target and time and date when registration was attempted is shown.

When the Progress page is displayed, it shows only the status for the selected event. If the selected event is registered, or had been submitted for registration on other targets, you can view the status of this event for those targets by selecting the desired target from the Target pull-down list. The status of the event displays for that target. To view the status of this event for all destinations simultaneously, select <All>.

The following options are available on the Progress page:

Target (pull-down list)

Select the destination of the event you want to view from the pull-down list. Select <All> for all destinations for which this event has either been registered or failed to be registered.

Status

Status for the event: Registered, Registration Pending, Modification Pending, or Registration Failed.

Target

Network destination for the event.

Date/Time

Date and Time the event was submitted for registration.

Show Output

Displays the Event Status Message dialog. This button is active only when you have selected a failed event registration. Selecting this option will allow you to view the reasons for the failure.

Save List

Saves the contents of the list to a text file.

Administrator Event Notification

Oracle Enterprise Manager allows you to specify administrators that are notified when a particular event condition occurs. Each administrator can be associated with an email ID and/or a pager number. When using a paging service or email notification, each administrator can be assigned responsibility for specific systems at specific days and times.

For more information on setting up Oracle Enterprise Manager administrators, see "Managing Enterprise Manager Administrators" on page 1-8.

Oracle Event Tests

This section lists the Event system event tests with their parameters and return values. See "Event Parameters Page" on page 6-49 for information on entering parameter values. A list of event tests with numeric pager event Ids is also provided. See "Numeric Pager Job/Event Ids" on page 6-57 for more information.

Event tests are specified for database, listener, http, and node services. The event tests are also divided into fault, space, resource, and performance management

categories. Only the UpDown event tests are included with Oracle Enterprise Manager. Additional advanced event tests are available with the optional Oracle Diagnostics Pack. See the Oracle Enterprise Manager Event Test Reference Manual for complete information on available event tests. Complete event test information is also available from online help.

Some of the database event tests, such as Chain Row, require access to system tables and require additional permissions. You need to set up preferred credentials for the monitored database with an administrator that has system privileges. See "Enterprise Manager Monitor Role" on page 6-3 and "Preferred Credentials" on page 1-25 for more information.

Numeric Pager Job/Event Ids

The Event Management System provides paging services that notify an administrator with a page when an event has occurred. Alphanumeric pagers provide a brief text message identifying the event. Numeric pagers provide the numeric pager event Ids to identify the event.

For job notifications, you will receive a 6 digit number. The first 3 digits indicate the job-id. The last 3 digits indicate job status.

For event notifications you will receive the event ID with the status code.

For a complete list of pager job/event IDs, see "Paging Status Codes for Numeric Pages" on page 1-20

Event System Features and Requirements

Because the Enterprise Manager framework is a three-tier system that can manage a heterogeneous environment, it is important to keep in mind various software version requirements necessary for proper event system operation. Table 6-4, "Event Features and Associated Requirements" lists event system features and associated software version requirements.

Table 6–4 Event Features and Associated Requirements

Feature Name	Description	Enterprise Manager Version	Required Agent	Management Server/Console Required	Works In Browser
Advanced Events	All events for databases, nodes, listeners. See Enterprise Manager online help for more information.	Diagnostics Pack 1.5.5 and higher	All supported agents, latest recommended	For Enterprise Manager 2.x, the Management Server and Console that corresponds with that Pack	yes
Event Handler	Component that allows you to log event information or execute custom commands in response to an event occurrence. See	9.0.1 and higher	n/a	9.0.1 and higher	n/a
Improved Node Up/Down Monitoring	Enhancement to the Node Up/Down event test. Provides more information on whether or not the node is down, the agent is down, etc.	2.2 and higher	all supported	2.2 and higher	yes
User-Defined SQL Test	Allows you to write your own custom SQL to monitor database events	Diagnostics Pack 2.1 and higher	8.1.6 and higher	2.1 and higher	yes
Enhanced monitoring for target subcomponents	"For events whose targets involve multiple subcomponents (e.g. monitor tablespace full for ALL tablespaces), information on which subcomponent is in alarm is now provided	"2.2 and higher	8.1.7 and higher	2.2 and higher	yes
Context sensitive help for Event tests	"In the Parameters tab of the Event dialog, invoking ""Help"" will bring up information pertinent to the current selected event test	"2.2 and higher	n/a	2.2 and higher	yes
Events with synonymous event tests	"Events can be created that have more than one of the same event test (e.g. a ""Tablespace Full"" test for ""SYSTEM"", another ""Tablespace Full"" event test for ""USER"")	"2.2 and higher	all supported by Enterprise Manager 2.2	2.2. and higher	yes
Job and Events Notification filters					

Table 6–4 Event Features and Associated Requirements

Feature Name	Description	Enterprise Manager Version	Required Agent	Management Server/Console Required	Works In Browser
■ Filters that apply to both paging & email	Allows you to filter pages & emails based on job and event status	2.1	all agents supported by Enterprise Manager 2.1	2.1	yes
■ Different filters for paging & email	Allows separate filters for pages & emails based on job and event status	2.2 and higher	all agents supported by Enterprise Manager 2.2	2.2 and higher	yes
Customization for paging & email messages	Allows you to customize the messages for email and pages	Diagnostics Pack 2.2. and higher	all agents supported by Enterprise Manager 2.2	2.2 and higher	yes
Advanced O/S event tests	New event tests that monitor operating system specific metrics	2.2 and higher	8.1.7 and higher	2.2 and higher	yes
User-Defined Events	Allows you to define events based on any user-specified monitoring script.	9.0.1 and higher Diagnostics Pack	9.0.1 and higher Intelligent Agent	9.0.1 and higher and higher	yes
Dynamic modification of registered events	Allows you to dynamically change attributes of registered events.	9.0.1 and higher	all Intelligent Agents supported by Enterprise Manager 9i. 9.0.1 and higher versions of the Intelligent Agent allows full modification	9.0.1 and higher	yes
Event Schedules	Allows you to specify event evaluations based on a schedule.	9.0.1 and higher	9.0.1 and higher	9.0.1 and higher	yes
Event Integration with Performance Manager charts	Allows you to create events from Performance Manager charts.	9.0.1 and higher Diagnostics Pack	9.0.1 and higher	9.0.1 and higher	yes
Oracle9iAS Events	Events to monitor Oracle9iAS	9.0.2 and higher	9.0.2 and higher	9.0.2 and higher	yes
Real Application Clusters Events	Events to monitor Real Application Clusters-specific metrics	9.0.1 and higher	9.0.1 and higher	9.0.1 and higher	yes

Table 6–4 Event Features and Associated Requirements

Feature Name	Description	Enterprise Manager Version	Required Agent	Management Server/Console Required	Works In Browser
Concurrent Manager Events	Events to monitor error conditions against the Oracle Applications Concurrent Processing Server	2.0.4 and higher	8.1.5 and higher	2.0.4 and higher	yes
Forms Server Events	Events to monitor error conditions against the Oracle Developer Forms Server	2.0.4 and higher (2.0.4 console needs Forms Extensions.	8.0.6 and higher (requires Forms Agent Extensions.	2.0.4 and higher (2.0.4 console needs Forms Extensions.	yes
Program Filtering in Concurrent Manager Events	Allows you to monitor particular Oracle Applications Concurrent Programs. Also allows you to exclude particular Concurrent Programs from being monitored.	2.2 and higher	8.1.7 and higher	2.2 and higher	yes

Event Handler

Enterprise IT practices may require that responses to certain event occurrences be handled in certain ways. For example, if the database updown event triggers, administrators may want to automatically open an in-house trouble-ticket so that the appropriate IT staff can respond to this event occurrence. The ability to provide customized automatic responses to event occurrences can be achieved by using the Event Handler. This chapter discusses the following topics:

- Event Handler Overview
- How the Event Handler Works
- Setting Up the Event Handler
- Known Issues
- Differences between UNIX and Windows NT
- Running the Event Handler in a multi-Management Server Environment
- Migrating from Prior Releases

Event Handler Overview

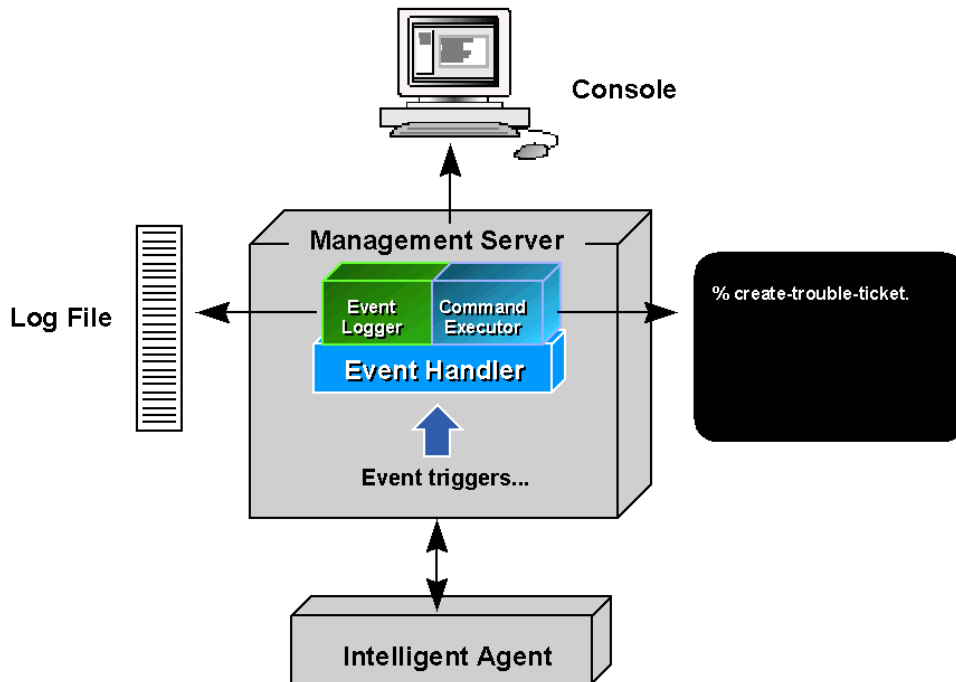
The Event Handler is an integral part of the Oracle Management Server. It listens for event notifications and responds to these events in ways specified by the administrator. The Event Handler's response capability is performed by its two components:

- *Event Logger* - allows the Event Handler to log events to designated log files
- *Command Executor* - allows simple operating system commands to be executed in response to an event occurrence.

You can use either or both components in response to a triggered event.

Prior to passing event notifications for further processing, the Event Handler also provides a simple filtering mechanism that allows system administrators to specify the conditions by which the events are passed to either or both components.


Figure 7-1 Event Handler Architecture



How the Event Handler Works

When the Event Handler starts, information from the *filters* are used to determine which events the Event Handler should select and pass on to the Event Logger and/or Command Executor. Once these components receive the event, they further process the event based on the *templates* set up for them. Templates provide a way to customize the behavior of the components. For the Event Logger, this means specifying which events should be logged and the specific log file to be used. For the Command Executor, this means specifying the types of events it should respond to and the operating system command it should execute in response to that event.

Figure 7-2 Event Handler Process



Setting Up the Event Handler

By default, the Event Handler has been pre-configured with default parameters and can be enabled immediately after installing Enterprise Manager. The default configuration implements the Event Logger only. To enable the Command Executor and/or customize the Event Handler's filtering capability, you will need to customize the Event Handler configuration. The two methods are outlined below.

Quickstart Method (Default)

1. Before enabling the Event Handler, the Management Server must first be stopped. To stop the Management Server, issue the following command:

```
% oemctl stop oms <superuser>/<password>
```

2. Enable the Event Handler by typing the following:

```
% oemctl enable eventhandler
```

3. Start the Event Handler by starting the Management Server

```
% oemctl start oms
```

To view the default configuration of the Event Handler type the following:

```
% oemctl dump eventhandler
```

The Event Handler is now ready to use. This default configuration will log all events to a file called `eventhandler.log`, which is located in the `%ORACLE_HOME/sysman/log` directory. By default, the format of an Event Logger message is:

```
<name>;<occurrence number>;<timestamp>;<assignee, if any>;<severity>
```

and is overwritten each time the Event Handler starts.

The format and behavior of the default log message can be modified slightly by using the following properties in the `omsconfig.properties` file.

```
/com/oracle/sysman/em/eventlogger/logfile=<full logfile path>
/com/oracle/sysman/em/eventlogger/separatorstring=<separator string for logfile entries>
/com/oracle/sysman/em/eventlogger/appendonstart=<true/false>
```

The `separatorstring` property is used by the default Event Logger template ONLY and cannot be used with user-defined Event Logger templates. This string separates placeholder entries within a single event occurrence, not different occurrences of the event.

The `appendonstart` property determines whether or not the log file will be overwritten each time the Event Handler starts. The default is false, which means overwrite. Set this to true if new log entries should be appended to an existing log file.

Customizing the Event Handler Setup

To change the default configuration, you can customize the Event Handler as described in this section.

1. Before enabling the Event Handler, check to make sure the Management Server is not running. To stop the Management Server, issue the following command:

```
% oemctl stop oms <superuser>/<password>
```

2. Enable the Event Handler by typing the following.

```
% oemctl enable eventhandler
```

3. Configure the Event Handler parameters.

The Event Handler parameters are stored in the Enterprise Manager repository. In order to change the parameter settings, you must export the parameters to a file, change the parameters, then import them back into the repository. These steps are illustrated below.

Note: Make sure the Event Handler is first enabled as explained in step 2.

Export parameters to a text file

At the command line enter the following:

```
% oemctl export eventhandler <filename>
```

Example: % oemctl export eventhandler myEventHandler

This example exports the parameters to a text file called myEventHandler.

Change the parameters

Use any text editor to edit the contents of the file containing the exported parameters.

Details on the parameters are explained in "Event Handler Configuration Parameters" on page 7-7. It is important to follow the exact syntax for the parameters.

Import the parameters

Once you complete the parameter changes, they must be imported back to the repository. At the command line enter the following:

```
% oemctl import eventhandler <filename>
```

Example: % oemctl import eventhandler myEventHandler

This example imports the Event Handler settings defined in the file myEventHandler back into the Enterprise Manager repository.

Important: Importing the parameter file will completely override any previous Event Handler settings.

4. Start the Management Server.

Starting the Management Server will start the Event Handler with the new settings.

```
% oemctl start oms
```

Event Handler Configuration Parameters

The Event Handler parameters are single-line entries that specify the events to which the Event Handler should respond and actions to be taken in response to these events.

Important: The syntax for each entry should be followed exactly. There must be NO linefeed or carriage returns within a single entry.

Event Handler parameters fall into three categories: Blackouts, Filters, and Templates.

Blackouts

Syntax:

```
eventhandler.respect_blackouts = <true/false>
```

The blackouts parameter tells the Event Handler whether or not to act on an event if the event triggered on a target that has paging and email blackouts set. This parameter can be toggled true or false.

A setting of False (default setting) instructs the Event Handler to continue processing events even if paging and/or email blackouts have been set for those events.

A setting of True prevents the Event Handler from processing events on targets that have paging/email blackouts.

Filters

When the Event Handler starts, it uses "filters" to determine which events are passed on to the Event Logger and/or Command Executor components. The system administrator can designate what events, if any, are sent to the Event Logger and Command Executor components. This is called "filtering." Events may be filtered on a global basis (applied to both components) or on a per component basis.

Filters apply to either the Event Logger or the Command Executor. Filters are defined in external files and imported into the Management Server by using the *oemctl import eventhandler* command. Each entry specifies the conditions by which the events are passed to the components. An event passes through the filter if it meets the conditions specified by the filter. Events that don't satisfy the condition are ignored by the Event Handler.

Each filter is uniquely identified by a name. All entries relating to a filter are grouped under this name. The filter conditions themselves are based on the following event properties:

Event properties for filters:

- `eventname` -- this is the name of the event
- `node` - this is the monitored node on which the event occurred
- `targetname` - name of the target
- `targettype` - type of the target. It is one of the following values:
 - `oracle_sysman_node` (target is a node)
 - `oracle_sysman_database` (target is a database)
 - `oracle_sysman_listener` (target is a net listener)
 - `oracle_sysman_cmanager` (target is a concurrent manager)
 - `oracle_sysman_ops` (target is an ops node)
 - `oracle_sysman_webserver` (target is an apache webserver)
 - `oracle_sysman_hotstandby` (target is a standby database)
- `owner` - owner of the event
- `severity` - status of the event. It is one of the following values: `alert`, `warning`, `error`, `clear`, `nodedown`

Filter Syntax The syntax for filters is as follows:

Syntax for global filters:

```
/com/oracle/sysman/em/eventHandler/global_filters/<filter-name>/<property-name> = <value>
```

Example:

```
/com/oracle/sysman/em/eventHandler/global_filters/myFilter/node = dlsun1234
```

Syntax for filters for the Event Logger only:

```
/com/oracle/sysman/em/eventHandler/eventlogger_filters/<filter-name>/<property-name> = <value>
```

Example:

```
/com/oracle/sysman/em/eventHandler/eventlogger_filters/myFilter2/owner = mary
```

Syntax for filters for the Command Executor only:

```
/com/oracle/sysman/em/eventHandler/commandexecutor_filters/<filter-name>/<property-name> = <value>
```

Example:

```
/com/oracle/sysman/em/eventHandler/commandexecutor_filters/myFilter3/severity = alert
```

Important: The 'value' specified must either be an exact match or a single wildcard character "*" to specify all possible values. Regular expressions are not supported. For example, 'myevent*' is not an acceptable value. The sense of the condition can be negated by prefixing a "!" in front of the value.

The following rules apply to all event filters:

1. All conditions specified by a given filter must hold true for the event.

Example: For the global_filter MyFilter:

```
/com/oracle/sysman/em/eventHandler/global_filters/MyFilter/eventname = cputest
```

```
/com/oracle/sysman/em/eventHandler/global_filters/MyFilter/node = prodserver.us.oracle.com
```

If the name of the event is 'cputest' and if it occurred on node 'prodserver.us.oracle.com', then the event will be passed to both Event Handler components.

2. The event must pass at least one filter.
3. An event is forwarded to a component (Event Logger or Command Executor) if it passes through either a global filter or a filter for that particular component.

By default, all events are suppressed, and at least one filter (global or per component) must be present in order for events to be forwarded to any of the adapters. See Sample Filters and Templates: on page 7-20 for examples of filters.

Assuming the event has passed through the filters, they are forwarded on to the event adapters, Event Logger or Command Executor.

Important: The default Event Handler configuration provides a global filter that allows all events to be passed to both components: Event Logger and Command Executor. In most cases, this should be sufficient. Further selection of the type of events to respond to can be specified via templates which are discussed in the next section.

Templates

Templates tell the Event Logger and Command Executor adapters how to respond to the event occurrence. For the Event Logger, a template specifies which events should be logged and how this information should be formatted and to which file the information should be logged. For the Command Executor, a template specifies the events the adapter should respond to and the operating system command it should execute in response to that event.

Event Logger Templates The Event Logger logs events that have passed through the event filters (as discussed in the previous section). You specify how the event information is logged via templates.

Templates are configuration entries that tell the Event Logger: the events to which the templates apply, the log file to use, and the format by which the event information should be written.

You can have multiple templates defined. Each template must be uniquely identified by name. Each template should specify the following:

The events to which the template applies

Use this format:

```
/com/oracle/sysman/em/eventlogger/templates/<template-name>/<property-name> = <value>
```

where:

- **<template-name>** is your name for this template
- **<property-name>** is the property of the event used to determine the events to which this template applies
- **<value>** is the value associated with the **<property-name>**

You can use any of the following event properties:

- *eventname* : name of the event
- *node* : the node on which the event triggered
- *targetname*: name of the target
- *targettype*: type of target
 - oracle_sysman_node (target is a node)
 - oracle_sysman_database (target is a database)
 - oracle_sysman_listener (target is a net listener)
 - oracle_sysman_cmanager (target is a concurrent manager)
 - oracle_sysman_ops (target is an ops node)
 - oracle_sysman_webserver (target is an apache webserver)
 - oracle_sysman_hotstandby (target is a standby database)
- *owner*: owner of the event
- *severity*: status of the event. Can be any one of the following: alert, error, warning, clear, nodedown

To specify more than one event property, use multiple entries.

The following example specifies the template "MyTemplate" should be used when an event is triggered on node dlsun1234 AND the event severity is alert:

```
/com/oracle/sysman/em/eventlogger/templates/MyTemplate/node = dlsun1234
/com/oracle/sysman/em/eventlogger/templates/MyTemplate/severity = alert
```

If, in a single template, you would like to use the SAME event property multiple times in a logical AND relationship, you need to append a number to the event property name such that they're unique.

For example:

```
/com/oracle/sysman/em/eventlogger/templates/HRTemplate/eventname1=!spaceEvent
/com/oracle/sysman/em/eventlogger/templates/HRTemplate/eventname2=!cpuEvent
/com/oracle/sysman/em/eventlogger/templates/HRTemplate/message=Event %eventname% triggered on node %node% and
has severity %severity%
/com/oracle/sysman/em/eventlogger/templates/HRTemplate/logfile=%ORACLE_HOME%/sysman/log/eventhandler.log
```

In the above example, the Event Logger will log a message of the specified format to a log file if the eventname is neither 'spaceEvent' nor 'cpuEvent'. This condition

is expressed as "eventname is not spaceEvent" AND "eventname is not cpuEvent". Note the addition of numbers to the eventname property (eventname1 and eventname2) as a way to uniquely identify the same property within the template HRTemplate.

To specify more than one property in a disjointed relationship (OR, for example), use separate templates as illustrated in the next example.

This example specifies "MyTemplate1" should be used if the event triggered on node dlsun1234, OR if the event triggered on target orcl817, use "MyTemplate2".

```
/com/oracle/sysman/em/eventlogger/templates/MyTemplate1/node = dlsun1234
```

...

```
/com/oracle/sysman/em/eventlogger/templates/MyTemplate2/targetname = orcl817
```

...

Format for Event Information

You must also specify the message format to use when logging the event information to a log file. The formatting string used can contain placeholders, which are symbolic representations for pertinent pieces of information about the event. Placeholders are enclosed within "%" characters to distinguish them from ordinary words in the template string. Available placeholders are:

Table 7-1 Event Handler Place Holders

Place Holder	Definition
%eventname%	The name of the event.
%severity%	The severity of the event as a string. (Alert, Clear, Warning, Node Down)
%timestamp%	The timestamp of the event occurrence. (MM-dd-yyyy hh:mm:ss Example: 05-22-01 05:22:00 AM)
%targetname%	Name of the target.
%targettype%	Type of the target. (oracle_sysman_node, oracle_sysman_database, oracle_sysman_listener)
%occ_no%	The occurrence number of the event.
%assignee%	The assignee for this event.
%node%	The node on which the event occurred, not to be confused with target.

Table 7-1 Event Handler Place Holders

Place Holder	Definition
%output%	The output associated with the event occurrence.
%owner%	The owner of the event.

To specify the message format that the template should use, the following format should be used:

```
/com/oracle/sysman/em/eventlogger/templates/<template-name>/message = <message format>
```

Example:

```
/com/oracle/sysman/em/eventlogger/templates/MyTemplate/message = %eventname% was triggered on %node% at
severity %severity%
```

This example generates a logged message of the form:

```
TablespaceUsage was triggered on dlsun1234 at severity warning
```

Log file

Each template must specify the log file to which it will write:

```
/com/oracle/sysman/em/eventlogger/templates/<template-name>/logfile=<logfilename>
```

Example:

```
/com/oracle/sysman/em/eventlogger/templates/MyTemplate/logfile = /u1/myhome/myevents.log
```

An optional property can be specified to indicate whether or not the logfile should be appended or overwritten for each session that the Event Handler starts:

```
/com/oracle/sysman/em/eventlogger/templates/<template-name>/appendonstart = <true/false>
```

The default value is false, meaning each time the Event Handler starts, any old logfile entries will be overwritten by new log entries.

Note: Forward slashes ("/") should always be used when specifying a path to ensure compatibility with both UNIX and Windows systems.

Multiple Event Logger Templates In cases where the conditions of more than one template are met, all matching templates will be executed. Since each template has its own log file, it is possible to log events to multiple log files on one event notification.

An example of an entry in a template file would be:

```
/com/oracle/sysman/em/eventlogger/templates/foo/eventname=cputest
```

```
/com/oracle/sysman/em/eventlogger/templates/foo/message=%eventname% fired on %node%: Cpu usage on %targetname% is high! occ_no: %occ_no% Severity: %severity% Time: %timestamp%
```

```
/com/oracle/sysman/em/eventlogger/templates/foo/logfile=%ORACLEHOME%/sysman/log/ev.log
```

```
/com/oracle/sysman/em/eventlogger/templates/foo/appendonstart=true
```

The template definition above assigns all events named `cputest` to the template named `foo`. The optional `appendonstart` entry has also been specified. Whenever a `cputest` event fires, the logged output will appear in the file `%ORACLE_HOME%/sysman/log/ev.log` as follows:

```
cputest fired on smptest16: Cpu usage on smptest16 is high! occ_no: 21 Severity: Alert Time: 10-21-2001 02:39:29 PM
```

The Command Executor Templates The Command Executor executes simple commands in response to event occurrences. The Command Executor looks at a set of user-defined templates to decide what command to execute in response to event occurrences.

Templates are configuration entries that tell the Command Executor: the events to which the templates apply and the command to execute.

You can have multiple templates defined. Each template must be uniquely identified by name. Each template should specify the following:

The events to which the template applies Use this format:

```
/com/oracle/sysman/em/commandexecutor/templates/<template-name>/<property-name> = <value>
```

where:

- **<template-name>** is your name for this template
- **<property-name>** is the property of the event used to determine the events to which this template applies, and

- **<value>** is the value associated with the <property-name>

You can use any of the following event properties:

eventname : name of the event

node : the node on which the event triggered

targetname: name of the target

targettype: type of target

owner: owner of the event

severity: status of the event. Can be any one of the following: alert, error, warning, clear, nodedown

As in the Event Logger, to specify more than one event property, use multiple template entries.

Command to Execute

Use this format:

```
/com/oracle/sysman/em/commandexecutor/templates/<template-name>/command = <value>
```

where:

<value> is the command to execute

To include information about the event in the command, use placeholders. Placeholders are symbolic representations about pieces of information about the event. Placeholders are enclosed within "%" characters to distinguish them from ordinary words in the command string. For a list of available placeholders, see Table 7-1, "Event Handler Place Holders".

Optional: length of execution time

Optionally, you can also specify the length of time (in seconds) that the Event Handler will wait for a process to terminate after execution. For example, if a command begins a process that takes longer to run than what is normally acceptable, the Event Handler will automatically terminate the process.

Use this format to specify the execution time:

```
/com/oracle/sysman/em/commandexecutor/templates/<template-name>/exectimeout =<value>
```

where <value> is expressed in seconds.

By default, if the `exectimeout` parameter is not specified, the Event Handler will terminate the process after 40 seconds.

A typical entry for a template would be:

```
/com/oracle/sysman/em/commandexecutor/templates/foo/eventname=cputest
```

```
/com/oracle/sysman/em/commandexecutor/templates/foo/command=net send my-machine-name %eventname%  
fired on %node%: CPU usage on %targetname% is high! occ_no %occ_no% Severity: %severity% Time:  
%timestamp%
```

where "my-machine-name" is the name of the machine to which you want to send a message.

The template definition above assigns all events named `cputest` to the template named `foo`. Whenever a `cputest` event fires, the following message will be sent to machine "my-machine-name":

```
cputest fired on smptest16: Cpu usage on smptest16 is high! occ_no: 21 Severity: Alert Time:  
10-21-1999 02:39:29 PM
```

There may be situations where the command you want to execute has an argument consisting of several items, for example:

```
foo a b c "this is a test" "d=e"
```

where:

`foo` is the executable command, and

- `a` is its first argument
- `b` is its second argument
- `c` is its third argument

`this is a test` is its fourth argument

`d=e` is its fifth argument.

In this case, you can use quotes (") to group items within a single argument together as indicated in the example above, e.g. "this is a test". To include literal quotes as part of the argument, you need to prefix these quotes with backslashes (\). For example:

```
foo a b c "d=\"e\" " "this is \"quoted\" "
```

will be pass the following arguments to `foo`:

a is the 1st argument

b is the 2nd argument

c is the 3rd argument

d="e" is the 4th argument

this is "quoted" is the 5th argument

Summary of Event Handler Configuration Commands

Typically, the way to configure the Event Handler would be to export the current configuration to a file, edit the file to change the entries, and then import the file back into the Enterprise Manager repository.

When a new Enterprise Manager repository is created, configuration entries are automatically created for all event notifications to be passed to the Event Handler, and for the Event Logger to log all event notifications into the file `$ORACLE_HOME/sysman/log/eventhandler.log`. However, the Event Handler is disabled by default.

The following is a summary of Event Handler Configuration Commands.

Enabling the Event Handler

To start the Event Handler, execute the following command:

```
oemctl enable eventhandler
```

Note: The Management Server must be stopped before executing this command.

Disabling the Event Handler

If at a subsequent time you need to disable the Event Handler, type

```
oemctl disable eventhandler
```

Viewing Current Event Handler Configuration Settings

To print out the current Event Handler configuration for viewing, type:

```
oemctl dump eventhandler
```

This will dump out the current Event Handler configuration to standard output. It will also indicate the Management Servers that are currently enabled with the Event Handler.

Creating a Configuration File from the Current Event Handler Configuration Registry Entries

To export the current Event Handler configuration registry entries into a file, type:

```
oemctl export eventhandler <filename>
```

You can then edit the configuration entries in the exported file. To import any changes, see the next section.

Importing a Configuration File

To import a file containing Event Handler configuration entries into the Enterprise Manager Repository, type:

```
oemctl import eventhandler <filename>
```

Importing the configuration file will completely override any previous Event Handler settings.

Important: None of the Event Handler configuration commands take credentials since all Event Handler configuration commands inherit the repository credentials from the `omsconfig.properties` file. For this reason, before using the Event Handler configuration commands, you must first configure the Management Server.

Troubleshooting Tips

If the Event Logger and/or Command Executor do not seem to respond to events, check the following:

- If the Command Executor does not run a specified shell script, perform the following:
 - Specify the shell name or interpreter name before the command to execute the shell script as shown in the following example.

```
> /bin/sh /db01/apps/oracle/eventresponse.sh
```
 - Ensure the first line of the shell script invokes the shell or interpreter. In the example, this is `/bin/sh`. The first line of the shell script should be `#!/bin/sh`.

- Make sure there are no linefeed characters within any single entry in the import file. The import utility will check for errors upon execution.
- Make sure you have set a filter that allows events to pass through to them. This is part of the default configuration.

Example:

```
/com/oracle/sysman/em/eventHandler/global_filters/allEvents/eventname=*
```

- Check the syntax of any templates associated with the Event Logger/Command Executor. Try using a simple template or any of the samples provided in this chapter.
- When using the Event Logger on UNIX, make sure you have the appropriate permissions to create the log file. The Event Logger uses the permissions of the operating system user who started the Management Server. Try creating a file in the directory where the Event Logger's log file should have been created.
- When using the Command Executor on UNIX, make sure you have the appropriate permissions to execute the command. The Command Executor uses the permissions of the operating system user who started the Management Server. Try running the command directly on the operating system to make sure it works.
- Make sure the Event Handler service is available by typing the following command at the command prompt (You can also use this command at any time to check the current configuration of the Event Handler.):

```
>oemctl dump eventhandler
```

- When using templates, make sure you specify both the event condition on which the template applies as well as the message format (for the Event Logger) or the command to be executed (for the Command Executor)
- More advanced troubleshooting can be achieved by tracing the Event Handler. Because the Event Handler is an integral part of the Oracle Management Server, the Event Handler is automatically traced when you set up tracing for the Management Server. Trace information about the Event Handler will be included in the Management Server's trace/log file. Refer to Appendix B of the *Oracle Enterprise Manager Configuration Guide* for more information about Management Server tracing.

Note: The Event Handler is an English-only release.

Sample Filters and Templates:

The following are sample filters and templates. Some of these entries can be found in the %ORACLE_HOME%/sysman/admin/ *EventHandler.examples* file.

Filters:

Pass all events to both Event Logger and Command Executor

```
/com/oracle/sysman/em/eventHandler/global_filters/allNodes/node = *
```

Pass all events to the Event Logger

```
/com/oracle/sysman/em/eventHandler/eventlogger_filters/allEvents/eventname=*
```

Pass all events to the Command Executor

```
/com/oracle/sysman/em/eventHandler/commandexecutor_filters/allEvents/node=*
```

Pass on all events except “cptest” to both Event Logger and Command Executor

```
/com/oracle/sysman/em/eventHandler/global_filters/not-cptest/eventname =  
!cptest
```

Pass on all events named cptest that did not originate on the node smptest16 to the Event Logger:

```
/com/oracle/sysman/em/eventHandler/eventlogger_  
filters/pass-some-cptest/eventname=cptest  
/com/oracle/sysman/em/eventHandler/eventlogger_  
filters/pass-some-cptest/node=!smptest16
```

Pass all events except cptest to the Command Executor adapter:

```
/com/oracle/sysman/em/eventHandler/commandexecutor_  
filters/pass-no-cptest/eventname=!cptest
```

Pass on all events whose name is foo or whose originating node is skini-pc to the Command Executor:

```
/com/oracle/sysman/em/eventHandler/commandexecutor_filters/passfoo/eventname=foo
/com/oracle/sysman/em/eventHandler/commandexecutor_filters/pass-skini/node=skini-pc
```

Templates

Sample template for the Event Logger:

```
/com/oracle/sysman/em/eventlogger/templates/allEvents/eventname=*
/com/oracle/sysman/em/eventlogger/templates/allEvents/message=%eventname% fired
on %node%: Target %targetname%: Output %output% Severity: %severity% Time:
%timestamp%
/com/oracle/sysman/em/eventlogger/templates/allEvents/logfile=ev.log
/com/oracle/sysman/em/eventlogger/templates/allEvents/appendonstart=false
```

A sample log entry using this template would look like:

```
cpptest fired on smptest16: Target smptest16: Output Cpu usage is high.
Severity: Alert Time: 10-21-2001 02:39:29 PM
```

Sample templates for the Command Executor:

For Windows NT:

```
/com/oracle/sysman/em/commandexecutor/templates/allEvents2/eventname=*

/com/oracle/sysman/em/commandexecutor/templates/allEvents2/command-net send
my-machine %eventname% fired on %node%: Target %targetname%: Output %output%
Severity: %severity% Time: %timestamp%
```

where my-machine is the name of the PC to which the message will be sent.

Whenever any event fires, a message using the above format is sent to machine my-machine.

For UNIX:

```
/com/oracle/sysman/em/commandexecutor/templates/allEvents/command=xterm
-display hqsun1:0 -e telnet
```

Here, whenever any event fires, a telnet session is opened.

Complex Example 1: If the event name is foo or the originating node is prod-pc, log the string "Event %eventname% occurred on node %node%" into the file *myevents.log*. If the event name is anything other than foo, execute the command "mail admin@acme.com -s %eventname% %node%":

```
/com/oracle/sysman/em/eventlogger/templates/foaname/eventname=foo
/com/oracle/sysman/em/eventlogger/templates/foaname/message=Event %eventname% occurred on
node %node%
/com/oracle/sysman/em/eventlogger/templates/foaname/logfile=myevents.log
```

```
/com/oracle/sysman/em/eventlogger/templates/prod/node=prod-pc
/com/oracle/sysman/em/eventlogger/templates/prod/message=Event %eventname% occurred on
node %node%
/com/oracle/sysman/em/eventlogger/templates/prod/logfile=myevents.log
```

```
/com/oracle/sysman/em/commandexecutor/templates/anythingelse/eventname=!foo
/com/oracle/sysman/em/commandexecutor/templates/anythingelse/command=mail skini@oracle.com
-s %eventname% %node%
```

Complex Example 2: If the event severity is Alert, execute the commands "pager %eventname% dbapager" and "mail dba@acme.com -s %eventname% Highest Priority!!!" If the severity is anything else, only execute the command "mail dba@acme.com -s %eventname% Normal Priority"

```
/com/oracle/sysman/em/commandexecutor/templates/alertsev1/severity=alert
/com/oracle/sysman/em/commandexecutor/templates/alertsev1/command=pager %eventname%
dbapager
```

```
/com/oracle/sysman/em/commandexecutor/templates/alertsev2/severity=alert
/com/oracle/sysman/em/commandexecutor/templates/alertsev2/command=mail dba@acme.com -s
%eventname% Highest Priority!!!
```

```
/com/oracle/sysman/em/commandexecutor/templates/sevanythingelse/severity=!alert
/com/oracle/sysman/em/commandexecutor/templates/sevanythingelse/command=mail dba@acme.com
-s %eventname% Normal Priority
```

Complex Example 3: If the eventname is neither 'spaceEvent' nor 'cpuEvent', then log a message of the format: "Event %eventname% triggered on node %node% and has severity %severity%" to a logfile called *eventhandler.log* in the directory %ORACLE_HOME%/sysman/log.

```
/com/oracle/sysman/em/eventlogger/templates/HRTemplate/eventname1=!spaceEvent
/com/oracle/sysman/em/eventlogger/templates/HRTemplate/eventname2=!cpuEvent
/com/oracle/sysman/em/eventlogger/templates/HRTemplate/message=Event %eventname% triggered
```



```
on node %node% and has severity %severity%  
/com/oracle/sysman/em/eventlogger/templates/HRTemplate/logfile=%ORACLE_  
HOME%/sysman/log/eventhandler.log
```

Known Issues

The Event Handler may not start if there are any carriage return characters within any of the required entries. Some errors are caught by the import utility. If a problem with the file is encountered during import, an error message is displayed.

Differences between UNIX and Windows NT

For Event Handler entries that require a directory path, use the syntax that is appropriate to the operating system. For example

UNIX:

```
/com/oracle/sysman/em/eventlogger/logfile=/app/oracle/8.1.6/sysman/log/MyEvents.  
log
```

Windows NT:

```
/com/oracle/sysman/em/eventlogger/logfile=c:\\orant\\sysman\\log\\MyEvents.log
```

Running the Event Handler in a multi-Management Server Environment

As part of the Management Server, the Event Handler has a high degree of reliability, failover, and load-balancing. However, certain rules apply when multiple Management Servers are connected to the same repository:

- By default, the Event Handler is **not** enabled in a newly configured Management Server and must be explicitly enabled for each Management Server you set up, by using the *oemctl enable eventhandler* command. Conversely, to disable the Event Handler on one or more Management Servers, you must explicitly disable each Management Server by executing the *oemctl disable eventhandler* command.

Note: The Management Server must be shut down when executing these commands.

- If more than one Management Server is Event Handler enabled, then each event notification is handled exactly once and can be handled by any of the enabled Management Servers. There is no way to specify which of the enabled Management Servers should pick up an event notification. This means that the commands executed by the Command Executor must be available in the same way to all enabled Management Servers. Likewise, if you are using the Event Logger, the log file must be seen by all enabled Management Server's in exactly the same way, i.e., the same pathname. On Unix systems, for example, this could be achieved by using NFS file systems.
- For situations where some Management Servers connected to a repository have the Event Handler disabled and others enabled, only Event Handler enabled Management Servers will pick up and execute Event Handler directives for event notifications (all Management Servers will continue to process event notifications in the usual way).

If one of the enabled Management Servers dies or is shut down, its Event Handler processing will failover automatically between the other enabled Management Servers. If all enabled Management Servers die (or are shut down) then Event Handler directives will not be processed by any of the remaining (Event Handler -disabled) Management Servers. In this situation, Event Handler operations will enter into the Management Server's reliability queue and will be processed as a batch once one or more of the Management Servers that were enabled for the Event Handler become operational.

Migrating from Prior Releases

If you used the Event Handler in Enterprise Manager 2.2 and wish to use the same configuration in Enterprise Manager 9i, then perform the following:

1. Copy all Event Handler entries from the omsconfig.properties file to a separate file. Use this separate file to make the changes described in the next steps.
2. If you're using the Event Logger

In Enterprise Manager 2.2, each template for the Event Logger shared the same log file. In Enterprise Manager 9i, each template has its own log file. To specify the logfile name, convert the template entry

From:

```
/com/oracle/sysman/em/eventlogger/logfile = <logfilename>
```

To:

```
/com/oracle/sysman/em/eventlogger/templates/<template-name>/logfile=<logfile
name>
```

Multiple templates can share the same logfile by specifying the same logfilename.

3. If you're using the Command Executor

No changes need to be done to the Command Executor templates.

Note: In Enterprise Manager 2.x, if multiple templates match the event condition, only one of the templates will be used. In Enterprise Manager 9i, if multiple templates match the event condition, ALL templates will be used by both the Event Logger and Command Executor.

4. Enable the Event Handler

The Management Server must first be stopped before doing this step. Stop the Management Server by:

```
% oemctl stop oms <superadmin/password>
```

Then enable the Event Handler:

```
% oemctl enable eventhandler
```

5. Export the current settings

Export the Event Handler's current default configuration to another file.

```
% oemctl export eventhandler <config_filename>
```

6. Update the configuration file

Update the configuration file created in step 5 by replacing the default templates with the template settings you created in steps 2 and 3.

7. Import the configuration settings

```
% oemctl import eventhandler <config_filename>
```

where <config_filename> is the file in step 6 that has the updated Event Handler configuration entries.

8. Restart the Management Server to start the Event Handler

```
% oemctl start oms
```

You can verify the Event Handler settings by:

```
% oemctl dump eventhandler
```

Enterprise Manager Reporting

Enterprise Manager Reporting provides administrators with an easy, yet powerful way to quickly view and analyze information about managed applications and systems.

This chapter covers the following topics:

- Key Concepts
- Configuring Enterprise Manager Reporting
- Enterprise Manager Reporting Website
- Creating a Report from an Existing Report Definition
- Editing a Report Definition
- Generating a Report from Enterprise Manager Applications
- Creating a User-defined Report Definition
- Navigating the Enterprise Manager Reporting Website

Enterprise Manager Reporting

The Enterprise Manager reporting system provides flexible reporting functionality to administrators, permitting quick and easy access to information about the status, configuration, and performance of all monitored systems in their enterprise. Administrators can create, schedule, and publish a wide variety of enterprise system reports. When published to a website, these reports can be accessed by a wider audience, enabling anyone from administrators to managers to executives to quickly access information regarding their monitored environment. The reporting functionality is fully integrated with the Enterprise Manager Job scheduling system, allowing reports to be generated automatically at specific times or at regular intervals. Reports can also be generated on-demand, such as when an administrator requests to view a specific report by clicking on a link within the reporting website.

Important: The reporting system is only available for Enterprise Manager Consoles connected to a Management Server. Consoles running standalone only have access to the minimal reporting functionality that is available through the standard database management tools.

See the *Enterprise Manager Configuration Guide* for information on reporting setup and configuration.

Key Concepts

Before using the reporting system to create reports, there are two fundamental concepts you must be familiar with: report definitions and report elements.

What is a Report Definition?

A key concept to understand is that all reports are generated from report definitions. An administrator creates a report definition that generates the desired report. A report definition defines what type of report is generated, its content and appearance, or whether it is published on the reporting website. It also includes various attributes such as an optional schedule. A report is uniquely identified by its definition name. Specifically, a report definition allows you to define the following attributes:

- Definition Name

- Owner
- Report type
- Report content
- When a report should be generated

Enterprise Manager supplies a wide array of predefined report definitions, allowing you to generate reports without having to create new report definitions. Select the Report Definition object in the Console Navigator to display a list of all report definitions in the detail view. See Figure 8–1, "Report Definitions in the Console Detail View".

If none of the predefined definitions meets your reporting requirements, you can modify definition parameters from an existing report definition and save it as a new report definition.

Figure 8–1 Report Definitions in the Console Detail View

Report Title	Owner	Category	Sub Category	Definition Name
Instance	REPORTS_USER	General	Configuration	EM_DB_InstanceConfig
Instance	REPORTS_USER	General	Current Status	EM_DB_InstanceStatus
OLAP	REPORTS_USER	General	Configuration	EM_DB_OLAP
Replication	REPORTS_USER	General	Configuration	EM_DB_REPLICATION
Database Information Related to Repository	REPORTS_USER	General	Configuration	EM_DB_REPOSITORY
Schema	REPORTS_USER	General	Configuration	EM_DB_Schema
Security	REPORTS_USER	General	Configuration	EM_DB_Security
Storage	REPORTS_USER	General	Configuration	EM_DB_StorageConfig
Storage	REPORTS_USER	General	Current Status	EM_DB_StorageStatus
Administrator Overview	REPORTS_USER	Setup	Administrators	EM_IMPL_ADMINISTRATORS
Intelligent Agent Overview	REPORTS_USER	Setup	Agents	EM_IMPL_AGENTS
Active Collections	REPORTS_USER	Setup	Collections	EM_IMPL_DG_COLL
All Outstanding Alerts	REPORTS_USER	Job/Event	Events	EM_IMPL_EVT_ALERTS
Outstanding Alerts in History	REPORTS_USER	Job/Event	Events	EM_IMPL_EVT_HISTORY
All Events in Library	REPORTS_USER	Job/Event	Events	EM_IMPL_EVT_LIBRARY
Event Notifications Sorted by Administrator	REPORTS_USER	Job/Event	Events	EM_IMPL_EVT_NOTIFY_ADMIN
Event Notifications Sorted by Name	REPORTS_USER	Job/Event	Events	EM_IMPL_EVT_NOTIFY_NAME
Registered Events Sorted by Name	REPORTS_USER	Job/Event	Events	EM_IMPL_EVT_REG_NAME
Registered Events Sorted by Target	REPORTS_USER	Job/Event	Events	EM_IMPL_EVT_REG_TARGET
Active Jobs	REPORTS_USER	Job/Event	Jobs	EM_IMPL_JOB_ACTIVE
Average Execution Time per Job	REPORTS_USER	Job/Event	Jobs	EM_IMPL_JOB_AVERAGE
Overview of Completed Jobs Run Yesterday	REPORTS_USER	Job/Event	Jobs	EM_IMPL_JOB_COMPL1
Overview of Completed Jobs Run Yesterday	REPORTS_USER	Job/Event	Jobs	EM_IMPL_JOB_COMPL7
Overview of Failed Jobs Run Yesterday	REPORTS_USER	Job/Event	Jobs	EM_IMPL_JOB_FAIL1
Overview of Failed Jobs Run Last Week	REPORTS_USER	Job/Event	Jobs	EM_IMPL_JOB_FAIL7
All Jobs in Library	REPORTS_USER	Job/Event	Jobs	EM_IMPL_JOB_LIBRARY

If your reporting needs go beyond the scope of the predefined report definitions, you can create completely new report definitions. The next section covers the different ways you can create a report.

As shown in Figure 8-1, "Report Definitions in the Console Detail View", existing report definitions are displayed in a multi-column list that provides information and current status for all report definitions. The columns shown are:

- *Status Icon*: Indicates whether a report definition is published (globe icon with a plus "+" subscript), not published (globe icon with a red "x" subscript), or scheduled (globe icon with a clock subscript).
- *Report Title*: Report title that appears on the generated report and on the reporting website.
- *Owner*: Owner of the report definition.
- *Category*: Relates to the generated report location on the reporting website. Major report categories appear as tabs along the top the reporting website. See "Navigating the Enterprise Manager Reporting Website" on page 8-27 for more information on reporting website organization.
- *Subcategory*: Relates to the generated report location on the reporting website. Each major category (tab page) has a series of subcategories listed along the left side of the category page.
- *Definition Name*: A unique report definition name that only appears in the in the Console list of report definitions and not in generated reports.

What is a Report Element?

Report elements are the building blocks of any report definition, and by extension, the generated report itself. By selecting and ordering report elements, you construct the format and content of your report. Enterprise Manager provides three general categories of report elements:

- *HTML*: Customize the HTML for a generated report and call SQL and Javascript functions for additional row-level processing.
- *Queries*: Specify SQL queries to generate charts or tables.
- *Service Levels*: Specify time-based statistics for monitored services.

The categories actually displayed depend on the chosen report type. See "Report Elements Page" on page 8-15 for an explanation of specific report elements for each of these categories.

Ways to Select Targets for Report Generation

Once your Enterprise Manager reporting environment has been configured, you can generate reports immediately by using the predefined report definitions. Regardless of whether you use a predefined report definition or a custom definition, there are two distinctions that must be made regarding report generation:

- Creating individual reports for all targets of a selected type
VS.
- Creating a single report for preselected targets

Creating Individual Reports for All Targets of a Selected Type

Reports generated using this method do not require that you choose specific targets before publishing or saving the report. In this situation, the administrator who views the report chooses the target(s). For example, a super administrator wants to publish a report listing the number of active jobs submitted against a target. They would like this report to be used by other administrators as a helpful utility to allow them to determine the job loads for targets they are managing. Since there are over 560 targets in their managed environment, it is not practical to generate a report for all 560 targets. Instead, the super administrator specifies in the report definition that an individual report be created for all targets of a selected type and then publishes the report to the Enterprise Manager reporting website. A regular administrator goes to the reporting website to find out how many jobs are active for the five targets they are managing. They click on the hyperlink to generate a report. They are then prompted by the system to specify the exact targets they wish to run the report against. They choose the targets and the report is generated for those targets only.

Creating a Single Report for Preselected Targets

Specific targets must be selected from the report definition before generating the report. Hence, the administrator who clicks on that report's hyperlink in the reporting website will see a report for all targets chosen by the owner of the report definition.

Configuring Enterprise Manager Reporting

You set up your reporting environment using the web server supplied with your Oracle installation. Upon installing Enterprise Manager components, all internal configuration of the Enterprise Manager Reporting web site is taken care of automatically.

Configuration requires a two-way exchange of information between the Reporting web server and the Oracle Management Server.

- The web server knows which Management Server it needs to connect to and the password of the REPORTS_USER administrator.
- The Oracle Management Server knows which web server is configured to run against an Enterprise Manager repository

For more information on Reporting setup and configuration, see the *Oracle Enterprise Manager Configuration Guide*.

Enterprise Manager Reporting Website

The primary advantage of the reporting system lies in its ability to publish reports to a website for other administrators and other users to access information on their environment. The reporting website (along with an Apache web server) is installed from the Server CD with Enterprise Manager.

As with any website, there is a home page. The Enterprise Manager reporting website home page provides both a summary of all managed target types and their status. Clicking on any target link displays all reports related to that target. From this page, an administrator, or anyone with intranet access can navigate to any published report for your managed environment.

Figure 8–2 Reporting Website Home Page

Enterprise Manager Reporting Home Page April 18, 2001 10:08:32 PM EDT

From this page you can access any Enterprise Manager report that has been published. Reports are accessed from the System State at a Glance table below and from [Additional Reports](#).

System State at a Glance

All Targets	Critical	Warning	Clear	Unknown	Error	Unmonitored
1 Concurrent Managers	0	0	0	0	0	1
42 Databases	3	0	0	7	0	32
12 HTTP Servers	0	0	0	1	0	11
32 Listeners	0	0	0	4	0	28
16 Nodes	0	3	1	2	0	10
1 Groups	0	0	0	0	0	1
All Targets	3	3	1	14	0	83

TIP To view a report for a specific target, click the appropriate target type in the first column above. Reports that include multiple targets are accessible from any of the included targets. Reports that are not target-specific are accessed from [Additional Reports](#).

By default, the reporting website uses the standard Enterprise Manager header. If you wish to use your own custom header, replace the oem.gif file found in the ORACLE_HOME/oem_webstage/sysman/reporting/gif directory. In order to ensure correct web page formatting, the replacement graphic must have the same dimensions as the original Enterprise Manager graphic. Because the oem.gif image is cached, you must purge your browser's cache to see the new image.

When creating a report definition, two pieces of information that you must specify are the Category and Subcategory. These parameters refer to the published report's location on the reporting website. See Figure 8–3, "Reporting Website Categories

and Subcategories" to see how these parameters graphically correspond to website navigation. This figure shows the Additional Reports page.

Figure 8–3 Reporting Website Categories and Subcategories

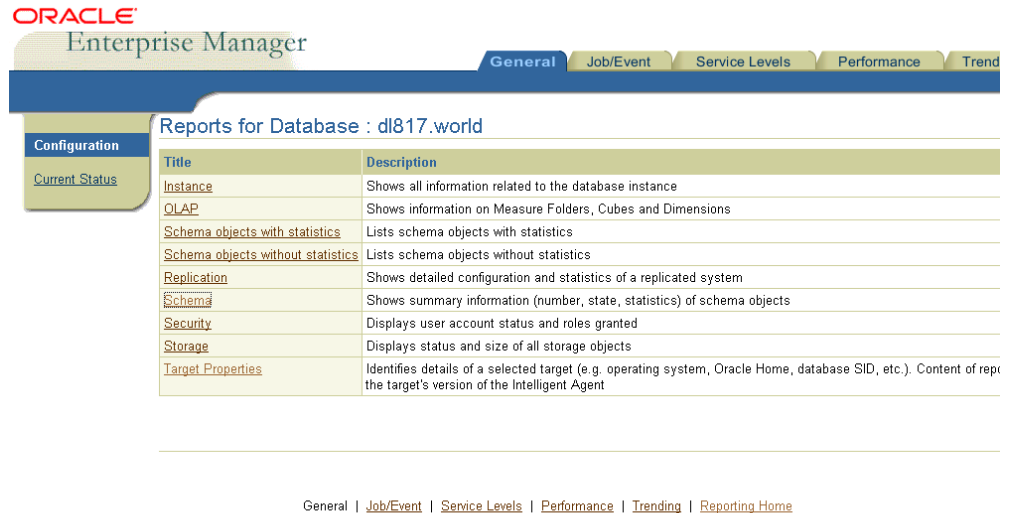
The screenshot shows the Oracle Enterprise Manager Reporting Website interface. At the top left is the Oracle logo and 'Enterprise Manager' text. On the top right is a 'Reporting Home' icon. Below the logo is a navigation bar with 'Category' and three sub-tabs: 'General', 'Job/Event', and 'Setup'. On the left side, there is a sidebar with 'Configuration' and 'Current Status' subcategories. The main content area displays 'Reports for Category : General' with a table listing various report titles and their descriptions. At the bottom, there is a breadcrumb trail: 'General | Job/Event | Setup | Reporting Home'.

Title	Description
Instance	Shows all information related to the database instance
OLAP	Shows information on Measure Folders, Cubes and Dimensions
Replication	Shows detailed configuration and statistics of a replicated system
Schema	Shows summary information (number, state, statistics) of schema objects
Schema objects with statistics	Lists schema objects with statistics
Schema objects without statistics	Lists schema objects without statistics
Security	Displays user account status and roles granted
Storage	Displays status and size of all storage objects

If you are accessing reports from the Console, typically you will see specific reporting web pages related to managed objects selected within the Navigator. The ability to display reports in context from the Console is particularly useful when managing large numbers of targets. For example, if you right-click on a database and choose View Published Reports from the context-sensitive menu, a list of reports pertaining to that database are displayed in your browser. Figure 8–4, "In-context Reports" shows the results of choosing the View Published Reports menu item for a database (dl817).

For a complete site map of the reporting website, see "Navigating the Enterprise Manager Reporting Website" on page 8-27.

Figure 8–4 In-context Reports



Creating a Report from an Existing Report Definition

The following report creation scenario assumes you are the owner of the report definition. In cases where you are not the owner (i.e. you are using one of the predefined reports supplied with Enterprise Manager), you must either click Save As to save the report definition under another definition name, or alternatively right-click on the report definition in the detail view and choose Create Like from the context-sensitive menu.

To create a report using an existing report definition:

1. Select Report Definitions from the Console Navigator. A full list of existing report definitions displays in the detail view.
2. Double-click on the desired report definition. The Edit Report property sheet displays
3. On the Parameters page, click on each of the report elements in the Selected Elements list and enter the requisite parameters, if necessary.

Note: If necessary, you can add or remove report elements on the Elements page.

4. On the Publish page, choose whether you want to generate the report only when viewed from the reporting website or only at a scheduled time.
5. Click View Report to preview the generated report. If necessary, you can further modify the report definition and perform this step again. When you view a report, the your own preferred credentials are used. When publishing a report, the preferred credentials of the report definition owner.
6. After you are satisfied with the results, click OK.

Editing a Report Definition

You edit a report definition by double-clicking on a definition entry in the detail view to display the Edit Report property sheet for that report definition. Alternatively, you can right-click on the definition in the detail view and choose Edit from the context-sensitive menu. Simply change the parameters in the property sheet pages as required. However, you must be the owner of the report definition or have Super Administrator privileges in order to modify it. If you are not the report definition owner (REPORTS_USER for predefined definitions supplied with Enterprise Manager), clicking OK to save the report definition will result in your being prompted to save the report definition under another name via the Save As dialog.

Generating a Report from Enterprise Manager Applications

The Enterprise Manager reporting system is integrated with many applications that are part of the Enterprise Manager framework. As such, certain applications can access the reporting system directly to publish data automatically to the web for use by others. No additional reporting system configuration is required. For example, Performance Manager (part of the Diagnostics Pack), utilizes the reporting system to publish charts and saved analyses to the Performance and Trending pages of the reporting website.

Note: Performance and Trending pages only appear on the reporting website if you have the Enterprise Manager Diagnostics Pack installed. You must first select an individual target. These pages do not appear under “Additional Reports.”

Features that are integrated into the Console Navigator, such as database management functionality, access the reporting system via the predefined report definitions supplied with Enterprise Manager. For example, right-click on a database in the Console Navigator to access the context-sensitive menu and choose View Published Reports (this menu option also appears in the Console Object menu). The reporting website page appears with a list of published reports that can be run against the selected database. Selecting a specific Navigator object, such as Schema, and choosing View Published Reports from the Console's Object menu, displays the Target report, which allows you to view information such as:

- Schema Objects Summary (non SYS and SYSTEM)
- Invalid Objects
- Procedural Object Errors
- Objects (non SYS and SYSTEM) with statistics
- Objects (non SYS and SYSTEM) without statistics
- Tables with disabled primary keys
- User objects in SYSTEM tablespace

The ability to create reports in-context provides a powerful way to monitor the status of managed targets from the Console.

Creating a User-defined Report Definition

In general, the standard predefined report definitions, or their user-customized variants, should provide sufficient reporting capability. However, there may be cases where you may want to create your own report definitions to meet specific monitoring needs. The reporting system provides you with a high degree of flexibility in extracting data and generating a properly formatted report.

Once you have created a custom report definition, it is added to the list of predefined report definitions, at which point you can view a report just as you would with any predefined report definition.

To create a new report definition:

1. Right-click on the Report Definitions object in the Console Navigator to display the context-sensitive menu.
2. Choose Create Report. The Create Report property sheet displays.

3. Enter the requisite information on the various property pages. For a detailed discussion on how to fill out the property sheet, see "The Report Property Sheet" on page 8-12.
4. Click OK to save the report definition.

The Report Property Sheet

The Report property sheet is the primary user interface to the Enterprise Manager reporting system. It is through this property sheet that you define and control the operation of the reporting system in addition to determining the security of the system. Changes made to the report definitions are automatically reflected on the reporting website if the *Publish to Enterprise Manager reporting website* option is checked on the Report General page.

Report General Page

The Report General page allows you to define the primary report identification and report type parameters. See Figure 8-5, "Report General Page".

Figure 8–5 Report General Page
Table 8–1 Report General Page Parameters

Parameter	Description/Usage
Definition Name	Name of the report definition. This name must be unique. This name is used internally (appears only in the Console and not a generated report).
Owner	Name of the report definition creator. This is the fully qualified username of the administrator logged into the Enterprise Manager Console.
Report Title	Title appearing in the header of the generated report and on the reporting website. The title also appears in the report definition multi-column list when Report definitions is selected in the Console Navigator.

Table 8–1 Report General Page Parameters

Parameter	Description/Usage
Report Description	Brief description for the report definition. The description appears on the reporting website.
Category	Specifies the report category in which the report will appear on the reporting website. The primary categories (e.g., General, Custom, Job/Event, Service Levels, Setup) appear as main tab pages on the reporting website. See Figure 8–3, "Reporting Website Categories and Subcategories".
Subcategory	Specifies the subcategory in which the report will appear on the report website. You may select one of the pre-defined subcategories or type in your own. See Figure 8–3, "Reporting Website Categories and Subcategories".
Publish to Enterprise Manager Reporting website	Selected by default (when the webserver is configured), specifies that the report generated from the report definition be published to the Enterprise Manager reporting website. This option must be selected in order to select options on the Report Publish page. Publishing the report allows you to view the report on the reporting website.
Report Type	Specifies the target types to be used for the report.
Create individual reports for all targets of selected type	Allows you to create a report definition that, when it is viewed, interactively prompts an administrator for the information sources. Alternatively, from the Reporting Website, the administrator can select a target and then access the report.
Create a single report for the targets selected below	Allows you to create a report definition using information sources chosen from the Available Targets list.
Available Targets	List of all available targets from which a report can be generated. Entries in this list vary according to the report type selected
Selected Targets	List of information targets from which the report is to be generated. To add entries to this list, select an entry in Available Targets list and click Add. You can remove entries from the Selected Targets list by selecting the desired entry and clicking remove

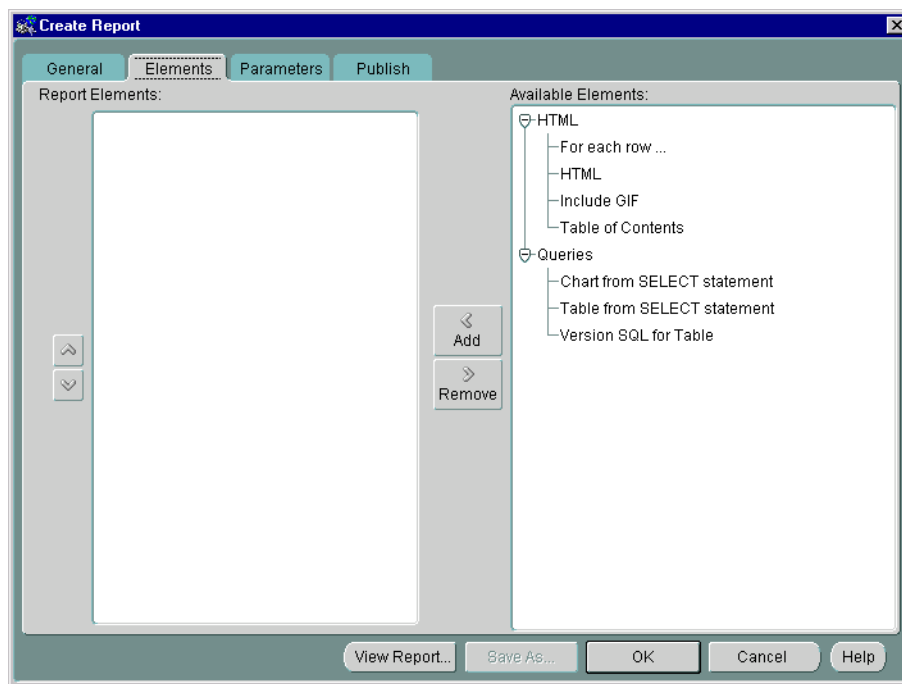
Table 8–1 Report General Page Parameters

Parameter	Description/Usage
View Report	Executes the report definition and displays the generated report in a web browser.
Save As	Saves the current report definition under a different report definition name. This button is disabled when creating a new report definition.
OK	Saves the current report definition to the repository.
Cancel	Close the Report Definition property sheet without implementing any changes.

Report Elements Page

The Elements page allows you to select one or more report elements to be included in the report definition. Report elements correspond to specific types of content to be included in a report. Essentially, you create a report by assembling and ordering elements. Figure 8–6, "Report Elements Page" shows the elements page and Table 8–2, "Report Elements Page Parameters" describes the user interfaces. See Table 8–3, "Report Element Descriptions" for a detailed description of all report definition elements and associated parameters.

Figure 8–6 Report Elements Page



For the predefined report definitions supplied with Enterprise Manager, these elements have already been selected. You need only specify the element parameters.

Table 8–2 Report Elements Page Parameters

Parameter	Description/Usage
Available Elements	<p>Tree list displaying a hierarchy of available report elements. Report elements in this tree list vary according to the Report Type chosen on the General Page. Major report element categories are:</p> <ul style="list-style-type: none"> ■ <i>HTML</i>: Customize the HTML for a generated report and call SQL and Javascript functions for additional row-level processing. ■ <i>Queries</i>: Specify SQL queries to generate charts or tables. ■ <i>Service Levels</i>: Specify time-based statistics for monitored services.

Table 8–2 Report Elements Page Parameters

Parameter	Description/Usage
Report Elements	Elements selected to be included in the generated report.
Add/Remove Buttons	Adds selected report elements from the Available Elements tree list to the Report Elements list or Removes report elements from the Report Elements list. Note: You can also add/remove elements by double-clicking on the specific element.
Up/Down Arrows	Positions a selected element within the Report Elements list. The up/down arrows allow you to determine the order in which the information corresponding to a specific element should appear in the generated report

Table 8–3 Report Element Descriptions

Category	Report Element	Description/Usage
HTML	For each row...	<p>The For each row report element provides you with the highest level of flexibility for your reporting needs; It allows you to generate custom report output without having to develop a new report element. By associating a Javascript procedure with each row of information returned by a SQL SELECT statement, you can perform a multitude of actions to create custom formatting or additional data processing.</p> <p>Parameters:</p> <p>Javascript function to call: Name of the Javascript used to process each row of data returned by the SQL SELECT statement. Javascript functions can be added to a report definition using the HTML report element.</p> <p>Select Statement: SQL SELECT statement needed to return data to be used as input parameters by the Javascript function.</p>
	HTML	<p>The HTML report element allows you to insert any sequence of alphanumeric characters directly into a generated report. This is a flexible report element in that you can use it to perform a variety of functions. For example, you can:</p> <ul style="list-style-type: none"> ■ Insert boilerplate HTML into your report ■ Insert HTML to override default formatting ■ Insert Javascript procedures for use by the For each row report element <p>Parameters:</p> <p>Specify HTML: HTML or Javascript coding to be inserted into the report.</p>

Table 8–3 Report Element Descriptions

Category	Report Element	Description/Usage
	Include GIF	<p>The Include GIF report element allows you to insert a GIF image into a generated report.</p> <p>Parameters:</p> <p><i>Specify Name of GIF file:</i> Name of GIF file. The specified GIF file will be copied from its original location to the appropriate report subdirectory.</p> <p><i>Browse:</i> Displays the Choose File dialog box.</p> <p><i>File Tag:</i> File is copied from the location of origin to the subdirectory of the generated report's index.html file. Specifying a File Tag renames the gif file when it is copied into its respective report directory. Flat directory references from within generated report. Not being used as a true alias. When referencing the "tag" from the HTML report element, you must specify the complete filename with extension.</p> <p>You can also embed an image within a report using the HTML tag ()</p>
	Table of Contents	<p>The Table of Contents report element allows you to insert a hyperlinked table of content at any point within a generated report. This report element uses headers defined in the report to generate an HTML hyperlinked list, each entry taking you directly to the report section associated with its respective header.</p> <p>Parameters:</p> <p>None</p>

Table 8–3 Report Element Descriptions

Category	Report Element	Description/Usage
Queries	Chart from SELECT statement	<p>This report element allows you to create one or more bar, line, and pie charts for a generated report using information obtained from a SQL SELECT statement. The data returned by the SELECT statement must conform to the input parameter requirements for each chart type.</p> <p>Parameters</p> <p><i>SQL:</i> SQL SELECT statement used to return data required to generate the desired chart type. See below for input data requirements.</p> <p><i>Chart Type:</i></p> <p>Bar: This selection generates a typical bar chart and requires the data source to consist of one key column (used for labeling) and one or more numeric data columns. The key column must be the first column in the table.</p> <p>Line: This selection generates a line chart that plots the tabular data against the X-Y axis. As with the bar chart, the data source must consist of a single key column and one or more numeric data columns.</p> <p>Pie: This selection generates a pie chart from a single row of data. If more than one row is returned by the SELECT statement, a new pie chart is generated for each row.</p> <p>Note: For all charts, only the key column can be non-numeric. Non-numeric data appearing anywhere except the first column generates a diagnostic message.</p> <p><i>Orientation:</i> (Applicable to Bar charts only)</p> <p>Horizontal (Default Setting): Produces bar charts with bars originating from the Y-axis.</p> <p>Vertical: Produces bar charts with bars originating from the X-axis.</p> <p>Advanced: Displays the Advanced Options dialog which allows you to change chart dimensions and levels of label cascading (number of label levels along an axis).</p>

Table 8–3 Report Element Descriptions

Category	Report Element	Description/Usage
	Table from SELECT statement	<p>This report element allows you to create one or more tables in a generated report using information obtained from a SQL SELECT statement.</p> <p>Parameters</p> <p><i>SQL:</i> SQL SELECT statement used to return one or more rows of tabular data.</p> <p><i>Orientation:</i></p> <p>Horizontal (Default Setting): Produces a typical row/column format for all tabular data. There is a maximum limit of 100 rows.</p> <p>Vertical: Produces a layout in which a single row table is generated for each column of data found in the information source. The column label appears on the first (leftmost) column with a single row of data in the second column. For each row returned by the SELECT statement, a new table is generated. The order in which the table columns are displayed is determined by the SELECT statement.</p>

Table 8–3 Report Element Descriptions

Category	Report Element	Description/Usage
	Version SQL for Table	<p>Versions of SQL used by older versions of the Oracle Server may not support certain SELECT statement options used to extract information for inclusion into a report. Parameters on this property page allow you to manually define specific SELECT statements to be used with different versions of the database. Defining alternate database version compatible SELECT statements allows you to choose information targets that may contain multiple versions of the database.</p> <p>Parameters:</p> <p><i>Database Version (>=):</i> Lists available SQL (database) versions. A version number selected from this list corresponds to the user-defined SQL statement that is compatible with database versions that are greater than or equal to the specified database version number.</p> <p><i>Add:</i> Displays the Database Version dialog box and allows you to add a new database version number.</p> <p><i>Remove:</i> Deletes the database version number that is selected in the Database Version (>=) list.</p> <p><i>SQL:</i> SQL statement used for the query.</p> <p><i>Database Option Name:</i> Specific options (Replication, OLAP, Generic) used for the query. If Generic is selected, an option check will not be performed for this report element. More options may appear depending on the options installed with the database.</p>
Service Levels	Availability by service	Shows availability information for selected services.
	Low Level Capture of Availability State Changes	Shows sequential service level availability state changes for selected services. This includes messages from the agent regarding outages and restored availability.
	Overall Availability	<p>Displays combined availability of the services as well as the availability of each individual service.</p> <p>Important: All Service Level reports require Up/Down events to be registered on monitored targets.</p>

Table 8–3 Report Element Descriptions

Category	Report Element	Description/Usage
	Downtime Details	Shows downtime details for selected services including percentage downtime, number of times down, total downtime in minutes, and comments/annotations pertaining to downtime.
	Simple Availability Element	Shows the percentages of uptime, downtime and unknown time as well as the number of downtimes and total downtime for selected services.
		<p>Note: All Service Level elements use the same time period parameters.</p> <p>From the Parameters page, you can set the time period for this report definition element. The default period is the current month.</p> <p>It is not necessary to set the time period for each report element (if you have multiple elements). The time period settings chosen for the first report element in a report will automatically apply to all elements in the report.</p> <p>The options are:</p> <ul style="list-style-type: none"> ■ <i>Current</i> - This option is for the current period, either Day, Week, Month, or Year. ■ <i>Previous</i> - This option is for the previous period, either Day, Week, Month, or Year.* ■ <i>Time Span</i> - This option lets you define the length of time for this report: Week, Month, or Year. You can then define the date at which the report is to start by editing the information in the Start Date field. ■ <i>Date Range</i> - This options lets you define a range of dates this report should encompass. In the Start Date field, enter the date that reflects the data you want to study. In the End Date field, enter the date that reflects the data you want to study. <p>Note: All times are in the time zone of the Intelligent Agent.</p> <p>If the current month is November and you select Previous Month as a reporting period, your report will contain data for the month of October. If the current year is 2000 and you select Previous Year as a reporting period, your report will contain data for the year 1999.</p> <p>Note: The reporting system does not distinguish between service downtime resulting from a scheduled target blackout and legitimate User service downtime. See the Intelligent Agent User's Guide for more information on target level blackouts.</p>

Report Parameters Page

The Parameters page allows you to define the information presented by a specific report element. Figure 8–7, "Report Parameters Page" shows the parameters for the *For each row* element, which allows you to define specialized data processing functions for each row of information returned by a SQL SELECT statement.

Figure 8–7 Report Parameters Page

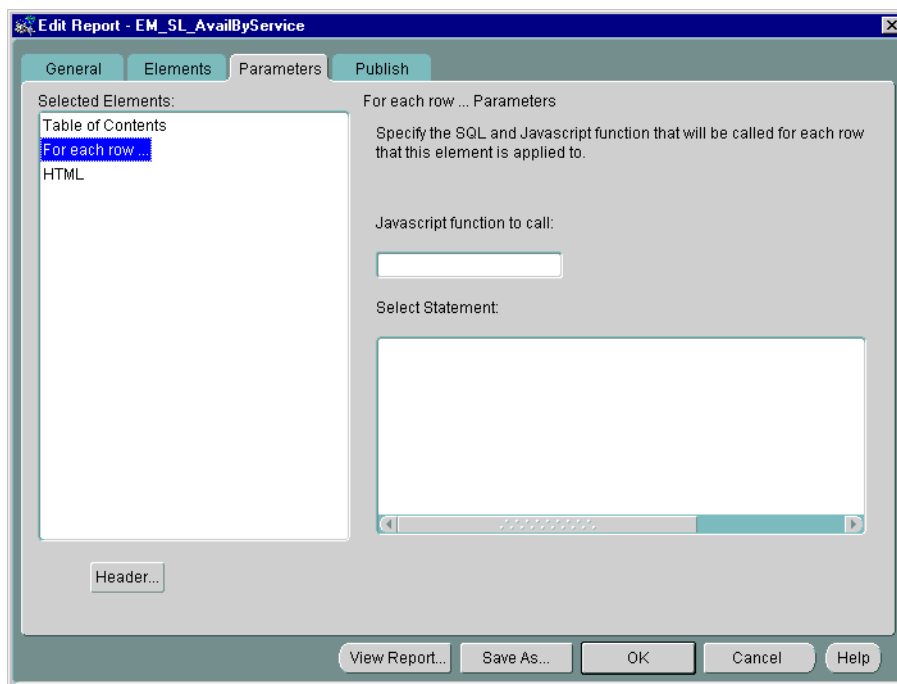


Table 8–4 Report Parameters Page Parameters

Parameter	Description/Usage
Selected Elements	Lists all report elements chosen on the Elements page.
Element Parameters	Displays all modifiable parameters (if any) for the selected report element. See Table 8–3, "Report Element Descriptions" for specific element parameters.

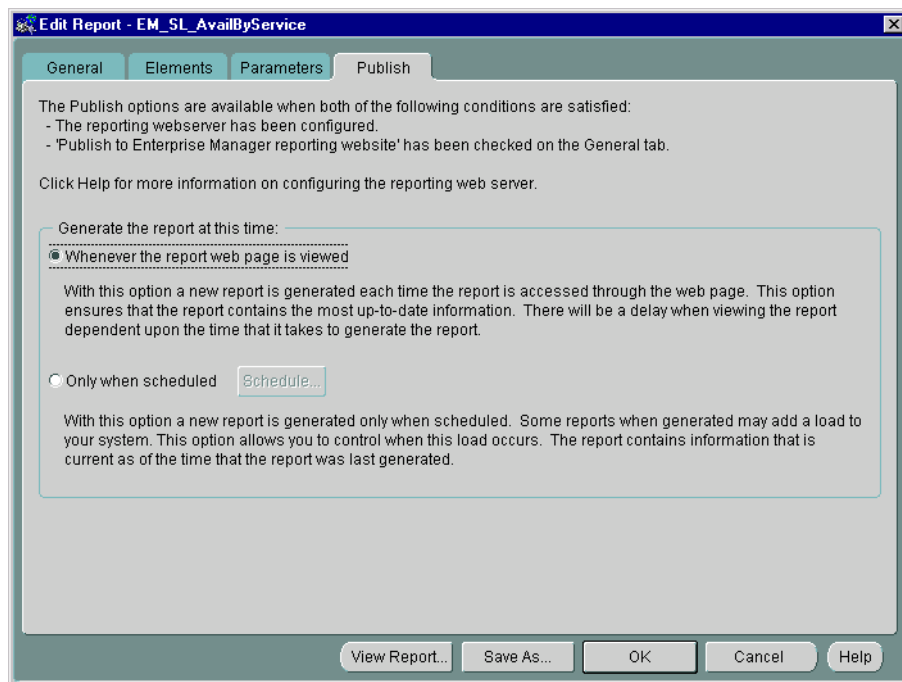
Table 8–4 Report Parameters Page Parameters

Parameter	Description/Usage
Header	Allows you to edit the default header associated with the selected report element. This option allows you to modify the text of the header, its hierarchical level (1 being the highest), or whether a header should be displayed at all in the generated report. Headers are used by the Table of Contents report element to generate an HTML hyperlinked list.

Report Publish Page

The Publish page allows you to publish generated reports to the Enterprise Manager Reporting website. Options on this page are only available when *Publish to Enterprise Manager reporting website* is selected on the General page. See Figure 8–8, "Report Publish Page".

Figure 8–8 Report Publish Page



This page allows you to select when the report should be generated: on-demand or according to a set schedule. Two options are available:

Whenever the report web page is viewed: New reports are generated whenever the report link on the reporting website is accessed from a web browser. Hence, the report is always current in that it presents real-time information.

Only when scheduled: Reports are generated based on a set schedule. Hence, report content does not present real-time information. To specify the schedule click the Schedule button to launch the Create Jobs property sheet. You can schedule report generation just as you would a normal Enterprise Manager job, however, the job task used to schedule a report cannot be accessed outside the reporting system.

Note: You must choose a target running a 9i version of the Intelligent Agent in order to schedule jobs. See "Job Schedule Page" on page 5-18 for more information about scheduling jobs.

Navigating the Enterprise Manager Reporting Website

The following tables provide a complete listings of all predefined reports available on the Enterprise Manager reporting website and in the Console Report Definition detail view. Availability of reports may vary depending on the options installed in your enterprise environment.

You can view the reporting website using one of the following methods:

From a browser:

1. Start a browser.
2. Enter the following URL: `http://<ReportingWebserverName>;port#/em/OEMNavigationServlet`. The default port number used is 3339.

From the Enterprise Manager Console:

1. Start the Enterprise Manager Console.
2. Choose View Published reports from the Console's Object menu.

Both methods assume that you have already configured your reporting environment.

Table 8–5 Database Reports

Category	Subcategory	Report Title	Description
General	Configuration	Instance	Shows all information related to the database instance.
		OLAP	Shows information on Measure Folders, Cubes and Dimensions
		Schema objects with statistics	Lists schema objects with statistics
		Schema objects without statistics	Lists schema objects without statistics
		Replication	Shows detailed configuration and statistics of a replicated system
		Schema	Shows summary information (number, and state) of schema objects
		Security	Displays user account status and roles granted
		Storage	Displays status and size of all storage objects

Table 8–5 Database Reports

Category	Subcategory	Report Title	Description
		Target Properties	Identifies details of a selected target (e.g. operating system, Oracle Home, database SID, etc.) Content of the report depends on the target's version of the Intelligent Agent.
	Current Status	Database Object Space Usage	Shows space usage reports for: <ul style="list-style-type: none"> ■ Objects unable to extend ■ List of Objects which are nearing Max Extents ■ Overextended segments (non SYS and SYSTEM)
		Database Top 10	For a selected database, shows the top 10: <ul style="list-style-type: none"> ■ SQL statements that have been executed. ■ Accessed tables. ■ Procedures that have been executed.
		Disk Space Used by Tables	Displays the disk space used by tables (for SYS and non-SYS).
		Instance	Displays instance statistics and process state
		Storage	Shows extent, segment and I/O information
Job/Event	Events	Outstanding Alerts	Shows information on outstanding alerts (with status of critical, warning, unknown, or error).
		Alert History	Shows information on alerts for a target that has moved to the event history.
		Registered Events	Lists all registered events for a target.
	Jobs	Active Jobs	Displays details for jobs scheduled on a target.
		Failed Jobs from the last 24 hours	Lists jobs for a target that failed in the last 24 hours.

Table 8–5 Database Reports

Category	Subcategory	Report Title	Description
		Failed Jobs from the Last 7 Days	Lists jobs for a target that failed in the last 7 days.
		Completed Jobs from the Last 24 Hours	Lists jobs for a target that completed in the last 24 hours.
		Completed Jobs from the Last 7 Days	Lists jobs for a target that completed in the last 7 days.
		Average Execution Time per Job	Shows information on execution times for jobs completed against a target.
Service Levels	Summary	Overall Availability	Displays combined availability of the services as well as the availability of each individual service.
		Availability by Service	Shows availability information for selected services.
	Details	Downtime Details	Shows downtime details for selected services including percentage downtime, number of times down, total downtime in minutes, and comments/annotations pertaining to downtime.
		Service Level Availability Timeline	Displays sequential service availability state changes.
	Diagnostics	Availability Diagnostics	Shows captured low-level service availability data.

Table 8–6 Additional Reports

Category	Subcategory	Report Title	Description
Job/Event	Events	Outstanding Alerts Sorted by Target	Displays details, sorted by target name, on all outstanding alerts with status of critical, warning, unknown, or error. Note: If there are no alerts, the Report Title “No Outstanding Alerts” is displayed.
		Outstanding Alerts Sorted by Event	Displays details, sorted by event name, on all outstanding alerts with status of critical, warning, unknown, or error.
		Alert History Sorted by Target	Shows information, sorted by target name, on alerts that have been moved to the event history.
		Outstanding Alerts in History Sorted by Event	Shows information, sorted by event name, on alerts that have been moved to the event history.
		Registered Events Sorted by Target	Provides information, sorted by target for all registered events
		Registered Events Sorted by Event	Provides information, sorted by event name, for all registered events.
		Events in Library	Displays details on all events saved to the event library.
		Notifications for Events Sorted by Administrator	Lists all paging and email notifications, sorted by administrator name, that were sent due to event status changes.
		Notifications Sorted by Event	Lists all paging and email notifications, sorted by event name, that were sent due to event status.
		Jobs	
Active Jobs Sorted by Job	Provides information, sorted by job name, for all jobs scheduled.		
Failed Jobs from the Last 24 Hours	Lists all jobs that failed in the last 24 hours.		
Failed Jobs from the Last 7 Days	Lists all jobs that failed in the last 7 days		
Completed Jobs from the Last 24 Hours	Lists all jobs that completed in the last 24 hours.		

Table 8–6 Additional Reports

Category	Subcategory	Report Title	Description
		Completed Jobs from the Last 7 Days	Lists all jobs that completed in the last 7 days.
		Jobs in Library	Displays details on all jobs saved to the job library.
		Notifications for Jobs Sorted by Administrator	Lists paging and email notifications, sorted by administrator name, that were sent due to job status changes.
		Notifications Sorted by Job	Lists paging and email notifications, sorted by job name, that were sent due to job status changes.
		Active Jobs	Shows information on execution times for all completed jobs.
Setup	Administrators	Administrator Overview	Provides information on all Enterprise Manager administrator accounts.
	Agents	Intelligent Agent Overview	Displays status and other details on discovered Intelligent Agents.
	Collections	Active Collections	Shows all active, operational collections defined in Oracle Capacity Planner.
	Targets	All Targets	Lists all targets that have been manually or automatically discovered.

Enterprise Security Management

This chapter describes the component of Oracle Enterprise Manager used to administer Enterprise User Security for the Advanced Security Option. The chapter explains use of Enterprise Manager within a simple scenario in which an Oracle Internet Directory Server is used as the central repository for users in a large organization. It contains the following sections

- Overview of Enterprise Security Manager
- Introduction to Directory Servers
- Entailing and Configuring Your Enterprise Security Environment
- Administering Users
- Administering Oracle Contexts
- Command Line Tool

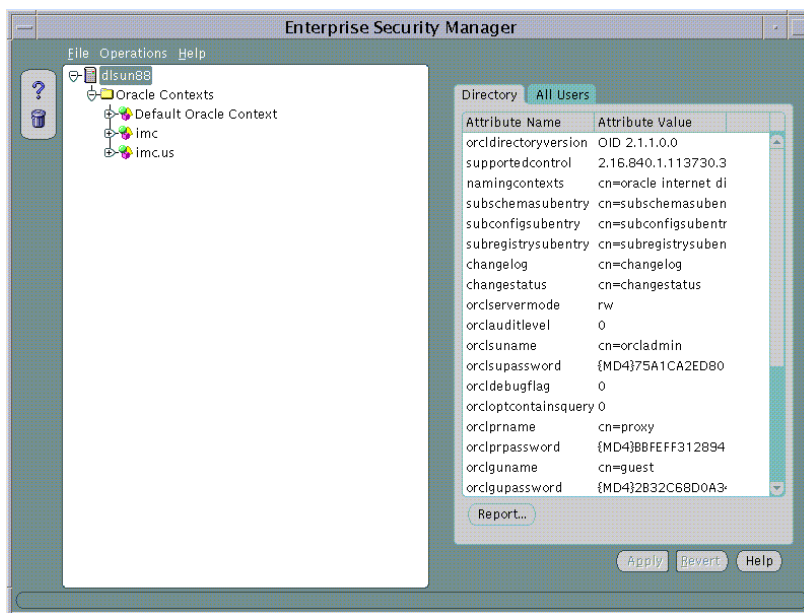
Overview of Enterprise Security Manager

Oracle Enterprise Security Manager provides an easy-to-use graphical interface to administer enterprise user security and access control for large numbers of databases in your enterprise environment through the Oracle Internet Directory server. You use Oracle Enterprise Security Manager to perform the following tasks:

- Manage Database Security Under Oracle Contexts in a Directory
- Manage Users in a Directory

You start this Enterprise Security Manager via the MS Windows Start menu, or by issuing the "esm" command on a UNIX command line. Upon logging in, Enterprise Security Manager appears as shown in Figure 9–1, "Enterprise Security Manager", given that the Directory contains at least the Oracle9i Default Oracle Context.

Figure 9–1 Enterprise Security Manager



Enterprise Security Manager manages one Directory Server, identified at the top of the main application tree. It has a series of menu operations that apply to this Directory Server.

Users are managed in the Directory using Enterprise Security Manager. The application shows the directory to which it is connected and allows you to add, delete and browse Users in that Directory. Enterprise Security Manager may also be used to manage Oracle Contexts in the Directory. An Oracle Context is an area of structured information in the Directory recognizable to Oracle8i and Oracle9i products as well as an administrator hierarchy for management of the data in Oracle Contexts for different Oracle product areas.

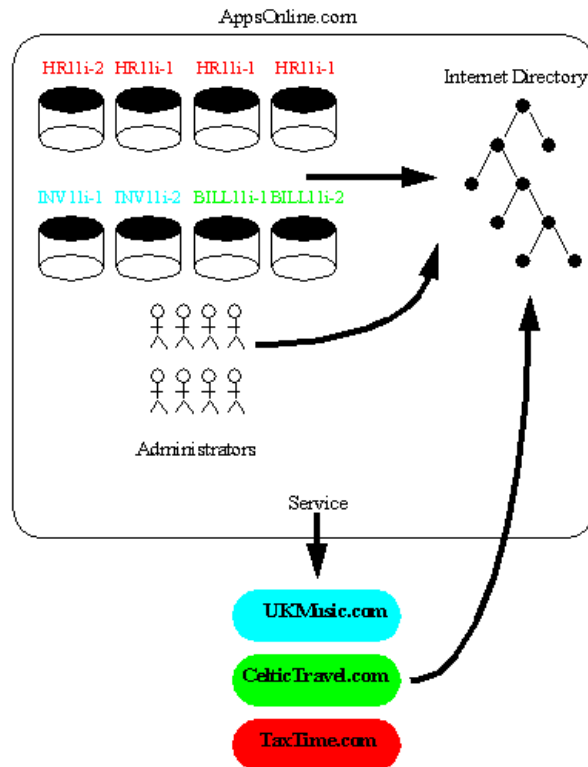
This chapter is presented in two parts; Administering Users and Administering Oracle Contexts. It will use the example the “AppsOnline” Application Service Provider to illustrate both facets of Enterprise User Security management.

Introduction to Directory Servers

A Directory Server may be used a general purpose means to centralize definitions of user and server access information over an entire network. As well as storing naming information, the Directory may be employed to centralize password definitions, digital certificates and application authorizations for the users that it defines. This is possible, in the particular case of Oracle Internet Directory, as it allows for secured access and modification of sensitive information held in the Directory such as passwords or application authorizations.

This chapter shall use as its example an Application Service Provider called, “AppsOnline”. AppsOnline has a large set of Oracle9i Databases that it uses to host different types of Application Software for its customers. AppsOnline needs to manage administrative access to these databases for its IT staff.

Figure 9–2 AppsOnline Hierarchy



AppsOnline maintains Oracle9i databases upon which are hosted three types of Application for its customers; Human Resources, Inventory and Billing. One customer, “TaxTime.com” subscribes to AppsOnline for its Human Resources Applications. A second customer, “CelticTravel.com” subscribes to the company for its Billing Applications. A third company, “UKMusic.com” subscribes to the company for its Inventory Management Applications.

AppsOnline dedicates some of its databases to each customer and manages these databases on behalf of the customer. The company has used an Oracle Internet Directory to hold information about their own employees, the databases on which they host Applications, and the customers for whom they provide a service. In the course of their business, they may wish to manage administrative access to their databases by their IT employees and manage access rights to information in these databases based upon each type of customer Application that they support.

This chapter will illustrate how Oracle Enterprise Manager may be used in this example scenario.

Entailing and Configuring Your Enterprise Security Environment

Task1: Configure an Oracle Internet Directory.

Task2: Install Oracle Enterprise manager

Task3: Configure Oracle Enterprise Manager for Enterprise User Security

Task4: Start Oracle Enterprise Security Manager

Task5: Log On To the Directory

Task 1: Configure an Oracle Internet Directory

Oracle9i Enterprise User Security is based wholly around an Oracle Internet Directory. The Directory Server must be properly installed and configured before Enterprise Manager may be used to manage Enterprise User Security. The following stages of Oracle Internet Directory configuration must be complete before proceeding

1. Either an Oracle8i or Oracle9i Internet Directory is installed, running and accessible over both standard LDAP and Secure Sockets Layer enabled LDAP (LDAP/SSL). For more information please refer to the *Oracle Internet Directory Administrators Guide*.
2. The Oracle Internet Directory has been configured to support Oracle9i Directory Schema Objects and contains an Oracle9i Default Oracle Context. In the case of a version 9i Oracle Internet Directory these requirements may already be in place. However, the Oracle9i Directory Schema Objects and Default Oracle Context may be configured on the Directory Server using The Oracle Net Configuration Assistant. For more information please refer to the Oracle Net Configuration Assistant Administrators Guide.

Task 2: Install Oracle Enterprise Manager

Oracle Enterprise Manager is automatically installed with the Oracle9i Enterprise Edition Server Install and includes all necessary functionality for Enterprise User Security. Oracle Enterprise Manager is also installed by default with the Oracle9i Infrastructure Install at the same time as Oracle Internet Directory. Oracle Enterprise Manager may also be installed separately in its own ORACLE_HOME using the custom install option.

Task 3: Configure Oracle Enterprise Manager for Enterprise User Security

Oracle Enterprise Manager may be used to manage Enterprise User Security in two modes of operation. The Oracle9i Enterprise Manager Console may be used to connect to the Oracle9i Management Server (OMS) and discover a Directory Server to manage. Alternatively, a dedicated application called, “Enterprise Security Manager” may be launched from the same ORACLE_HOME as Enterprise Manager and used to connect directly to the Directory Server. In either mode of operation functionality is identical. Only the latter mode, using the Enterprise Security Manager application, will be used in this chapter.

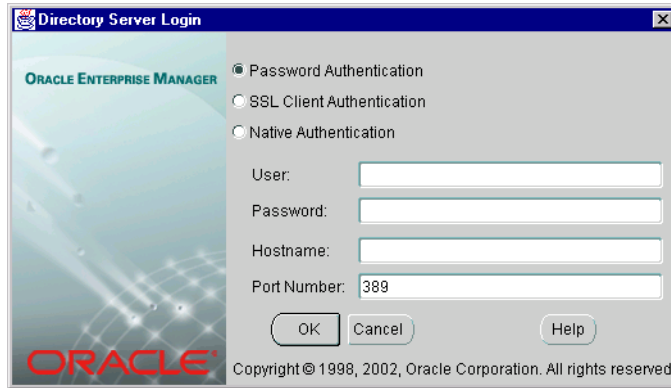
Enterprise Security Manager does not require any special configuration for it to run. However all Oracle Databases in the enterprise that need to avail of Enterprise User Security should be accessible over Oracle Net from the Enterprise Manager ORACLE_HOME.

Task 4: Start Oracle Enterprise Security Manager

To launch Enterprise Security Manager from the Enterprise Manager ORACLE_HOME, enter the following at the command line:

```
> esm
```

This will cause the Directory Log On box to appear

Figure 9–3 Directory Login Dialog

Task 5: Log On To the Directory

Enterprise Security Manager offers three ways to connect to a Directory Server by selecting the appropriate option in the Log On Box. These options are listed in the table below

Table 9–1 Directory Connection Methods

Authentication Type	Description
Password Authentication	Uses Simple Authentication requiring a distinguished name or a known directory nickname and a password
SSL Client Authentication	Uses two-way SSL Authentication in which both the client and server use Oracle Wallets containing digital certificates. The subsequent connection will then be encrypted.
Native Authentication	Applies only to Microsoft Windows NT or Windows 2000 and uses Operating System level authentication to log on to a Microsoft Active Directory

For example, Password Authentication may be selected when using the *orcladmin* Oracle Internet Directory super user name and password to log on.

Administering Users

Enterprise Security Manager may be used to Create Users in the Directory. This is done by selecting “Create Enterprise User...” from the Operations Menu.

Figure 9–4 Operations Menu



The Create User Window will appear in which to enter the name and location of the new User in the Directory.

Oracle Wallets

Oracle wallets are data structures that contain a user private key, a user certificate, and a set of trust points (the list of root certificates the user trusts). Enterprise Security Manager functionality pertaining to Oracle Wallets will only appear in the Create User or Edit User screens when running in an ORACLE_HOME that has been configured for this purpose to use Oracle PKI Products. First, you must generate a Certificate Signing Authority for Enterprise Security Manager. This is done by running "esm -genca" on the command line. The following example displays the expected output from running this utility.

```
> esm -genca
Generating CA Private Key. Please Wait..
```

```
Enter a Wallet Administrator Password to protect access
to your CA private key: test_password
```

```
A CA has been created for Enterprise Security Manager.
You must remember your Wallet Administrator Password. It is required
by Enterprise Security Manager to generate new Oracle Wallets.
```

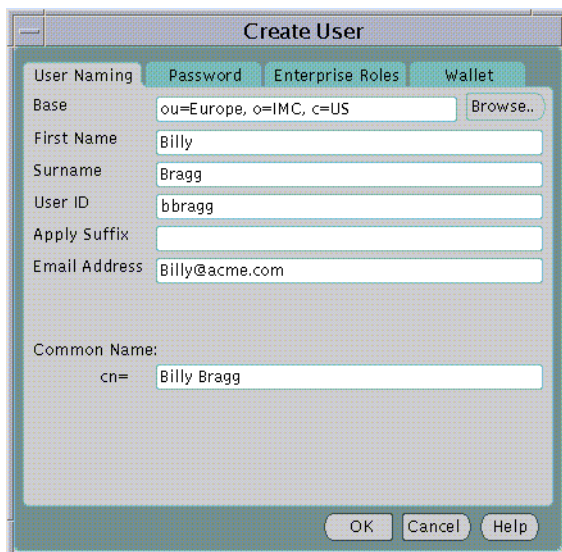
Note: When you use `esm -genca` to generate a new CA for Enterprise Security Manager, a default identity will be created for your CA using 'ORACLE' as the value for all X500 name components of the CA certificate.

You may define your own individual values to be used for the CA identity by editing the `ORACLE_HOME/sysman/admin/esmca.properties` file.

You must run `esm -genca` again after modifying this file.

Specifying a new User Name

Figure 9–5 Create User Property Sheet: User Naming Page



The screenshot shows a 'Create User' dialog box with four tabs: 'User Naming', 'Password', 'Enterprise Roles', and 'Wallet'. The 'User Naming' tab is active. The fields are as follows:

Field	Value
Base	ou=Europe, o=IMC, c=US
First Name	Billy
Surname	Bragg
User ID	bbragg
Apply Suffix	
Email Address	Billy@acme.com
Common Name: cn=	Billy Bragg

Buttons at the bottom: OK, Cancel, Help.

The following fields are mandatory for creation of a new User in the Directory:

Table 9–2 Create User Property Sheet: User Naming Page Mandatory Fields

Field	Description
Base	The entry point in the Directory at which the new User will be created
First name	First half (Christian Name) of the new User's full name
Surname	Second half (Surname) of the new User's full name
User ID	The Logon identifier that the user may use to access databases and applications

The following additional fields are not mandatory for creation of a new User in the Directory but may be recorded for the new User if desired.

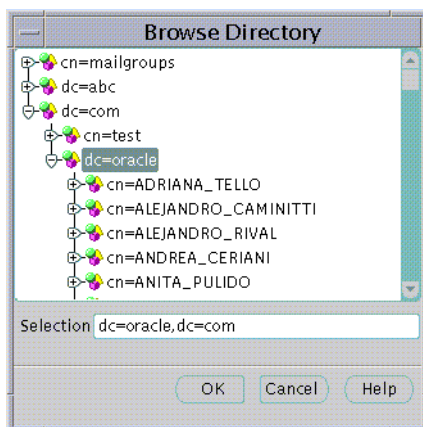
Table 9–3 Create User Dialog: Non-mandatory Fields

Field	Description
Apply Suffix	This is the current value of any common user ID suffix that is always applied to the end of the User ID for a new User. For example, <User ID>.us.acme.com
Email Address	The email address to record in the Directory for the new User, if desired.
cn=	This is the Common Name component (cn=) of the Distinguished Name of the new User in the Directory. By default it is set to the full name of the new User, however you can override the value if you wish to force a particular value for the "cn=" portion of the User's Distinguished Name.

Specifying a Directory Base

All Users in the Directory must exist at a particular "Base" within the Directory. The Base can be any existing Directory Entry such as Country Entry (e.g "c=US") or an Organization Entry (e.g "o=Acme, c=US". Many Users would typically share the same Base. This Base identifies all the Users contained under it as belonging to the same high level organization.

The Base at which to create a new User can be entered in the Base field in the Create User screen. However, you may explore the entire Directory to choose a suitable Base by clicking on the Browse... button. The Browse Directory dialog will appear.

Figure 9–6 Browse Directory Dialog

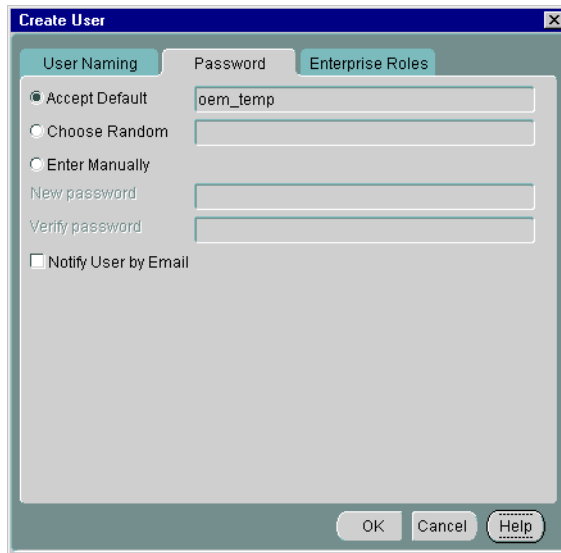
The Browse Directory screen lets you navigate the directory by drilling down into each entry from the top of the Directory Tree. When a Directory Entry is selected its Distinguished Name is placed in the Selection field. To accept the selected Distinguished Name choose the OK button. This value will then be returned as the selected Base for a new Directory User.

Note: This value will be preserved for all subsequent operations that create or search for Users in the Directory. However you may change it as many times as you like.

Specifying a new User Password

The second Tab Panel of the New User screen allows you to set an initial password for the new User in the Directory. This will be the new User's initial password for:

- Directory log on
- Database log on to databases that support Password Authenticated Global Users
- A new Oracle Wallet, if created for the new User at this time.

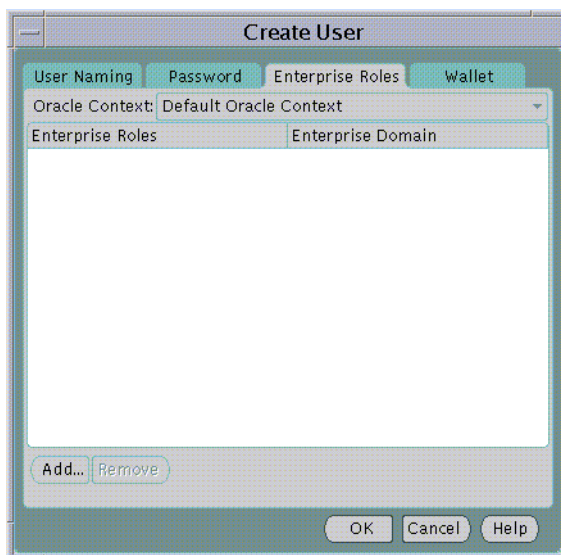
Figure 9–7 Create User Property Sheet: Password PageThe image shows a screenshot of the 'Create User' dialog box, specifically the 'Password' tab. The dialog has a title bar with 'Create User' and a close button. Below the title bar are three tabs: 'User Naming', 'Password', and 'Enterprise Roles'. The 'Password' tab is active. It contains three radio buttons: 'Accept Default' (selected), 'Choose Random', and 'Enter Manually'. Below these is a text field containing 'oem_temp'. Underneath are two more text fields labeled 'New password' and 'Verify password'. At the bottom left is a checkbox labeled 'Notify User by Email'. At the bottom right are three buttons: 'OK', 'Cancel', and 'Help'.

When Entering a password you may choose to accept a default first time password for the new User or manually enter the first time password for the new User. In either case, the new User must change their own password immediately after its first use.

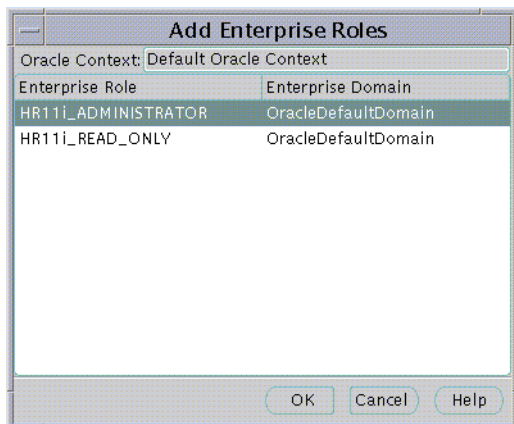
Specifying an Initial Enterprise Role Assignment

Enterprise Roles are discussed later in this Chapter. At the time of User creation you may select any previously configured Enterprise Roles and grant them to the new User.

Figure 9–8 Create User Property Sheet: Enterprise Roles



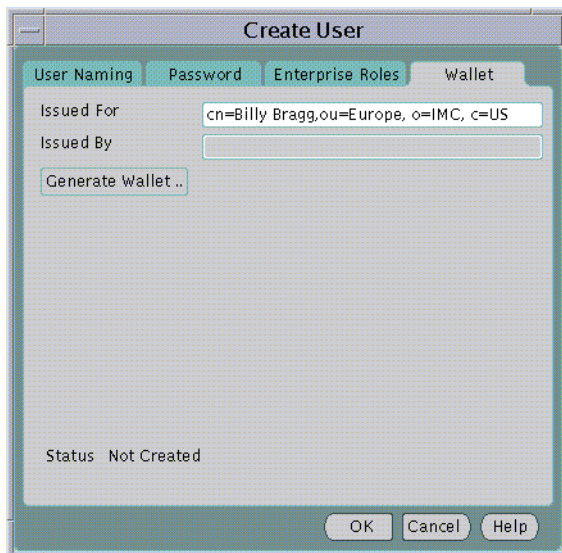
To select one or more Enterprise Roles to grant to the new User at this time choose **Add...** in the Enterprise Roles page of the Create User screen. The Add Enterprise Roles Page will appear from which you can choose any Enterprise Roles in your Oracle Context to assign to the new User.

Figure 9–9 Add Enterprise Roles Dialog

Specifying an Oracle Wallet

An Oracle Wallet containing a new Digital Certificate, Private Key and Certificate Trustpoints may be generated for the new User in an encrypted binary format. The Oracle Wallet will be stored with the new User in the Directory Server as part of the Directory Entry for the User.

Note: This functionality is only available AFTER you have run the `esm -genca` command in your environment

Figure 9–10 Create User Property Sheet: Wallet

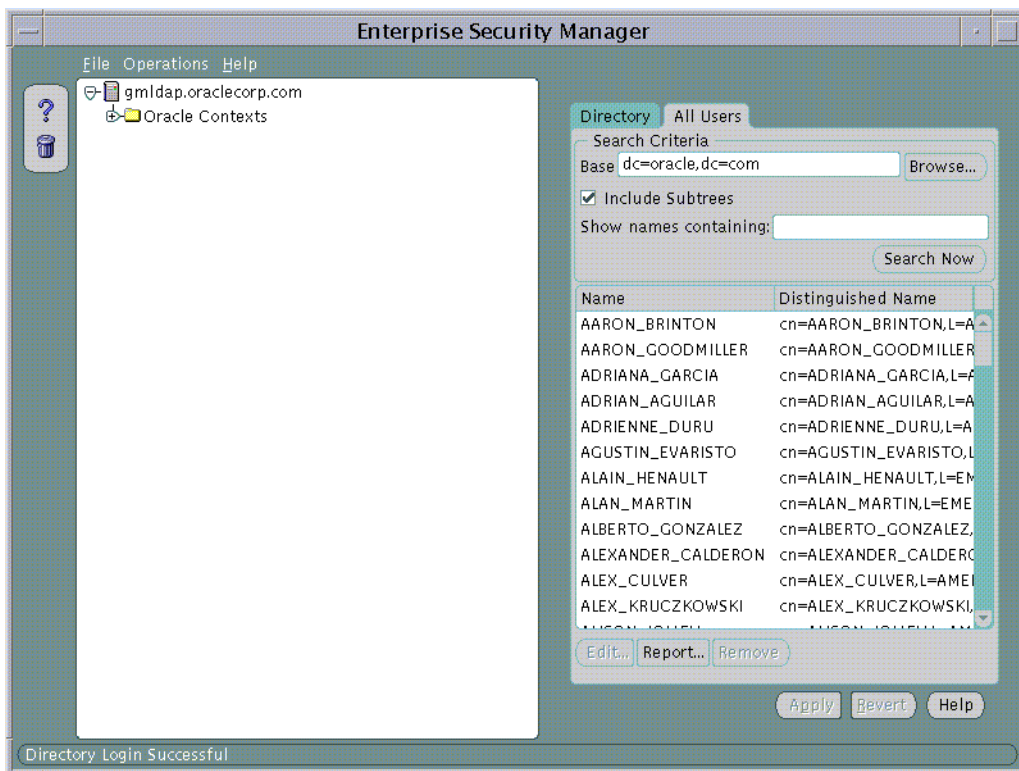
The Distinguished Name under which the new User will be created is used by default as the Distinguished Name for the Digital Certificate to be contained in the new User's Oracle Wallet. It is always good practice to let the Distinguished Names of User Certificates correspond to their Distinguished Names in the Directory. However, you may edit the Distinguished Name to be used for the Certificate before generating the Wallet by editing the contents of the Issued For: field.

An Oracle Wallet will be created when you click on the Generate Wallet... button.

Browsing Users in the Directory

Enterprise Security Manager allows you to browse all Users that are currently stored in the Directory. This is done by selecting the All Users page from the Directory at the top of the main application tree

Figure 9–11 All Users Page



To Search for one or more users the directory, the Search Criteria must be set and the Search Now button used to perform a new search for Users based upon the given Search Criteria. The All Users page will refresh to show the results of this search. There are three factors to User Search Criteria:

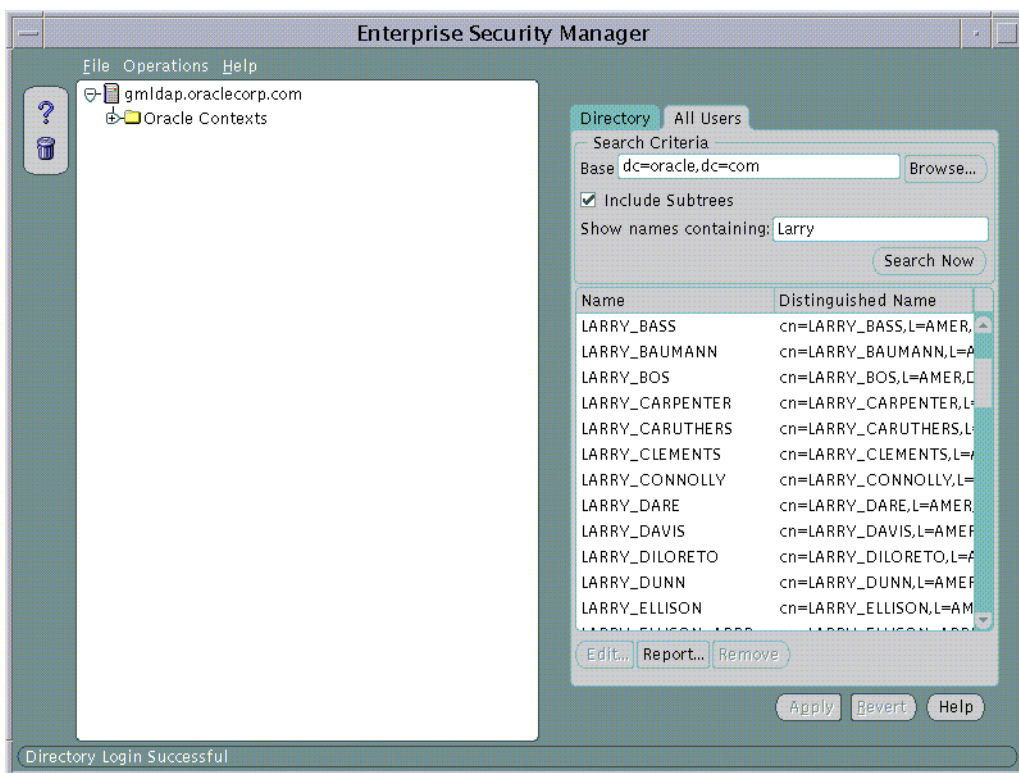
Table 9–4

Search Criteria	Affect on the Search
Base	This is the Base Entry in the Directory at which the search will be performed. Any Users returned in the search will exist under this Base in the Directory.

Table 9–4

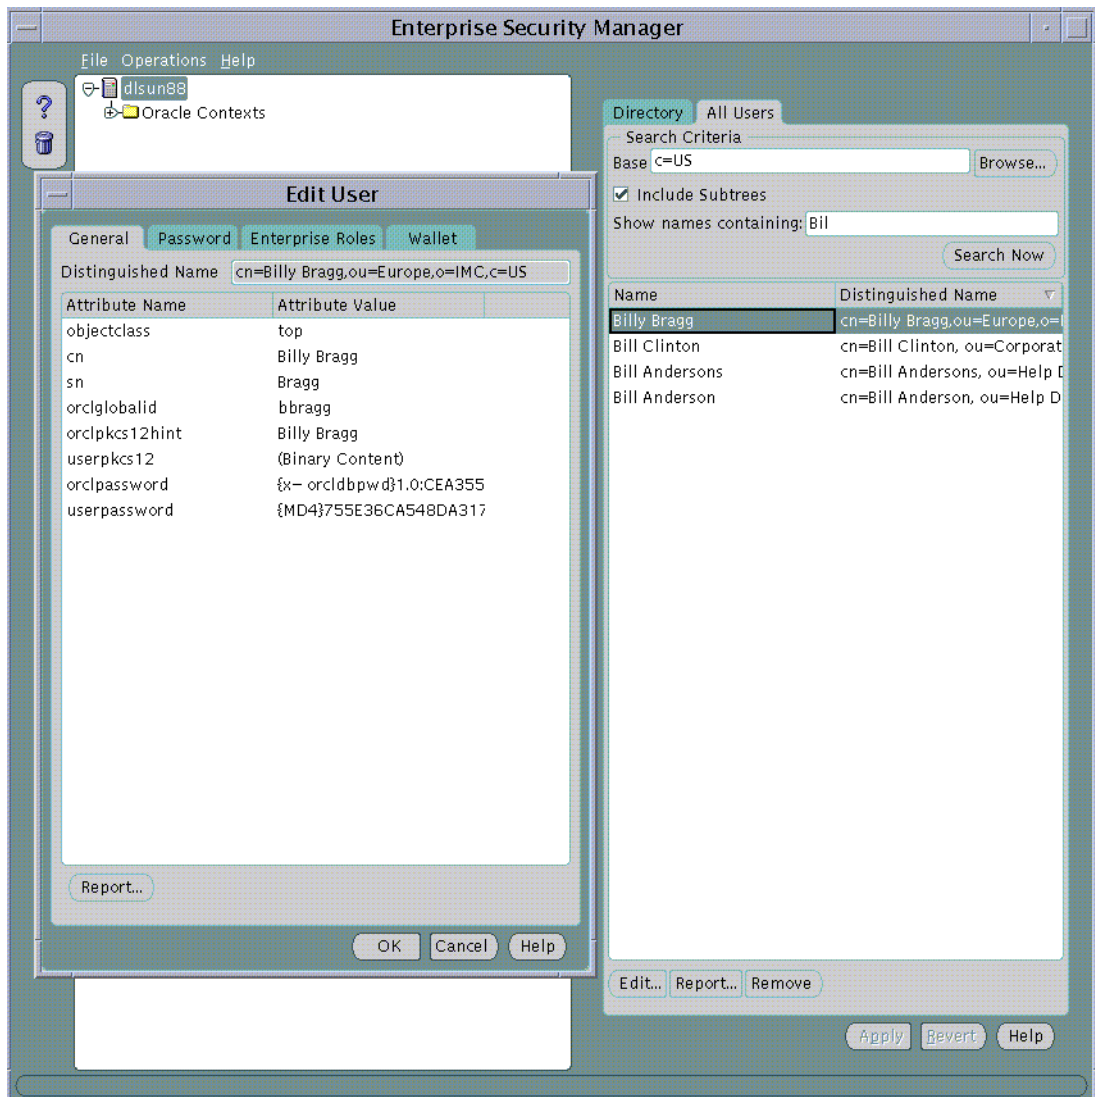
Search Criteria	Affect on the Search
Include Subtrees	This determines whether to show all Users in the Directory anywhere under the selected Base or to only show those Users that exist specifically at that Base location in the Directory.
Show Names Containing	This limits the entire search to contain only those Users whose Directory Entries have a Common Name that starts with a specified pattern. This is useful if the exact name or Base of the desired User is not known.

For example, the Search Criteria may be set to search this Directory for a User given only that the Base is dc=oracle, dc=com and the first name is “Larry”

Figure 9–12 Base Search Criteria

After searching for Users in the Directory, any one user can be chosen from the list and edited. This is achieved either by selecting the User from the list in the All Users page and choosing the Edit... button or by double clicking on that User in the list.

Figure 9–13 Editing a User



When a User in the Directory is selected for Edit, its password, Enterprise Role assignments and Oracle Wallet can be modified in the same way as discussed during creation of a new User in the Directory.

Administering Oracle Contexts

An Oracle Context is a top level Entry in the Directory underneath which is contained the data used by any Directory aware Oracle product. Enterprise Security Manager allows you to manage database and security related information in the Directory under an Oracle Context.

Note: Users do not need to be contained in the Directory within an Oracle Context. It is assumed that the Directory may define its Users for a wide variety of purposes. Oracle does not require that Users in a Directory to be created within an Oracle Context though it is still possible to do so.

Oracle Context Versions

An Oracle Context in the Directory may either be a version *8i* or version *9i* Oracle Context. For Enterprise User Security there is some functionality that can only be managed using a *9i* Oracle Context, for example, “Password Authenticated Global Users”. Enterprise Manager for Oracle *9i* may be used to manage version *9i* Oracle Contexts as well version *8i* Oracle Contexts in the Directory.

Oracle Enterprise Security Manager displays in its main application tree all the Oracle Contexts that exist in the Directory Server. It will display both version *9i* and version *8i* Oracle Contexts, should they exist. In the example below Enterprise Security Manager is connected to an Oracle Internet Directory that has been configured to support the Oracle *9i* Directory Schema and an Oracle *9i* Default Oracle Context.

Specifying Properties of an Oracle Context

An Oracle Context has a number of general properties that can be viewed and managed in the General page when an Oracle Context is selected on the tree:

Figure 9–14 Viewing an Oracle Context Properties

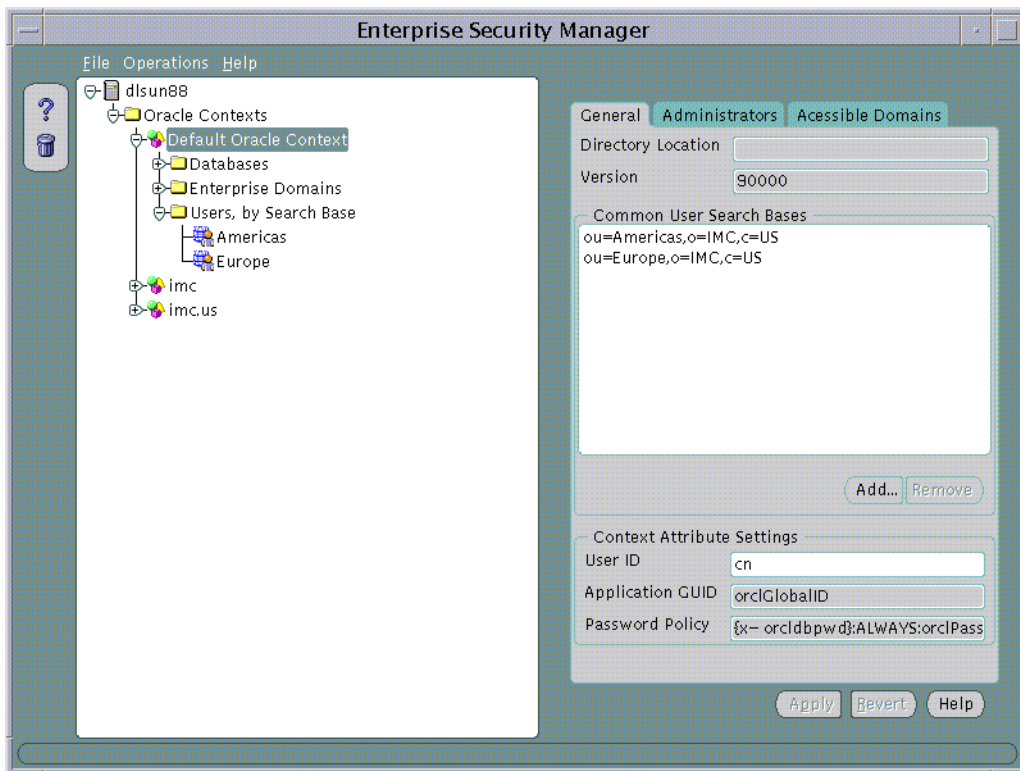


Table 9–5 Context Property Description

Property	Description
Directory Location	This is the Directory Base of the Oracle Context. In the case of the Default Oracle Context this value is empty as the Directory Base is the root of the Directory tree
Version	This identifies whether the Oracle Context supports 8i or 9i functionality

Table 9–5 Context Property Description

Property	Description
Common User Search Bases	This is the list of Base locations in the Directory at which Users may commonly exist. Identifying a list of User Search Bases allows you to quickly browse the users at those Directory Locations and also indicates to 9i Databases in the Oracle Context where they may find Directory Users that connect to them.
User ID	This is the name of the Attribute in a User Entry that determines the value of that Users's User ID. User Entries have many different attributes. This setting controls the User ID with which Users can authenticate to Oracle9i databases, Directory Servers or Directory enabled Applications. Its default value is, "cn", the Common Name of the Directory User.
Application GUID	This is the name of the Attribute in a User Entry in which unique Application GUID values will exist. It cannot be modified in this release
Password Policy	This is the Password Policy syntax used by Oracle9i database when authenticating Password Authenticated Global Users. It cannot be modified in this release.

Specifying User Search Bases

User Search Bases can be added to or removed from a version 9i Oracle Context using the Oracle Context General page.

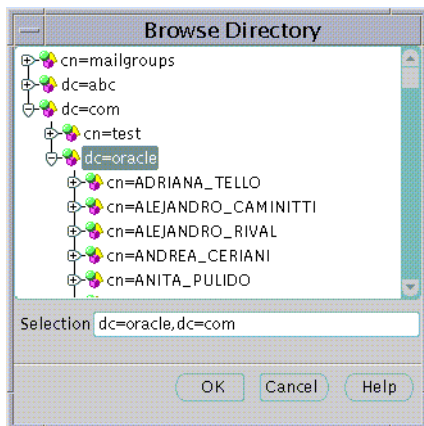
Note: This functionality is not available in version 8i Oracle Contexts.

To remove a User Search Base from the Oracle Context:

1. Select a Search Base in the Common User Search Bases List and choose Remove... The Search Base will be removed from the List.
2. Choose Apply; the User Search Base will be removed from the Oracle Context in the Directory

To add a new User Search Base to an Oracle Context:

1. Choose Add... The Common User Search Bases screen will appear.

Figure 9–15 User Search Base Dialog

2. Navigate the Directory to select a desired Directory Entry as a User Search Base. You may also edit the contents of the Selection field in this screen to manually define the User Search Base.
3. Choose OK in the Common User Search Bases screen. The selected Entry will be added to the list of User Search Bases in the Oracle Context General Page.
4. Choose Apply; the User Search Base will be added to the Oracle Context in the Directory

Specifying Oracle Context Administrators

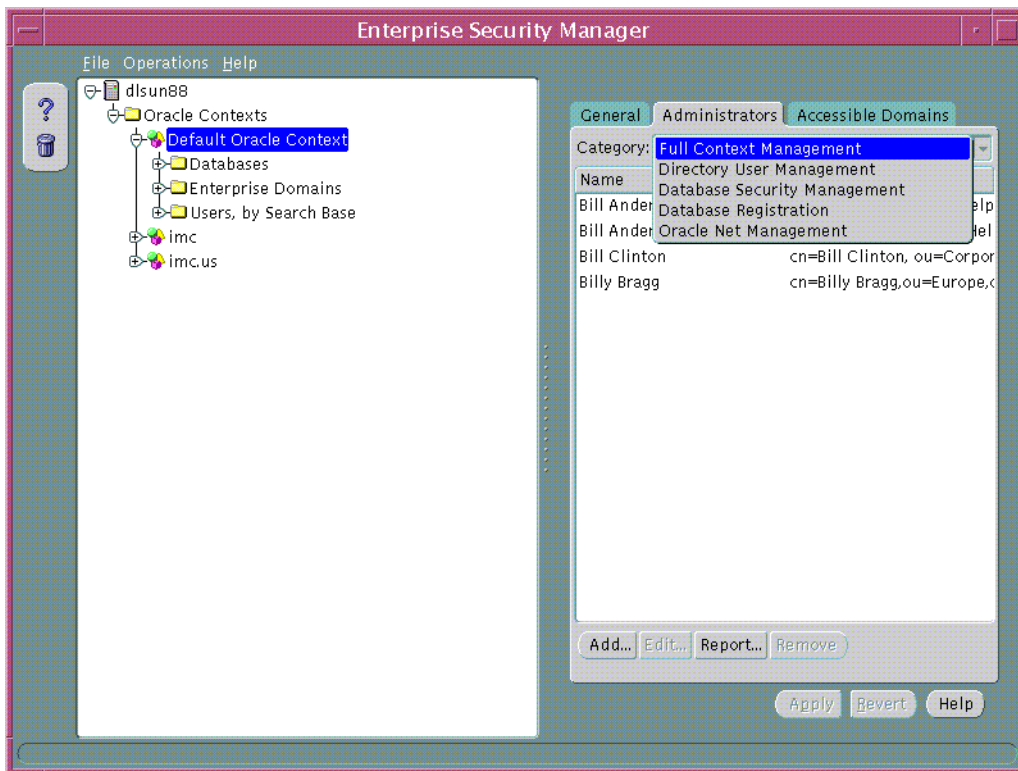
An Oracle Context may define sets of Directory Users that are enabled as different categories of Administrator. Each category has varying levels of privilege for operations within an Oracle Context. Some administrator categories are only available to version 9*i* Oracle Contexts and some are available to both version 8*i* and version 9*i* Oracle Contexts. The Administrator Categories for an Oracle Context are as follows:

Table 9–6 Oracle Context Administrator Categories

Administrator Category	Definition	Version 9i	Version 8i
Full Context Management	All possible Administrator privileges for all product areas in the Oracle Context	YES	NO
Directory User Management	Ability to view Directory User password reminders	YES	NO
Database Security Management	Ability to manage all Enterprise Domains and Enterprise Roles in the Oracle Context	YES	YES
Database Registration	Ability only to register a new database in the Oracle Context	YES	YES
Oracle Net Management	Ability to manage Oracle Net objects in the Oracle Context	YES	NO

Oracle Context Administrators are managed using the Administrators Page of an Oracle Context selected on the main application tree.

Figure 9–16 Oracle Context: Administrators Page



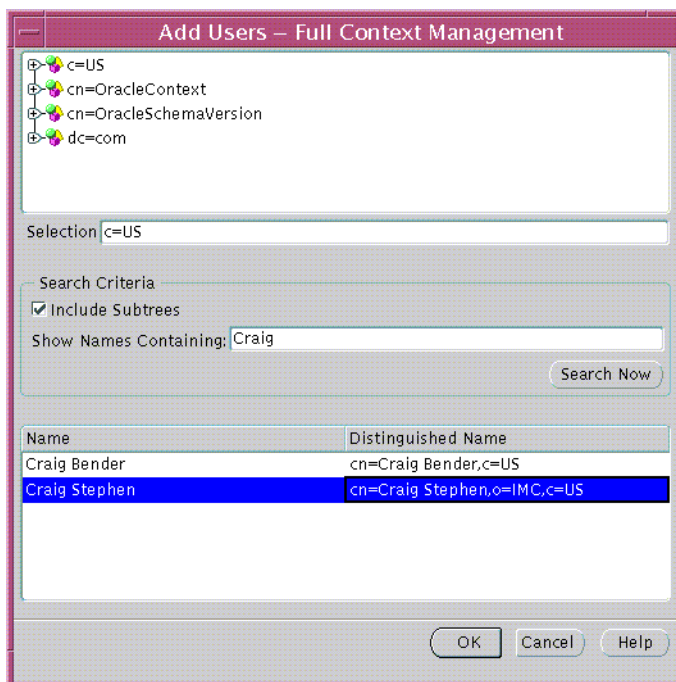
To remove a User from a list of Oracle Context Administrators:

1. Choose the type of Administrator to remove from the Categories combo box. The list of Administrators will refresh to show those of the type that you have selected.
2. Select a User by clicking on that User in the list of Administrators.
3. Choose Remove. The selected User will be removed from the list.
4. Choose Apply; the User will be removed as an Oracle Context Administrator of the category that you have selected.

To add a new User a list of Oracle Context Administrators:

1. Choose Add... The Add Users screen will appear. This page is used to locate and select one or more Users in the Directory. There are three components to the page. At the top is a Directory Search Tree. In the middle are Search Criteria controls that identify the Users to be returned by the search. At the bottom of the page is the result of the search from which one or more desired Users may be selected.

Figure 9–17 User Search Results



2. Navigate the Directory to select a desired Directory Entry as a User Search Base. You may also edit the contents of the Selection field in this screen to manually define the User Search Base.
3. Set the “Include Subtrees” Search Criteria option. The effect of selecting this option will be to search for Users not only as the specified Base but also in all possible levels underneath that Base in the Directory.
4. Enter any known User Name in the Show Names Containing field to which Users returned by the search must conform. The effect of using the Show

Names Containing field is to limit the search only to Users in the Directory who have a Common Name value that is or starts with the specified text.

5. Choose Search Now. If there are any Users in the Directory at the Base you have selected that match your Search Criteria they will be listed in the screen.
6. Select the desired User either by clicking on the it in the list and choosing OK or by double clicking on it. Multiple Users can be selected from the list by selecting a range of Users and choosing OK. The new Users will then appear in the list of Administrators under the category that you have selected.
7. Choose Apply; the new Administrators will be added to the Oracle Context in the Directory under the category that you have selected.

Note: This screen is used at all points in Enterprise Security Manager where it is necessary to choose one or more Users from the Directory.

Accessible Domains

When an Oracle Context is selected in the main application tree you may manage the list of Enterprise Domains within that Oracle Context whose databases may accept password authenticated connections from users that have their "Database Access Restriction" enabled. To add an Enterprise Domain to the list choose "Add.." and select one of the current Enterprise Domains from the resulting dialog. To remove an Enterprise Domain from the list, select it in the Accessible Domains page and choose "Remove.."

A "Database Access Restriction" may be applied to whole subtree of Users in the Directory when it is selected under the "Users, by Search Base" tree under an Oracle Context. With this option is set, all users under that subtree may only use their passwords to access databases that exist in Enterprise Domains that have been included in the list of Accessible Domains for the Oracle Context.

The default condition for any Enterprise Domain is not to be a member of the Accessible Domains for its Oracle Context. By identifying any Enterprise Domain to be one of the Accessible Domains and also by electing certain Users to have a Database Access Restriction, you are enforcing that it is only certain known databases that may access those Users' database logon settings in the Directory.

Note: This feature is only available to version 9 Oracle Contexts.

Managing Database Security

The Directory may be used as a central repository that controls authentication and authorization on multiple databases for Users. Enterprise Security Manager allows you to manage an Oracle Context in the Directory for the purpose of database security.

Oracle8i or 9i Databases are published to the Directory within an Oracle Context using the Oracle Database Configuration Assistant. For more information see the Oracle DBCA Guide. Once databases have been published to the Directory, Enterprise Security Manager may be used to manage User access to those databases. This is achieved using the following Objects in the Oracle Context:

Table 9–7 Oracle Context Objects

Object in the Oracle Context	Description
Database	This is a Directory Entry representing a published database.
Enterprise Domain	This is a grouping of databases published in the Directory upon which a common User access model for database security can be implemented
Enterprise Role	This is an Authorization that spans multiple databases within an Enterprise Domain. It is an “Enterprise Level” Role to which can be granted individual roles on each of the databases in an Enterprise Domain.

Enterprise Security Manager displays Databases and Enterprise Domains in its main application tree. Using our example of the AppsOnline Application Service provider, each of the company’s databases have been published into the 9i Default Oracle Context in the Directory.

Registering a Database with an Oracle Context

You register a database with an Oracle Context by selecting Register Database from the Enterprise Security Manager Operations menu. Selecting this menu option displays the Register Database dialog (shown in the following figure).

Figure 9–18 Register Database Dialog

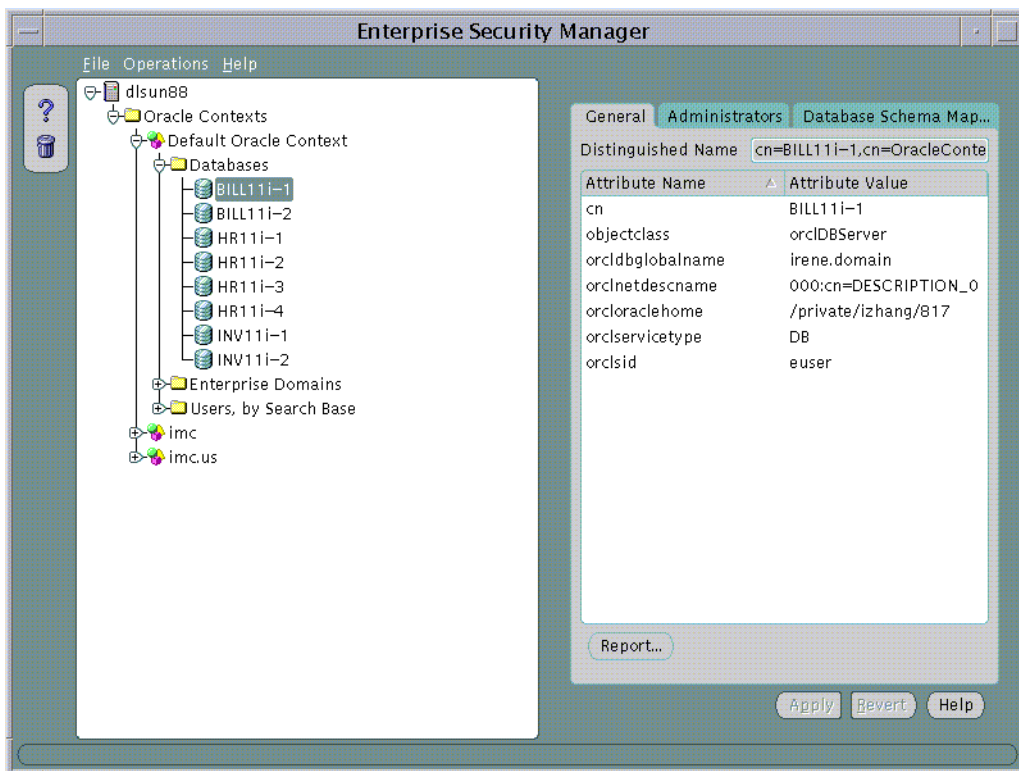
The screenshot shows a dialog box titled "Register Database". It contains the following fields and controls:

- Oracle Context: Root Context (dropdown menu)
- Oracle SID: (text box)
- Oracle Home: (text box)
- Oracle Database Version: 9.2.0.0.0 (dropdown menu)
- Hostname: (text box)
- Port Number: 1521 (text box)
- Store TNS Connect String
- Connect String: (text box)
- No Certificate Authority Available
- Password: (text box)
- Verify Password: (text box)

Buttons at the bottom: OK, Cancel, Help.

In addition to selecting the Oracle Context in which the database will reside, database registration also entails supplying the requisite connect information, as shown in the figure above. From this dialog, you can specify a database in one of three ways:

- Specify the SID, ORACLE_HOME, and hostname. Enterprise Security Manager will not verify the validity of these entries, hence, they must be valid, known values for a remote database.
- You can store a TNS connect string. Oracle Net LDAP Naming will be created.
- You can store a wallet for the database. This implies that a database can download that wallet using Oracle Wallet Manager while configuring for Enterprise User Security. You must enter a password to protect that wallet.

Figure 9–19 Security Manager Application Tree

In this example AppsOnline manages Oracle9i databases that host Applications for three customers; “UKMusic.com”, “CelticTravel.com” and “TaxTime.com”. Applications for UKMusic are hosted using databases INV11i-1 and INV11i-2. Applications for CelticTravel are hosted using databases BILL11i-1 and BILL11i-2. Applications for TaxTime are hosted using databases HR11i-1, HR11i-2, HR11i-3 and HR11i-4.

Given that the types of application hosted for each customer are different, only those databases that are used to support a common application type implement the same security model for their User Access. AppsOnline has decided to define three Enterprise Domains, one for each customer that it services.

Administering Databases

After a database has been published to an Oracle Context in the Directory, Enterprise Security Manager may be used to view and modify security characteristic of that database.

Managing Database Administrators

An Database Administrator is a Directory User that only has privileges to modify that Database in the Oracle Context. Database Administrators may be managed using the Administrators Page when a Database is selected under an Oracle Context in the main application tree.

To remove a User from the list of Database Administrators:

1. Select a User by clicking on that User in the list of Administrators.
2. Choose Remove. The selected User will be removed from the list.
3. Choose Apply; the User will be removed as an Database Administrator for that database in the Oracle Context.

To add a new User to the list of Enterprise Domain Administrators:

1. Choose Add... The Add Users screen will appear. This page is used to locate and select one or more Users in the Directory as discussed earlier. Select one or more desired users from the Directory to add as Database Administrators. The new Users will then appear in the in the Administrators Page.
2. Choose Apply; the new Administrators will be added to the database in the Oracle Context.

Managing Database Schema Mappings

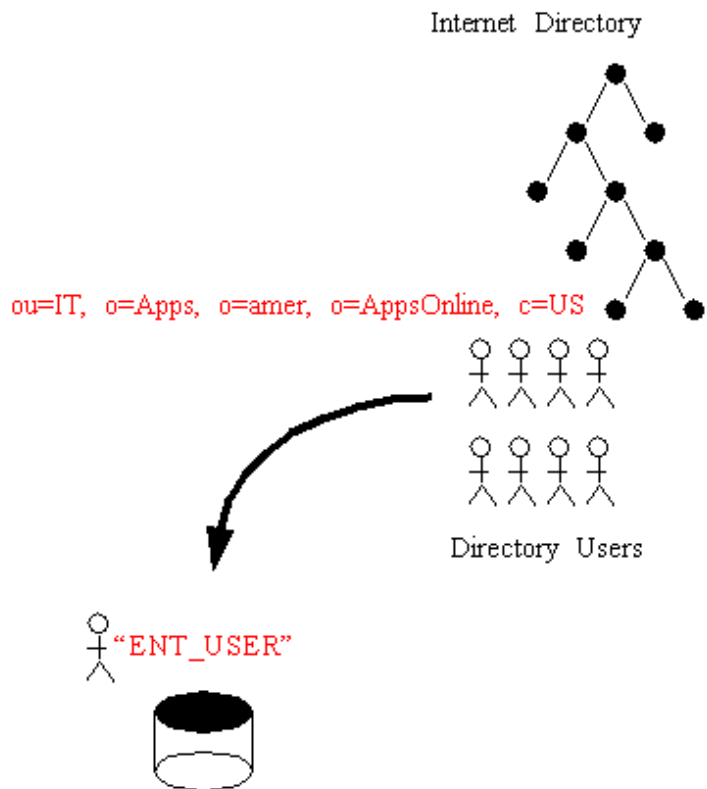
Database Schema Mappings allow databases that are registered in the Directory to accept connections from users without having any dedicated database schemas for them. For example, when user SCOTT connects to a database there must actually exist a database schema called "SCOTT" for that log on to be successful. This becomes difficult to maintain if there are thousands of Users and perhaps hundreds of databases in a very large enterprise.

Users that exist in the Directory do not need to have dedicated schemas on every Oracle8i or 9i database to which they might connect.

A database may use a "Schema Mapping" to share one database schema between any number of Users that exist in the Directory. The Schema Mapping is a pair of

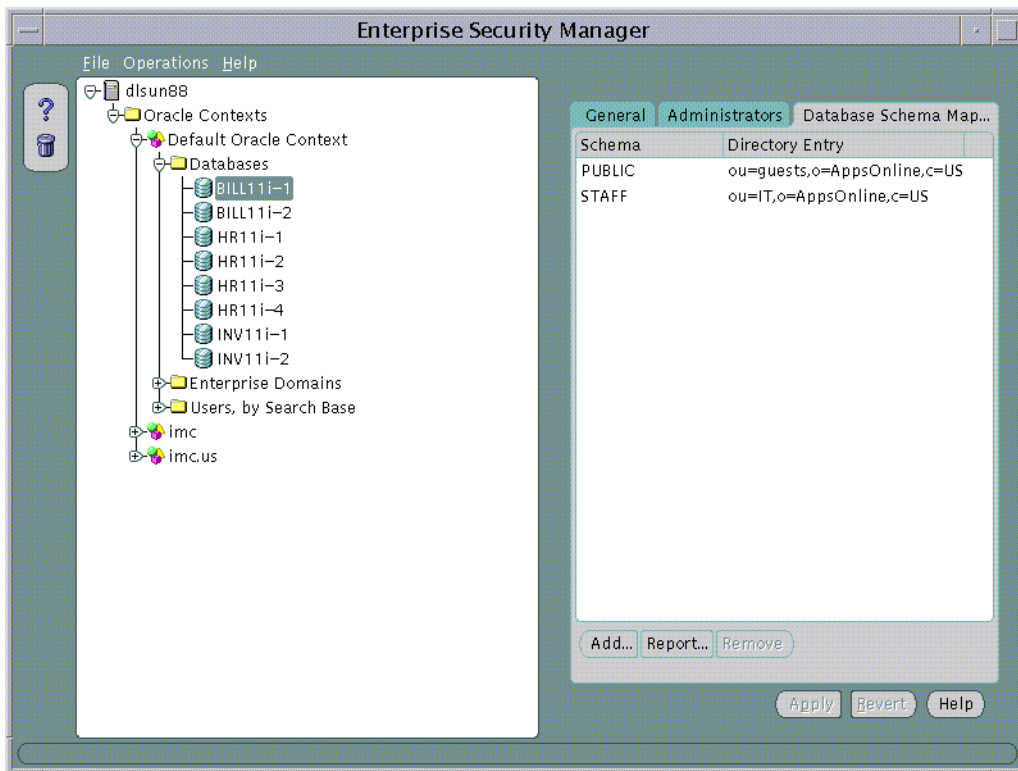
values; the Base in the Directory at which Users exist and the name of the database schema that they will use.

Figure 9–20 Database Schema Mappings



Database Schema Mappings may be managed using the Database Schema Mappings Page when a database is selected under an Oracle Context in the main application tree. This page contains a list of database schema name and Directory Base pairs.

Figure 9–21 Database Schema Mapping Page



To remove a Mapping from the list of Database Schema Mappings in the Enterprise Domain:

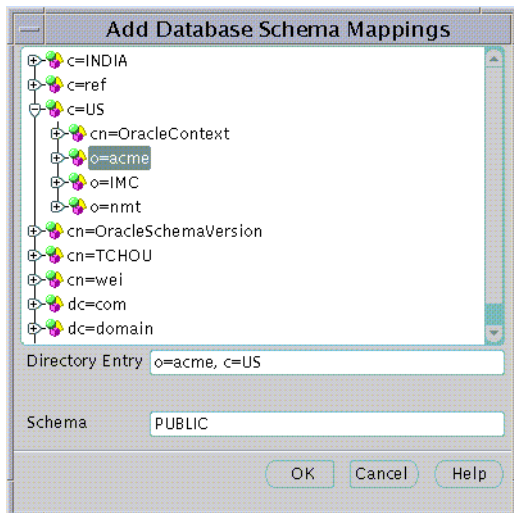
1. Select a Mapping by clicking on that Mapping in the list.
2. Choose Remove. The selected Mapping will be removed from the list.
3. Choose Apply; the Mapping will be removed from the Enterprise Domain and no longer used by any databases in the Enterprise Domain.

To add a new Mapping to the list of Database Schema Mappings in the Enterprise Domain:

1. Choose Add... The Add Database Schema Mappings screen will appear. This page is used to locate and select one Base in the Directory and pair it with a database schema name to make a Database Schema Mapping. There are two

components to the page. There is a Directory Search Tree from which to select a Base and a field in which to enter a schema name.

Figure 9–22 Add Database Schema Mappings



2. Navigate the Directory to select a desired Directory Entry as a Base for the Database Schema Mapping. This may be any Directory Entry but should be above the subtree of Users in the Directory for which you want to perform the mapping. You may also edit the contents of the Selection field in this screen to manually define this Base.
3. Enter the name of the database schema for which this Mapping will be made and choose OK. This must be a valid name for a schema that already exists on that database. The new Database Schema Mapping will then appear in the Database Schema Mappings Page.
4. Choose Apply; the new Database Schema Mapping will be added to the selected database in the Oracle Context.

Administering Enterprise Domains

An Oracle Context will always contain at least one Enterprise Domain called, “OracleDefaultDomain”. The OracleDefaultDomain is part of the Oracle Context when it is first created in the Directory. When a new database is registered into an

Oracle Context it automatically becomes a member of the OracleDefaultDomain in that Oracle Context. You may create and remove your own Enterprise Domains but you cannot remove the OracleDefaultDomain from an Oracle Context.

To create a new Enterprise Domain:


An Enterprise Domain can be created in an Oracle Context either from the Operations Menu or by using a Right Mouse Button click on an Oracle Context selected in the main application tree:

Figure 9–23 *Creating an Enterprise Domain*



The Create Enterprise Domain screen will appear.

Figure 9–24 *Create Enterprise Domain Dialog*



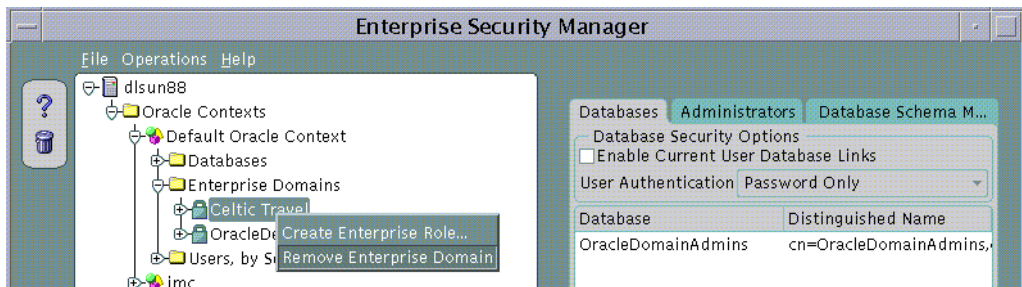
1. Choose the Oracle Context in which the Enterprise Domain is to be created from the Oracle Context drop down list. If the Create Enterprise Domain screen has been invoked using a Right Mouse Button click from an Oracle Context in the main application tree then the name of that Oracle Context will already be selected.
2. Enter the name of the new Enterprise Domain in the Domain Name field.

3. Choose OK. The new Enterprise Domain will be created in the Oracle Context and appears on the main application tree.

To remove an Enterprise Domain:

1. Click on the Enterprise Domain to remove in the main application tree.
2. Choose Remove Enterprise Domain either from the Operations Menu or by using a Right Mouse Button Click on the Enterprise Domain in the main application tree.

Figure 9–25 Remove Enterprise Domain Menu Option



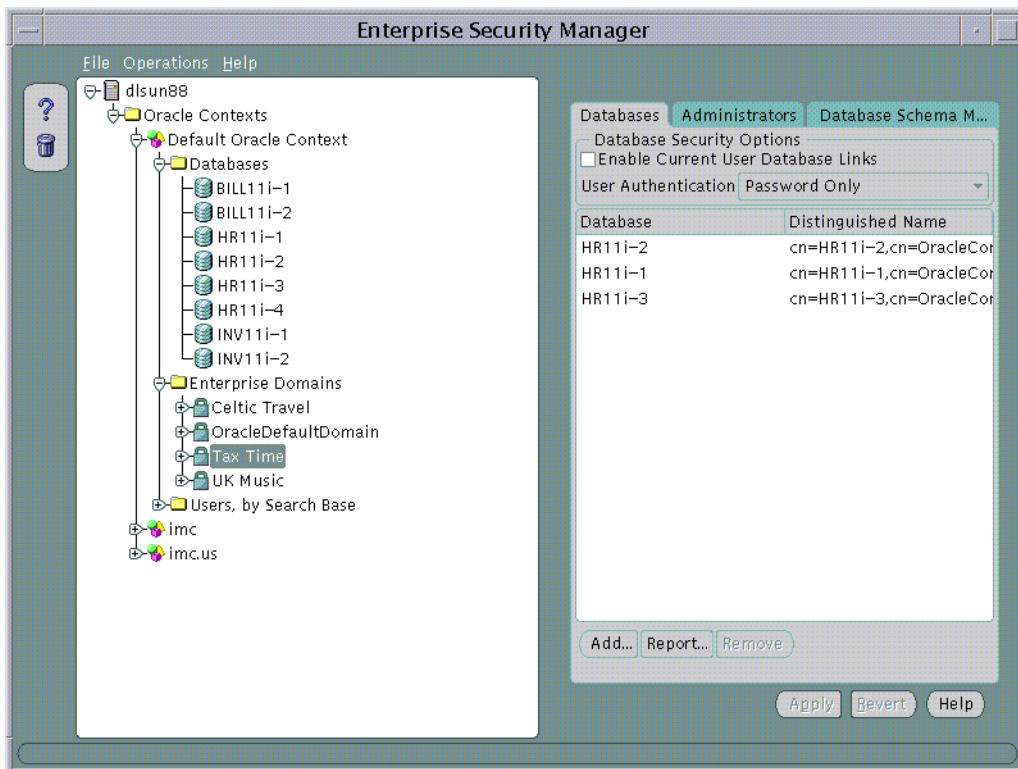
3. Enterprise Security Manager will ask you to confirm the operation before the Enterprise Domain is removed from the Oracle Context.

Note: You cannot remove an Enterprise Domain from an Oracle Context if that Enterprise Domain still contains any Enterprise Roles.

Specifying Database Membership of an Enterprise Domain

Database membership of an Enterprise Domain in the Oracle Context may be managed using the Databases Page when an Enterprise Domain is selected on the main application tree:

Figure 9–26 Security Manager Databases Page



To remove a database from an Enterprise Domain:

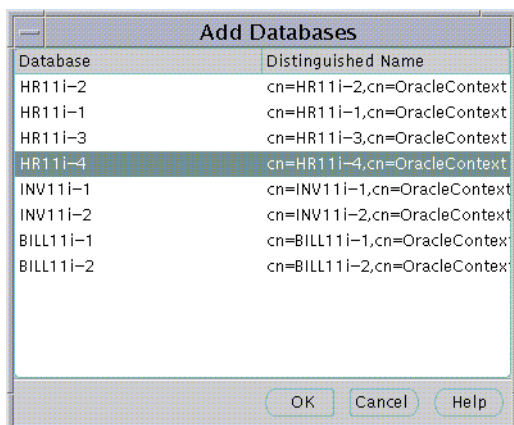
1. Select a database in the list and choose Remove... The database will be removed from the list.
2. Choose Apply; the database will be removed from the Enterprise Domain in the Oracle Context.

To add a database to an Enterprise Domain:

Note: You may only add databases as members of an Enterprise Domain that exist in the same Oracle Context as the Enterprise Domain. An Enterprise Domain cannot contain as its members, databases published in a different Oracle Context. Neither can any database in an Oracle Context be added as a member of two Enterprise Domains.

1. Choose Add... The Add Databases screen will appear. This screen lists all the databases in the Oracle Context

Figure 9–27 Add Databases Dialog



2. Select a database to add as a new member of the Enterprise Domain.
3. Choose OK in the Add Databases screen. The selected database will be added to the list of databases in the Databases Page.
4. Choose Apply; the new database will be added to the Enterprise Domain in the Oracle Context.

Managing Database Security Options for an Enterprise Domain

The Databases Page may be used to manage database security options that will apply to all the databases that are members of the Enterprise Domain. These options are as follows:

Table 9–8

Database Security Option	Description
Enable Current User Database Links	Any pair of databases will only allow use of Current User Database Links if they exist in an Enterprise Domain in which this setting is enabled.
User Authentication	All databases in the Enterprise Domain will enforce the type of authentication that its clients must use based on this property. Its values are: <ul style="list-style-type: none"> ▪ Password Authentication only. ▪ Oracle Net SSL Authentication only using Oracle Wallets. ▪ Either Password or Oracle Net SSL Authentication.

Managing Enterprise Domain Administrators

An Enterprise Domain Administrator is a Directory User that only has privileges to modify the content of that Enterprise Domain. Enterprise Domain Administrators may be managed using the Administrators Page when an Enterprise Domain is selected under an Oracle Context in the main application tree.

To remove a User from the list of Enterprise Domain Administrators:

1. Select a User by clicking on that User in the list of Administrators.
2. Choose Remove. The selected User will be removed from the list.
3. Choose Apply; the User will be removed as an Enterprise Domain Administrator for that Enterprise Domain in the Oracle Context.

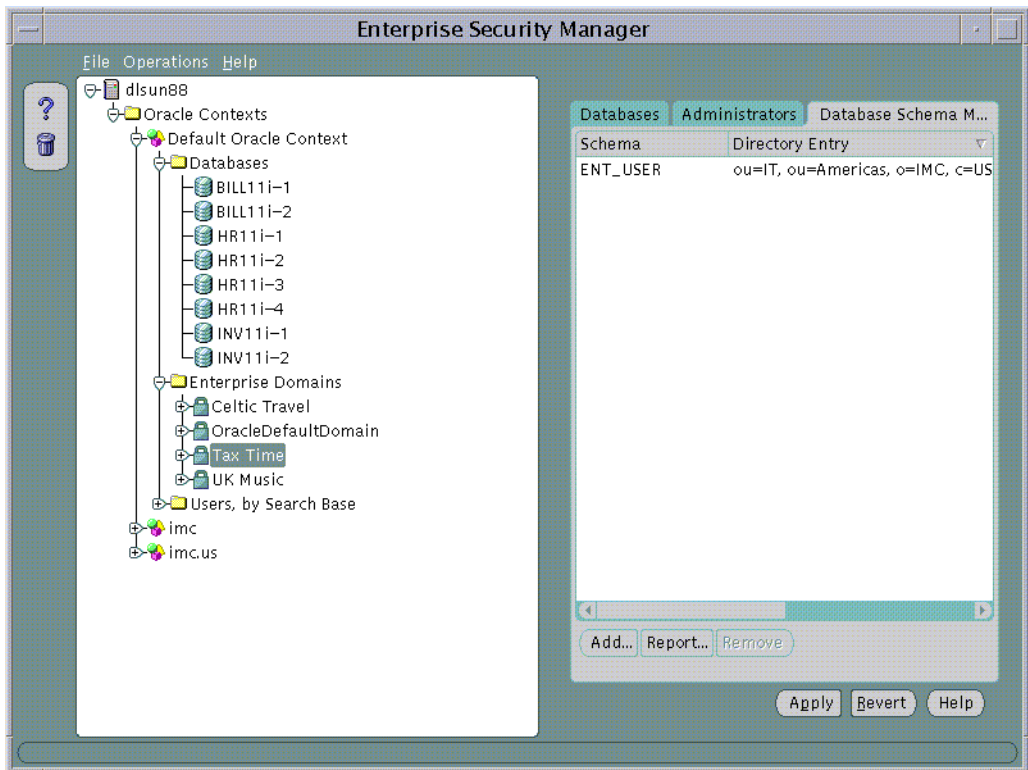
To add a new User to the list of Enterprise Domain Administrators:

1. Choose Add... The Add Users screen will appear. This page is used to locate and select one or more Users in the Directory as discussed earlier. Select one or more desired users from the Directory to add as Enterprise Domain Administrators. The new Users will then appear in the Administrators Page.
2. Choose Apply; the new Administrators will be added to the Enterprise Domain in the Oracle Context.

Managing Enterprise Domain Database Schema Mappings

Database Schema Mappings may be managed for each database in an Oracle Context as discussed earlier. Schema Mappings may also be performed for each Enterprise Domain in an Oracle Context using the Database Schema Mappings Page with an Enterprise Domain selected in the main application tree. These Mappings apply to all databases that are members of the Enterprise Domain. Therefore, each database in the Enterprise Domain must have a schema of the same name used in the Mapping.

Figure 9–28 Matching Database and Schema Names Used in the Mappings



To remove a Mapping from the list of Database Schema Mappings in the Enterprise Domain:

1. Select a Mapping by clicking on that Mapping in the list.

2. Choose Remove. The selected Mapping will be removed from the list.
3. Choose Apply; the Mapping will be removed from the Enterprise Domain and no longer used by any databases in the Enterprise Domain.

To add a new Mapping to the list of Database Schema Mappings in the Enterprise Domain:

1. Choose Add... The Add Database Schema Mappings screen will appear. This page is used to locate and select one Base in the Directory as discussed earlier. Enter a new Database Schema Mapping to add to the Enterprise Domain.
2. Choose Apply; the new Database Schema Mapping will be added to the Enterprise Domain selected in the Oracle Context.

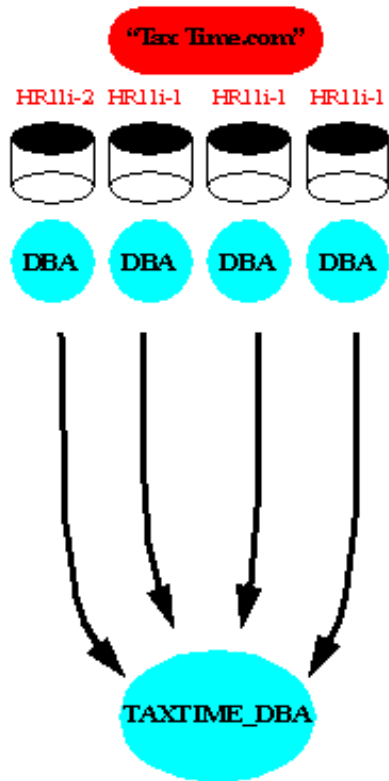
Administering Enterprise Roles

An Enterprise Domain within an Oracle Context may contain one or more Enterprise Roles.

In the example discussed earlier, AppsOnline has created three Enterprise Domains that group the databases it uses to serve each of its customers. This permits the company to define Enterprise Roles for each Enterprise Domain. An Enterprise Role is a set of Oracle Role based authorizations across on or more databases in an Enterprise Domain.

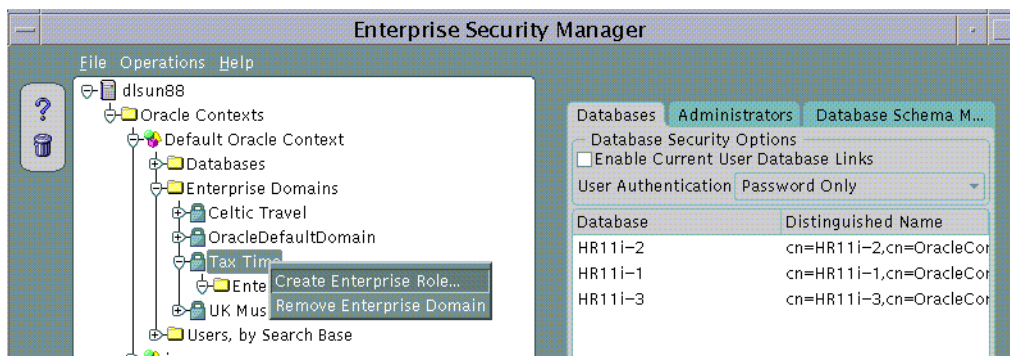
A simple Enterprise Role is defined by AppsOnline for DBA privileges on its databases in the “Tax Time” Enterprise Domain:

Figure 9–29 "Tax Time" Enterprise Domain

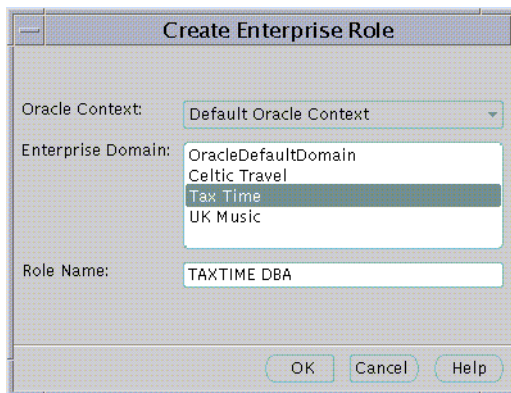


Creating a new Enterprise Role:

An Enterprise Role can be created in an Enterprise Domain either from the Operations Menu or by using a Right Mouse Button click on an Enterprise Domain selected in the main application tree:

Figure 9–30 Enterprise Role Creation

The Create Enterprise Role dialog appears.

Figure 9–31 Create Enterprise Role Dialog

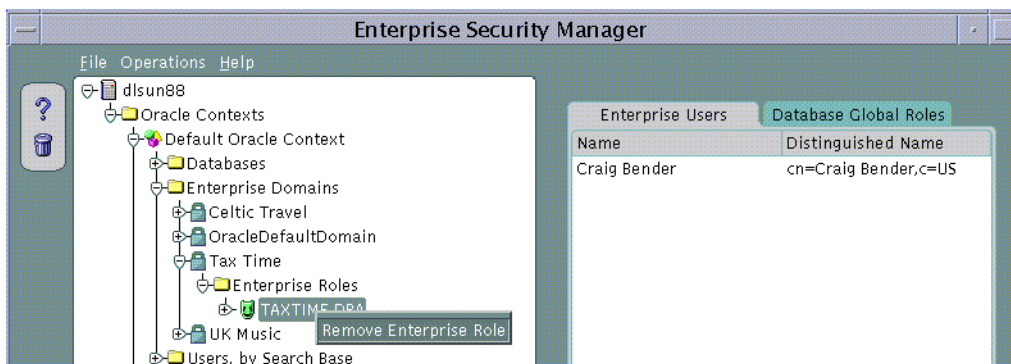
1. Choose the Oracle Context containing the Enterprise Domain in which the new Enterprise Role is to be created from the Oracle Context drop down list. If the Create Enterprise Role screen has been invoked using a Right Mouse Button click from an Enterprise Domain selected in the main application tree, then the name of that Oracle Context will already be selected.

2. Choose the Enterprise Domain in which the new Enterprise Role is to be created from the Enterprise Domain list. If the Create Enterprise Role screen has been invoked using a Right Mouse Button click from an Enterprise Domain selected in the main application tree, then the name of that Enterprise Domain will already be selected.
3. Enter the name of the new Enterprise Role in the Role Name field.
4. Choose OK. The new Enterprise Role will be created in the Enterprise Domain and appears on the main application tree.

Removing an Enterprise Role:

1. Click on the Enterprise Role to remove in the main application tree.
2. Choose Remove Enterprise Role either from the Operations Menu or by using a Right Mouse Button Click on the Enterprise Domain in the main application tree.

Figure 9–32 Removing an Enterprise Role



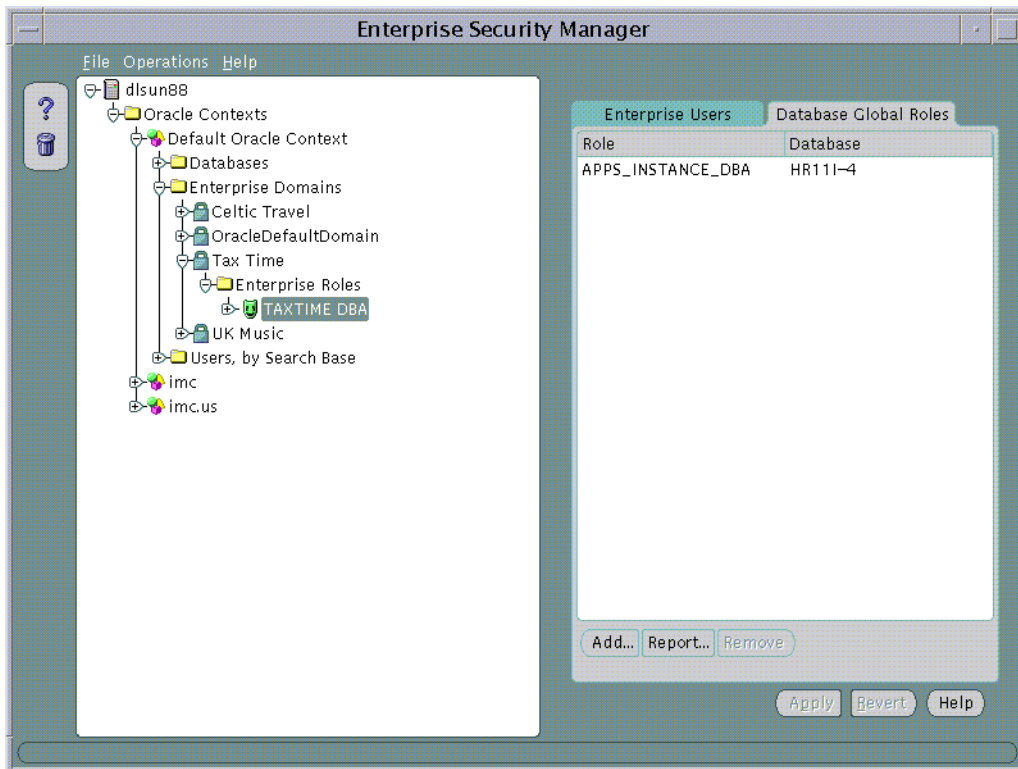
3. Enterprise Security Manager will ask you to confirm the operation before the Enterprise Role is removed from the Enterprise Domain.

Specifying Database Global Role Membership of an Enterprise Role

Database Role membership of an Enterprise Role in an Enterprise Domain may be managed using the Database Global Roles Page when an Enterprise Role is selected on the main application tree. This page lists the names of each Global Role that

belongs to the Enterprise Role along with the name of the database on which that Global Role exists.

Figure 9–33 Database Global Roles Page



When populating an Enterprise Role with different database roles it is only possible to reference roles on databases that are configured to be “Global Roles” on those databases. A Global Role on a database is identical to a normal Role, except that the administrator of the database has elected it only to be authorized via the Directory. A database administrator cannot locally grant and revoke Global Roles to users of the database.

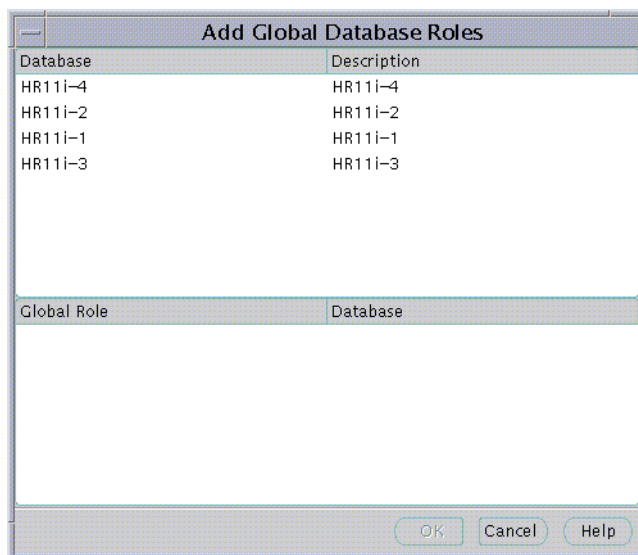
Removing a Database Global Role from an Enterprise Role:

1. Select a Global Role in the list and choose Remove... The Global Role will be removed from the list.
2. Choose Apply; the Global Role will be removed from the Enterprise Role in the Enterprise Domain.

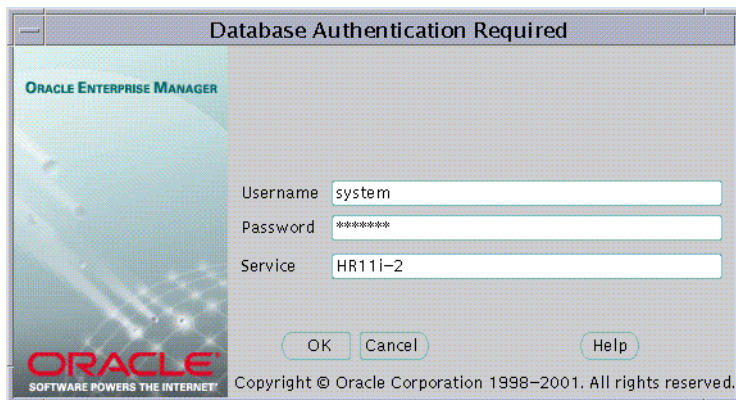
Adding a Global Role to an Enterprise Role:

1. Choose Add... The Add Database Global Roles screen will appear. This screen lists all the databases in the Enterprise Domain from which Global Roles may be selected to add to this Enterprise Role

Figure 9–34 Add Global Database Roles Dialog



2. Select a database from which to obtain Global Roles. A screen will appear in which you must enter logon details to authenticate to the database and fetch Global Roles. Typically this would be a DBA logging on to that database.

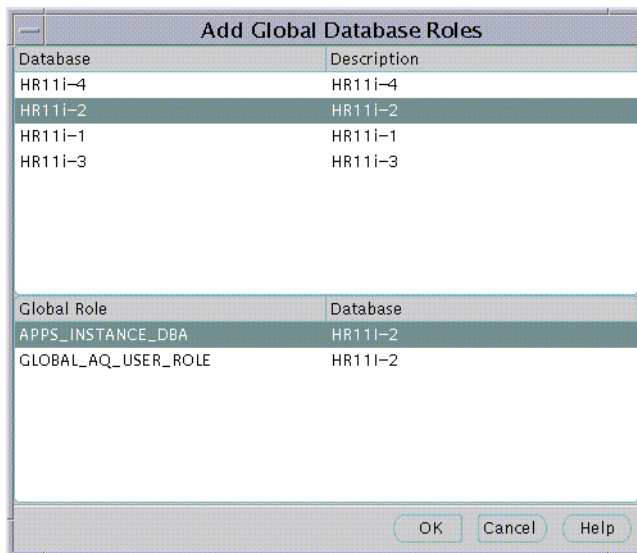
Figure 9–35 Database Logon

The name of the database appears in the Service field by default. You may use this name to connect to the database if your ORACLE_HOME has LDAP enabled as it Oracle Net Naming method or if this name appears as a TNS alias in your local Oracle Net configuration. Otherwise you may overwrite the content of the Service field with any other TNS alias configured for that database or by a connect string in the format:

```
<host>:<port>:<oracle sid>
```

For example, "cartman:1521:broncos"

3. Choose OK. Enterprise Security Manager will connect to the given database and fetch the list of Global Roles supported on that database. The list of values, if any, will appear in the Add Database Global Roles screen.

Figure 9–36 Add Global Database Roles Dialog

4. Select on or more of the Global Roles from the list of returned values and choose OK. These Global Roles will then appear in the Database Global Roles Page
5. Choose Apply; the new Global Roles will be added to the Enterprise Role in the Enterprise Domain.

Managing Enterprise Role Grantees

An Enterprise Role Grantee is a Directory User to whom has been granted an Enterprise Role and therefore all database Global Roles contained within that Enterprise Role. Enterprise Role Grantees may be managed using the Enterprise Users Page when an Enterprise Role is selected under an Enterprise Domain in the main application tree.

Removing a User from the List of Enterprise Role Grantees:

1. Select a User by clicking on that User in the list of Grantees.
2. Choose Remove. The selected User will be removed from the list.

3. Choose Apply; the User will be removed as a Grantee for that Enterprise Role in the Enterprise Domain.

Adding a New User to the list of Enterprise Role Grantees:

1. Choose Add... The Add Users screen will appear. This page is used to locate and select one or more Users in the Directory as discussed earlier. Select one or more desired users from the Directory to add as Enterprise Role Grantees. The new Users will then appear in the Enterprise Users Page

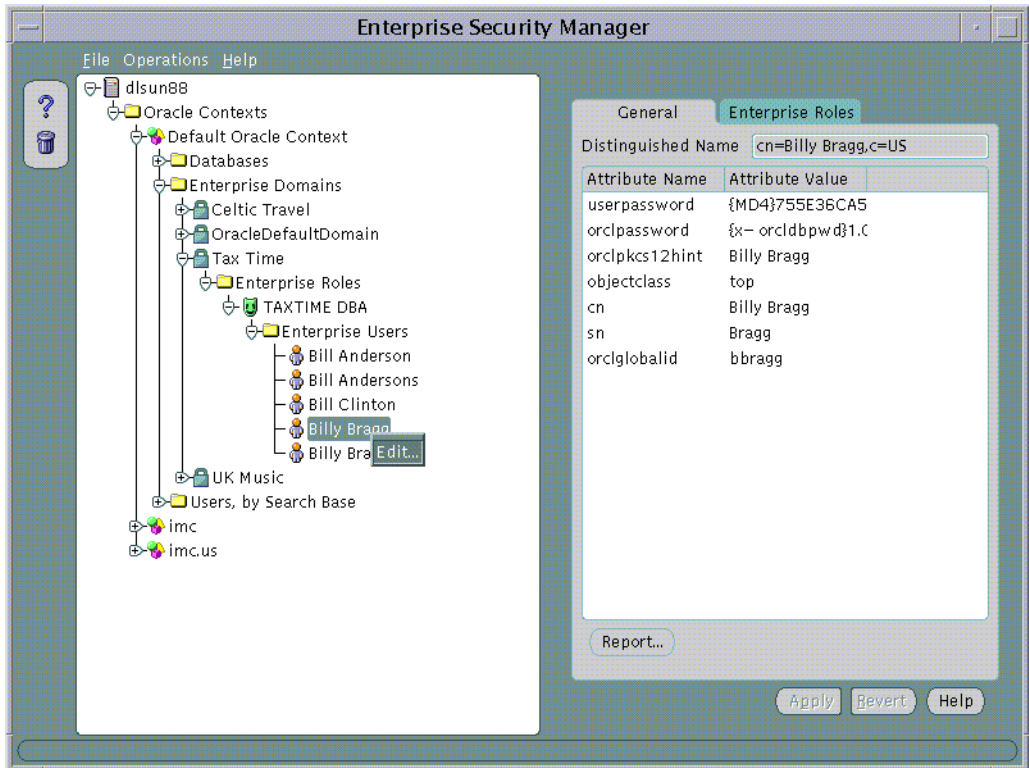
Figure 9–37 Enterprise Users Page



2. Choose Apply; the new Grantees will be added to the Enterprise Role in the Enterprise Domain.

Enterprise Role Grantees will also appear in the Enterprise Users tree under a selected Enterprise Role. A User selected on this tree can be edited as discussed in Part 1.

Figure 9–38 Enterprise Role Grantees



Command Line Tool

In addition to the graphical user interface, Enterprise Security Manger also provides a full-featured command line interface that allows you to perform security administration operations from other applications or from custom scripts.

Command line operations are invoked using the "esm" control utility. For example:

```
esm -cmd [operation] [options]
esm -cmd help [operation]
```

A complete list of operations and options, including definitions and usage syntax, is available online by invoking the tool's online help from the command line.

```
>esm -cmd help
```

The following example shows the help output.

Example 9-1 Command Line Help Output

```
Usage :
esm -cmd [operation] [options]
esm -cmd help [operation]

Operations :
[search, createUser, deleteUser, createWallet, createDomain, deleteDomain, createRole,
deleteRole, grantEnterpriseRole, revokeEnterpriseRole, addGlobalRole, removeGlobalRole,
addContextAdministrator, removeContextAdministrator, addPasswordAccessibleDomain,
removePasswordAccessibleDomain, addDomainAdministrator, removeDomainAdministrator,
addDomainDatabase, removeDomainDatabase, addDatabaseAdministrator,
removeDatabaseAdministrator, createMapping, removeMapping]

Options :
(* mandatory)
-U          SSL Authentication Mode - Should be SIMPLE / SSL / NATIVE
-h          LDAP Server *
-p          LDAP Server Port *
-D          Bind DN (required for SIMPLE Login)
-w          Bind Password (required for SIMPLE Login)
-W          Wallet Location (required for SSL Login)
-P          Wallet Password (required for SSL Login)
-dn        DN * (for user, domain, enterprise role, context or database)
-objectType Type of Object for search [user | database | domain | enterpriseRole |
context | schemaMapping | database | domainDatabase | fullContextAdministrator |
directoryUserAdministrator | oracleNetAdministrator | databaseSecurityAdministrator |
databaseRegistrationAdministrator | databaseAdministrator | domainAdministrator]
-firstname      User First Name
-lastname      User Last Name
-userID         User ID
-password      User Directory Authentication Password
-wcheck        User Wallet Check (true / false)
-context       Oracle Context DN
-userDN        User DN (required for assign / revoke operations)
-domainDN      Domain DN (required for assign / revoke operations)
-adminType     Administrator Type [context | user | databaseSecurity |
databaseInstall
| network]
-databaseRoleDN Database Global Role (in format <Database
DN>,GlobalRole=<GlobalRoleDN>)
-databaseDN    Database DN (required for assign / revoke operations)
```

```
-walletPwd      Wallet Password (required for Wallet creation)
-rootPwd       Root Password (required for Wallet Creation)
-target        Mapping Target Schema
-value         Mapping Directory Entry
-level         Mapping Level [1 (Entry)| 0 (Subtree)]
```

You may also display help pertaining to a specific Enterprise Security Manager operation by specifying the exact operation along with the help command:

```
esm -cmd help <operation>
```

Help for specific operations provides a usage sample. For example, executing "esm -cmd help createUser" from the command line displays the following help text:

```
Usage :
esm -cmd [operation] [options]
esm -cmd help [operation]

Operations :
createUser

Options :
(* mandatory)
-U          SSL Authentication Mode - Should be SIMPLE / SSL / NATIVE
-h          LDAP Server *
-p          LDAP Server Port *
-D          Bind DN (required for SIMPLE Login)
-w          Bind Password (required for SIMPLE Login)
-W          Wallet Location (required for SSL Login)
-P          Wallet Password (required for SSL Login)
-dn        User DN *
-firstname  User First Name *
-lastname   User Last Name *
-userID     User ID
-password   User Directory Authentication Password
-wcheck     User Wallet Check (true / false)

Example :
esm -cmd createUser -U SIMPLE -D orcladmin -w welcome -h dlsun1279.us.oracle.com -p 389
-dn
cn=TestUser -firstname Test -lastname User -userID RM -password testpass -wcheck false
```

An example of how the command line tool can be used is provided with your Enterprise Manager installation. The shell script "esmdemo" is located in the ORACLE_HOME/sysman/admin directory and showcases Enterprise Security Manager command line usage. Running this script performs sample operations using the command line tool. View the contents of this script to see working examples of how the command line tool can be used.

Part II

Database Administration Tools

- Chapter 10, "Database Administration"
- Chapter 11, "Managing Backup and Recovery"

Database Administration

The database administration features and wizards are integrated into Oracle Enterprise Manager. You can access the database administration features through the Console. The Console can either be launched with a connection to an Oracle Management Server, which utilizes Oracle Enterprise Manager's three-tier framework, or launched standalone, which connects directly to a database.

The Console's Databases folder allows you to administer database instances, schemas, security, and storage, and other database features from a unified tree view. The unified access to administration functions offered by the Databases folder makes it easy to switch between tasks and to gain an accurate overall view of the database configuration status. When you expand a database in the tree, a list of database features appears below.

- Instance, including startup, shutdown, and initialization.
- Schema, including tables, indexes, and all other schema objects.
- Security, including user accounts, roles, and privileges.
- Storage, including tablespaces, datafiles, rollback segments, redo log groups, and archive logs.
- Distributed, including in-doubt transactions, database links, Streams, Advanced Queues, and Advanced Replication
- Warehouse, including summary management and OLAP management. Summary Management includes tools for improving the performance of a data warehouse. OLAP management includes tools for creating and editing OLAP metadata based on a star or snowflake schema.
- Workspace, including a virtual environment that one or more users can share to make versioned changes to data.

-
- XML Database, including tools for storing and retrieving XML objects and optimizing access and updates to XML objects.
 - Other database features, depending on the products installed with your database.

You can use the Database folder features with or without connecting to an Oracle Management Server.

The Backup and Recovery wizards are also available to help you back up or restore and recover various objects such as the tablespaces, datafiles, or archive logs. With the Backup wizard you can also make an image copy of the datafiles and the current controlfile. Beginning with Oracle Enterprise Manager 9.2, the Backup Wizard allows the setting of additional options, such as backup retention policy, deleting obsolete backups and specifying the archive log deletion policy. With Oracle 9.2, recovery includes Block Media Recovery which improves the speed of recovery significantly in the case of block corruptions.

Oracle Enterprise Manager now features the SQL Scratchpad, which provides a user interface for you to enter, edit, and execute SQL quickly and easily.

Common Features of Database Management Features

This section discusses the common features shared by the database administration features in Enterprise Manager.

Tree Views

The Console displays a tree view of connected databases, which can be expanded to show subordinate objects.

General Information about Databases

When you select a database node in the tree, a non-editable General page appears on the right where you can view information about the host, port, SID, TNS descriptors, setup information (Oracle_Home and Listeners), and Operating System information.

Comprehensive Overview Pages

When you select any of these database features, a brief description of the feature appears in the comprehensive overview page on the right side of the Console. Depending on the feature, the page may contain a link to obtain more information, or to start a process, or a button to launch the related Quick Tour or Help screen.

Property Sheets

When relevant, if you select an object in the tree, a property sheet appears on the right where you can view or edit database properties. Wizards also display tree views and property sheets as necessary.

Multi-Column Lists

In most cases when you select any of the database features such as Schema (with the exception of Advanced Queues), Instance, and so on, a multi-column list of all the folder's objects appears on the right side of the Console, providing a quick summary of information about each object in the selected folder.

Database Version Awareness

All database features and wizards are aware of the features that are available in each database version. When you select a database in a tree view, the tool only displays objects and properties that are enabled in that database version.

Database Reports

You can extract information from the database such as object definitions, object dependencies, database configuration, or reports, including custom SQL queries.

Logging of Database Changes

You can now log all Data Definition Language (DDL) and Data Manipulation Language (DML) changes made by an application when connected to a database.

Showing Object DDL

Data Definition Language (DDL) commands set up the data such as creating and altering databases and tables. You can display the Data Definition DDL for objects.

Show SQL

Though one of the benefits of Enterprise Manager is that DBA tasks can be performed without manually entering SQL, you do have the option of viewing the SQL code generated for you. By selecting Show SQL button, you can review this code before implementing any changes, as well as copy and paste it into your own SQL scripts if you wish.

Show Dependencies

Database object dependencies and dependents can be viewed by right-mouse clicking an object in the tree view and choosing Show Dependencies. Dependencies show what the selected object depends on, such as the tablespace location and the owner of the selected object. Dependents rely on the selected object, such as which indexes will be dropped and which synonyms will be affected if you drop the selected object.

Right-Mouse Commands

With the database features, you can right-mouse click any folder or object in a tree list to perform administrative tasks. Right-mouse clicking an object shows all the tasks that can be performed on the object, such as connecting to or disconnecting from the database, creating users, adding or removing profiles, assigning privileges, showing dependencies, and bringing up wizards.

DB Search Capabilities

Database Search allows you to search for any object in a database given a flexible set of criteria:

- Names of the objects that you want to find. You can enter wild card characters for the object name.
- Database that you want to search.
- Object types that you want to include in the search.
- Schemas that you want to include in the search.

The object definitions that match the search criteria are displayed in a multi-column table.

Database Management Features and Wizards

This section describes Database Management features and wizards.

DBA Management Features	Task	See Page
Instance Management	Manages instances and sessions	10-7
Schema Management	Manages schema objects	10-14
Security Management	Manages security parameters	10-19
Storage Management	Manages database storage	10-21
Distributed Management	Manages in-doubt transactions, database links, streams, advanced queues, and advanced replication	10-24
Warehouse Management	Manages the performance of a data warehouse (Summary Management) and OLAP metadata using the CWMLite Release 1 APIs (OLAP management manages).	10-25
Workspace Management	Allows you to version-enable tables and create, modify, refresh, and merge workspaces	10-27
XML Database	Manages the storing and retrieving of XML objects and optimizing access and updates to XML objects.	10-28
SQL*Plus Worksheet	Executes SQL and PL/SQL commands	10-29
SQL Scratchpad	Provides a user interface for you to enter, edit, and execute SQL quickly and easily	10-30
Wizards	Assist with importing, exporting, loading, backing up, and recovering data, as well as analyzing and creating tables and views	10-32

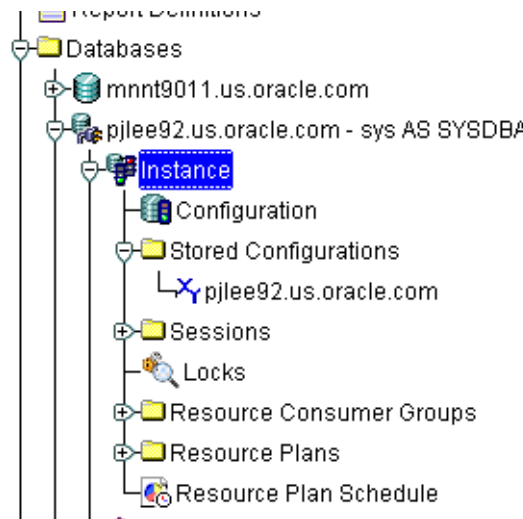
Instance Management

The Instance Management feature helps you manage database instances and sessions in your Oracle environment. With the Instance Management feature you can:

- Start up and shut down a database.
- View and edit the values of instance parameters.
- Tune your database resources for optimal use with the help of Memory and Mean-Time-To-Recover (MTTR) Advisors.
- Manage users' sessions, and view currently running SQL and its explain plan.
- Administer locks and sessions consuming the highest amounts of resources (if the Diagnostics Pack is installed).
- Monitor long-running operations.
- Control processing resources via Resource Plans.
- Perform backup, recovery and maintenance operations on the database files.

When you expand the Instance node under the database in the tree view, the following list of objects and folders appears:

- Configuration
- Stored Configurations (only when connected to Oracle Management Server)
- Sessions
- Locks
- Resource Consumer Groups
- Resource Plans
- Resource Plan Schedule

Figure 10–1 Instance Management

Configuration Operations

When you select the Configuration node under Instance, a property sheet of tabbed pages appears on the right for viewing information about the database instance and editing database properties.

General Page The General Page shows the following information which can be viewed and/or edited:

- Status of the instance, including the database version and any installed options, and allows you to start up and shut down a database.

For information on cluster databases, refer to *Oracle9i Real Application Clusters Administration*.

- Location of the spfile if the database had been started with an spfile
- Persistent parameters which allow you to modify, apply changes, and re-initialize the database.

Note: The Configured mode for a cluster database displays an additional "Instance Name" column, which helps distinguish the parameters as either database-wide or instance-specific. If this column is blank for a parameter; then, it implies that the parameter's value is available database wide (i.e. for all

instances). A string in this column specifies the SID of a specific cluster database instance to which the parameter's value is applicable.

Memory Page The Memory Page allows you to perform the following tasks:

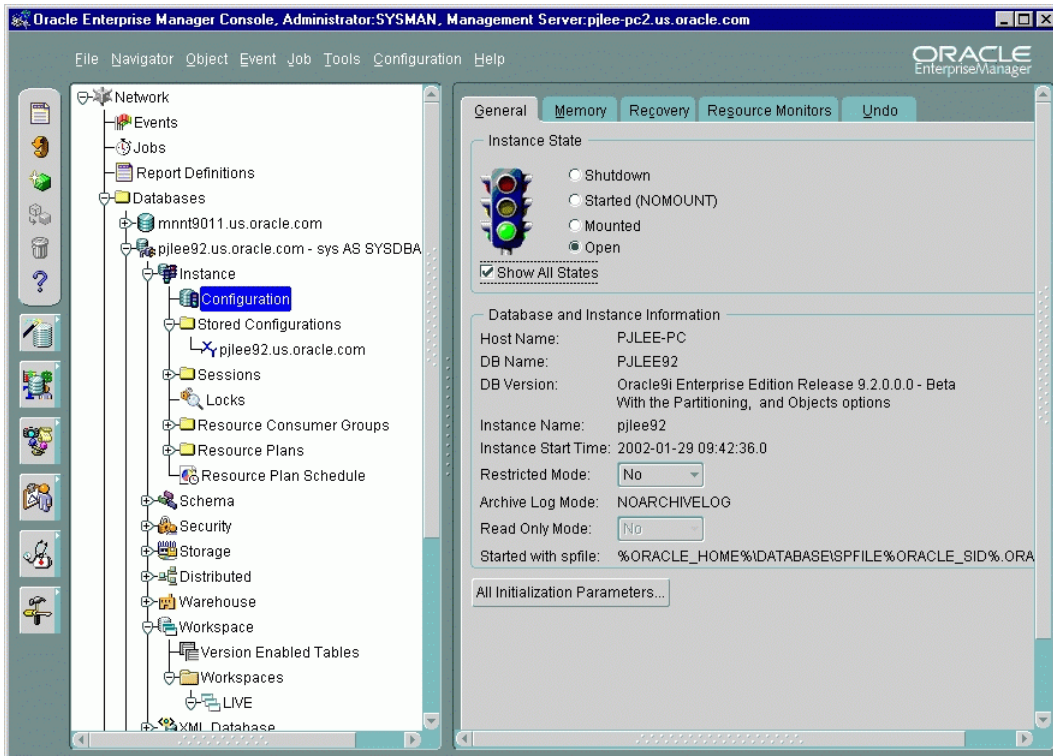
- View information about the memory usage of the current database.
- Use the Shared Pool Size Advisor to determine optimal Shared Pool size by tracking its use by the Library Cache. This advisor is available starting with Oracle 9.2.
- Use the Buffer Cache Size Advisor to determine the optimal size of the buffer cache.
- Use the Program Global Area (PGA) Advisor to tune PGA memory allocated to individual server processes. This advisor is available starting with Oracle 9.2.

Recovery Page The Recovery page allows you to perform the following tasks:

- Use the Mean-Time-To-Recover (MTTR) Advisor to determine the optimal value for the maximum time needed to recover the database. This advisor is available starting with Oracle 9.2.
- View the current state of the redo log archival and take the database in archivelog or noarchivelog mode.

Resource Monitors Page The Resource Monitors page allows you to view the performance statistics of an active plan and of each consumer group associated with the active resource plan.

Undo Page The Undo Page contains information about the undo tablespace including the name of the active undo tablespace and the current undo retention time. Through this page, you can modify the retention time based on your largest transaction time and immediately view the space required for the undo tablespace. Undo generation rates are calculated based on statistics available for undo space consumption for the current instance. From this page, you can decide on an optimal size for your undo tablespace and ensure that even your longest transactions always complete.

Figure 10–2 Instance Management Window

Stored Configurations

When the database is connected to the Oracle Management Server, the Stored Configurations folder appears in the tree view with which you can create multiple database start-up configurations without the need to track initialization parameter files (INIT<SID>.ORA). Stored configurations exist in the Oracle Enterprise Manager repository (they are not external files) and can be created, edited, and deleted. You can also add and delete parameters and export a configuration to a file.

Note: If you are connected to an Oracle9.x database, you can also start up the database by using the SPFILE on the server side. The database knows the location of the SPFILE and will look for it when it starts up to find the startup parameters. An SPFILE is similar to an init.ora file but located on the server-side and maintained by the server.

Sessions List

The Sessions List page displays the top number of sessions that you specify using database instance resources in real time. Sessions are displayed in descending order based upon the delta value of the statistic chosen as the sort statistic. You can use the information in the chart to isolate executing SQL or to kill a problem session. The resumable sessions are highlighted.

Sessions Folder

The Sessions folder lists all users connected to the discovered database. When you select a user in the list, the Sessions property sheet appears with which you can edit user properties, view information about the status of each user, view current SQL or the last run SQL for the database session, and view the database session explain plan.

Long Running Operations

A small clock appears on the session icon in the tree view for sessions with currently running in-progress operations. Select the Long Operations tab on the top of a Session detail view to view the status of long-running operations on Oracle8i or 9i databases. You can monitor the type of operation it is, how long it has been running, and the estimated time of completion.

Locks

The Locks list contains information about the locks currently held by the Oracle server and outstanding requests for a lock or latch. Locks are mechanisms that prevent destructive interaction between transactions accessing the same resource--either user objects such as tables and rows or system objects not visible to users, such as shared data structures in memory and data dictionary rows. In all cases, Oracle automatically obtains necessary locks when executing SQL statements, so users need not be concerned with such details. Oracle automatically uses the lowest applicable level of restrictiveness to provide the highest degree of data concurrency yet also provide fail-safe data integrity. Oracle also allows you to lock data manually.

Note: Background sessions holding locks are not problematic and should not be killed.

In-Doubt Transactions

The In-Doubt Transactions folder contains information about distributed transactions that failed in the PREPARED state. You can sort the Transactions list on each of the columns by clicking on the column heading.

The In-Doubt Transactions property sheet displays information about distributed transactions in which a commit was interrupted by a system, network, or any failure resulting from external factors.

Resource Consumer Groups

The Resource Consumer Groups folder lists sets of users who have similar resource usage requirements. When you select a resource consumer group object in the folder, a property sheet appears in which you can view or specify properties and assign or remove users from the resource consumer group.

Resource Plans

The Resource Plans folder lists objects that represent resource plans, which are ways of allocating resources among consumer groups. Resource plans contain directives that specify the resources to be given to each group and can be specified in hierarchical fashion using subplans.

Note: The activated resource plans are highlighted in the navigator.

The Resource Plans property sheet, which appears when you select an object representing a Resource Plan, allows you to choose available groups/subplans to include in the resource plan, select the percentage of CPU resources allocated to a group, specify the maximum number of parallel execution servers associated with a single operation for each resource consumer group, specify the maximum number of concurrently active sessions allowed within a consumer group, specify a maximum in kilobytes on the total amount of undo generated by a consumer group, specify a maximum execution time in seconds allowed for an operation if there were no other work on the system, specifying criteria that causes the automatic switching of sessions to another consumer group, and then activate the plan.

Resource Plan Schedule

The Resource Plan Schedule property sheet allows you to automate when to activate a resource plan.

Note: The scheduling job is implemented using DBMS_JOB.

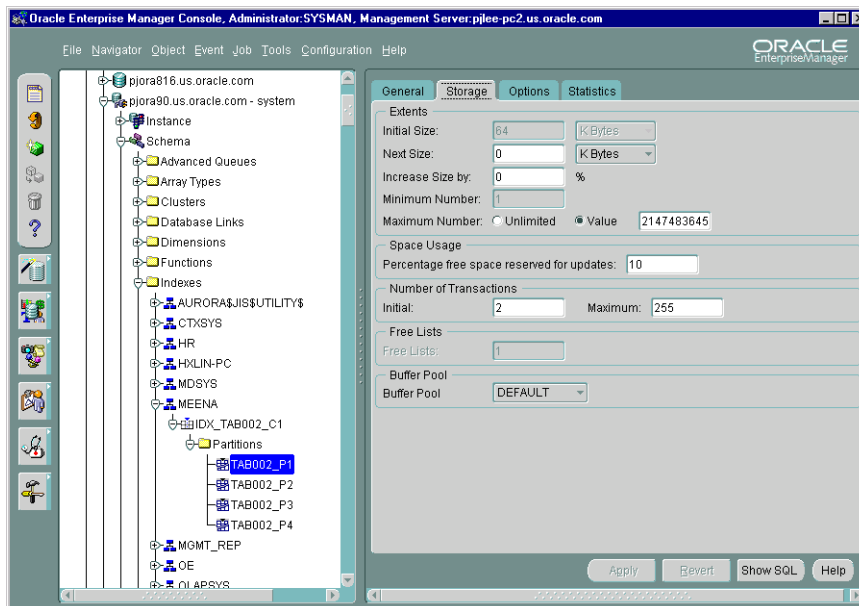
Schema Management

With the Schema Management feature, you can create, alter, or drop database schema objects such as clusters, indexes, materialized views, tables, and views, as well as view dependencies of schema objects. Storage layout information is also available for a table or an index if they are on EMC devices.

The Schema Management feature also supports index organized tables, partitioned tables and indexes, advanced queues, Java classes and sources, and unicode. Advanced queuing offers message transformation which can be used to transform and validate message communication amongst different business processes. The unicode feature allows you to select a column of "character" type and specify the length in bytes or characters.

You can also compile multiple objects such as functions, packages, package bodies, and triggers from their Summary View panel, and you can edit the storage and options information for clusters, indexes, and tables if the values apply to multiple objects.

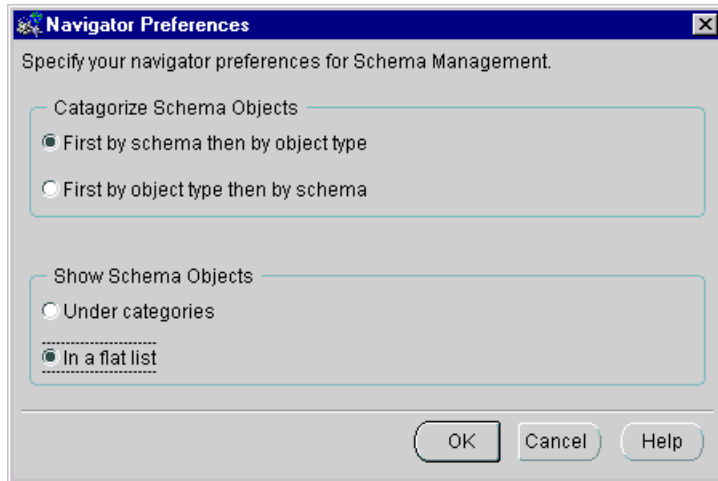
Figure 10–3 Schema Management Window



Tree List by Schema or Object

Beginning with 9.2, you can specify your navigator preferences for Schema Management.

Figure 10–4 Navigator Preferences



Databases contain at least one named schema for each database user. Regardless of object type, each schema object belongs to one of these named schemas.

If you need to edit several objects belonging to the same schema, select to view objects "First by schema and then by object type". Refer to Figure 10–5, "Schema Objects and Flat List" and Figure 10–6, "Schema Objects and Categories".

If you manage the same type of objects from different schemas, select to view them "First by object type and then by schema". Refer to Figure 10–7, "Objects and Flat List" and Figure 10–8, "Objects and Categories".

Depending on a second selection, the tree view will reorder all objects accordingly, presenting object folders in a categorized list or a flat list.

Figure 10–5 Schema Objects and Flat List

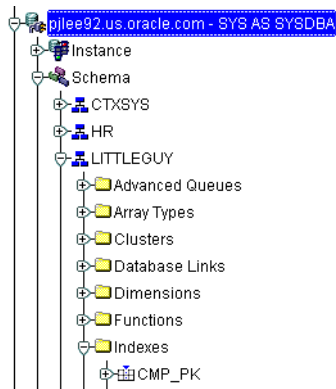


Figure 10–6 Schema Objects and Categories

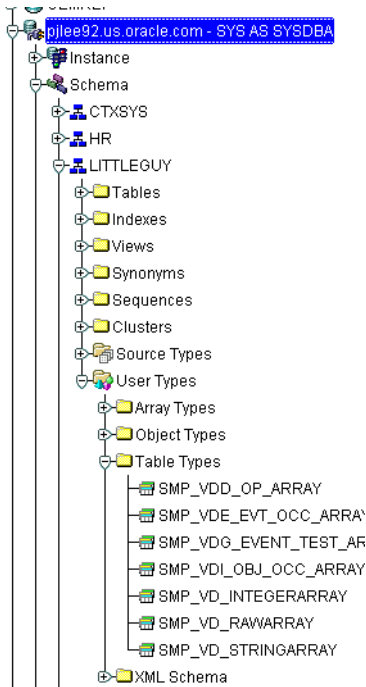


Figure 10–7 Objects and Flat List

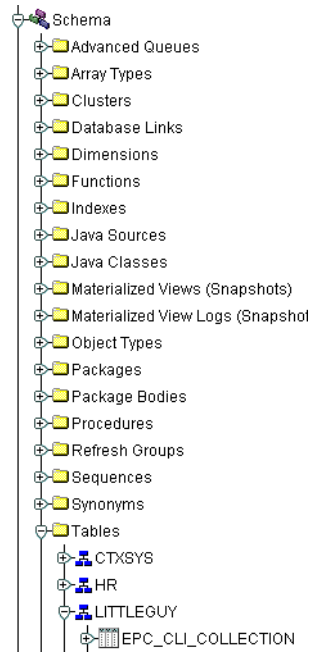
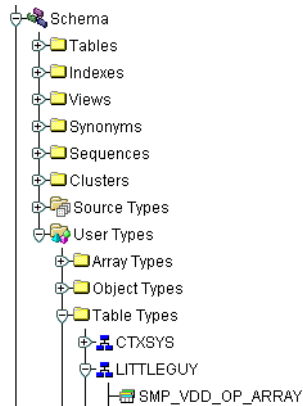


Figure 10–8 Objects and Categories

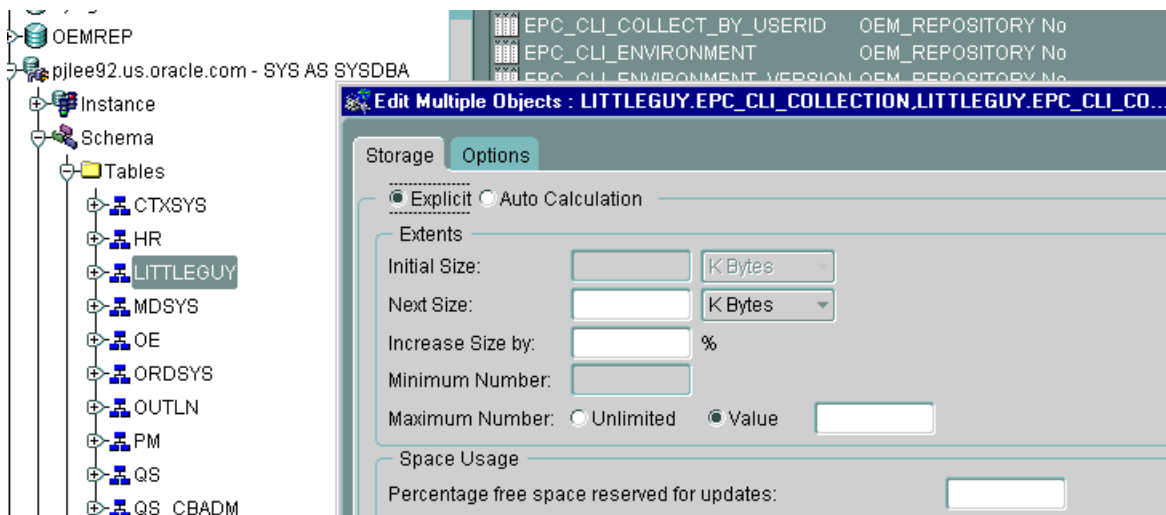


Editing an Object

To view/edit an object, you use the object's property sheet, which appears when you select the object in the tree view. You can then modify the object's parameters.

For clusters, indexes, and tables, you can use the Edit Multiple Objects feature to edit the storage and options information that apply to multiple objects at the same time.

Figure 10–9 *Editing Multiple Objects*



The Schema Management feature also includes the Table Data Editor content viewer, which allows you to view, update, and delete the contents of a table and display the contents of a view or synonym by selecting a right-mouse command on a table in the tree view.

Creating Objects

Schema Management allows you to create an object or a clone of an object by selecting Create or Create Like from the Object menu. When creating a clone of an object, all attributes are identical except for the name. Parameters for new objects and cloned objects are specified in property sheets which appear when you select Create and the object from the Create dialog or Create Like from the Object menu.

Security Management

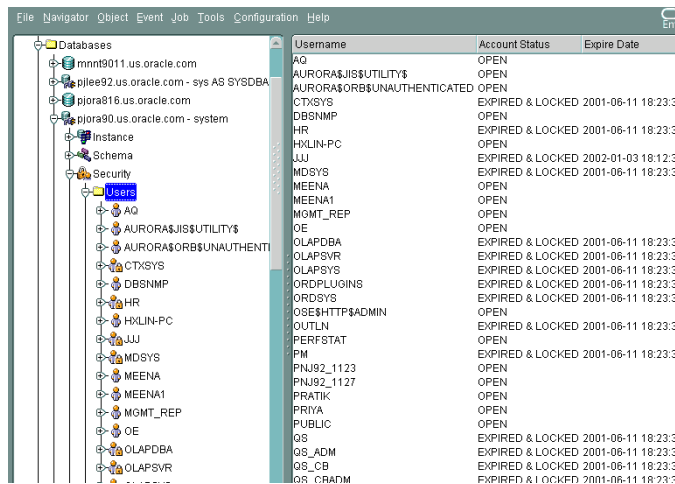
In a large network environment, security parameters for objects, administrators, and users are in constant change. With the Security Management feature, an administrator can make these necessary changes quickly and efficiently.

When you expand the Security node under the database in the tree view, folders for users, roles, and profiles appear.

Figure 10–10 Managing Users, Roles, and Profiles



Figure 10–11 Security Management Window



User Operations

The Security Management feature helps you manage the database users in your network by helping you create users and clones of users, add and remove user

permissions and roles, grant or revoke the switch privilege of resource consumer groups for a user or role, alter user properties, including account status and default profiles, and set up database users to act as proxy for a user. Security Management capabilities also allow you to easily see users' dependents and dependencies.

Role Operations

With the Security Management role operations feature, you can modify role properties as easily as user properties. You can also create roles and clones of roles, add and remove permissions from roles, and see grantees of roles, including consumer groups.

Profile Operations

A profile is a set of limits on a user's database resources. As with users and roles, you can create a profile or a clone of a profile, alter a profile's properties, and assign and remove profiles from users. You also have the ability to see profile dependents and dependencies, as well as grantees of profiles.

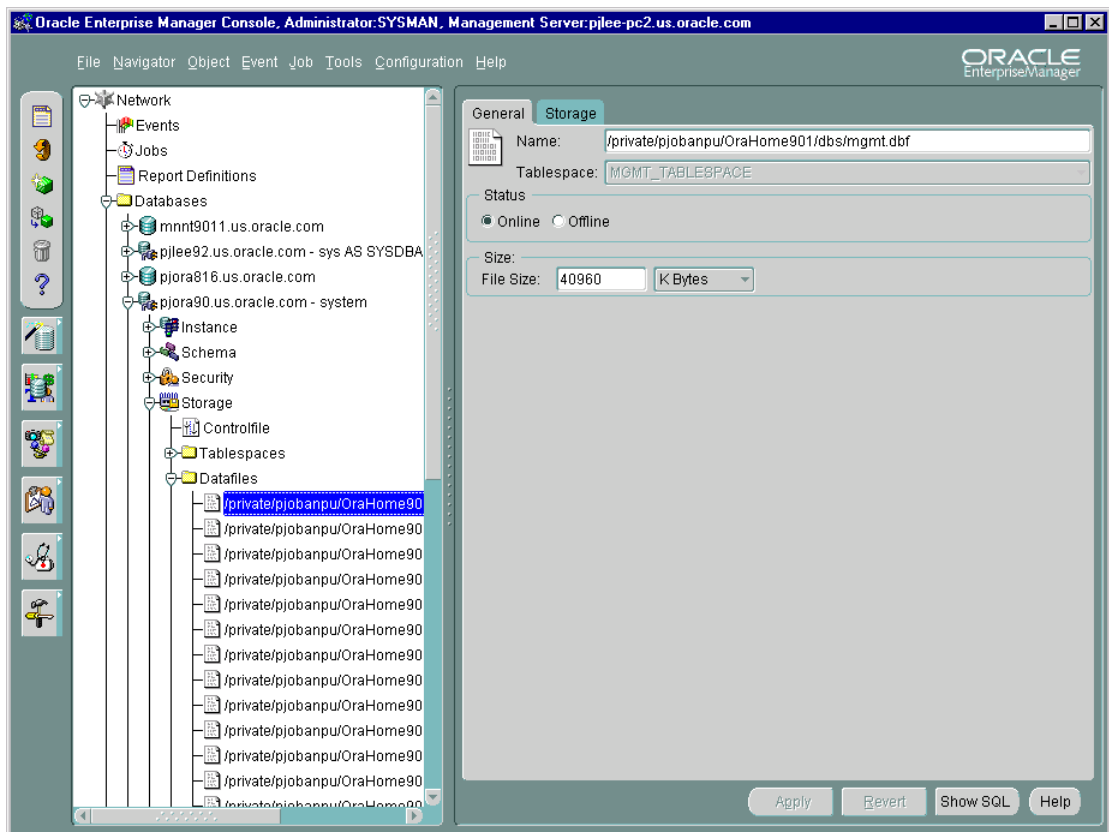
Profiles, roles, and the users to which they are assigned can easily be seen in security lists provided by Security Management. Administrators can then use property sheets to determine security parameters, simplifying the process of making changes.

The Security Management feature also supports Oracle password management, which increases system security. Supported features includes: account locking, password lifetime and expiration, password history, password complexity, verifications, and export/import of passwords.

Storage Management

The Storage Management feature helps you administer tablespaces (permanent, temporary, and undo), datafiles, redo logs, archive logs, and rollback segments for optimum database storage.

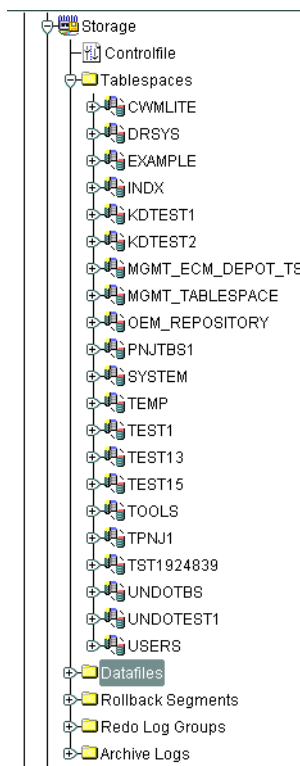
Figure 10–12 Storage Management Window



When you connect to a database, the Storage Management branch of the tree view lists an icon for the Controlfile and five folders which contain all the storage objects in the selected database. The five folders include:

- Tablespaces
- Datafiles
- Rollback Segments
- Redo Log Groups
- Archive Logs

Figure 10–13 *Managing Database Storage Parameters*



The following sections describe the Storage Management operations that can be performed with the objects in each of these folders:

Controlfile Operations

When you select the Controlfile icon, a property sheet appears where you can see the number of controlfiles created for the database and other statistics.

Tablespace Operations

Using the contents of the Tablespaces folder, you can create, edit, or drop a tablespace, switch to a new default temporary tablespace, add a datafile or rollback segment, take a tablespace off- or on-line, make a tablespace read-only or writable, and set tablespace storage parameters, including multiple block sizes which is key in facilitating the transportability of tablespaces from one database type to another.

Beginning with Enterprise Manager 9.2, you can choose to use bitmaps to manage the free space within segments. Bitmaps allow Oracle to manage free space more automatically, and offers high performance for free space management.

In addition, you can click the tablespace to see the used and free space of the tablespace or datafile.

Beginning with Oracle9i, you can also allocate your undo space in a single undo tablespace, instead of distributing them into a set of statically allocated rollback segments. For each Oracle instance, you will only have to allocate enough disk space for the workload in that instance in an undo tablespace. In addition, using Enterprise Manager, you can create or alter an undo tablespace.

Datafile Operations

With the contents of the Datafiles folder, you can create a datafile or a clone of a datafile, edit a datafile, and take a datafile off- or on-line. You can also click the datafile to see the used and available space. In 9.2, you can delete datafiles when dropping a tablespace. Also, The storage layout feature, which includes a page to display the storage layout information for a datafile and a page to display the overview of the file map, is available for datafiles on EMC devices.

Rollback Segment Operations

Using the Rollback Segments folder, you can create, alter, drop, or shrink a rollback segment as well as take one off- or on-line.

Redo Log Group Operations

Using the Redo Log Groups folder, you can switch the current redo log group, trigger a checkpoint in a redo log group, create a new redo log group, and rename, remove, or add new redo log group members.

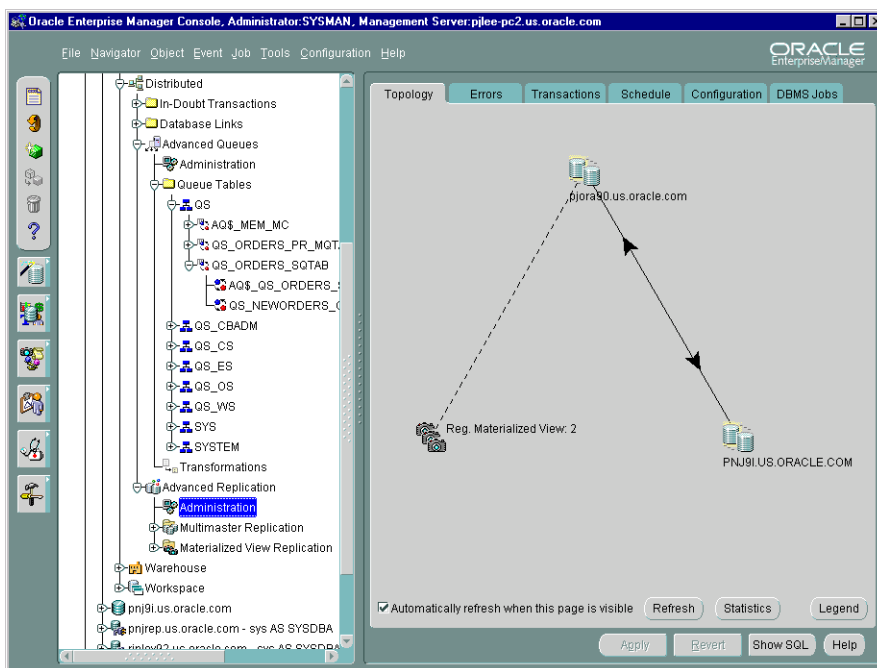
Archive Log Operations

The archive log folder allows you to view the current archive logs in the database.

Distributed Management

Oracle supports data replication and messaging technologies to support distributed applications and distributed database systems.

Figure 10–14 *Distributed Database Management*

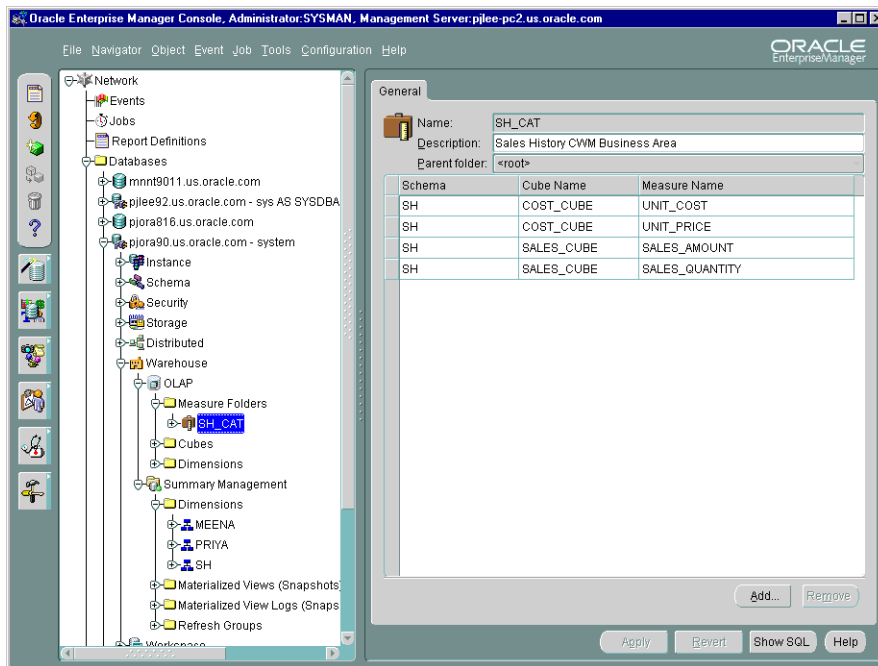


With Distributed Database Management, you can:

- use the In-Doubt Transactions folder to view failed two-phase commit transactions.
- use Database Links to create, edit and delete links between databases.
- use the Streams folder to replicate data over a network of databases. You can monitor and administer your Streams environment. This folder is available for Oracle version 9.2 and above.
- use the Advanced Queues folder to integrate applications by using messaging technology. You can monitor and administer your Advanced Queues environment.
- use the Advanced Replication folder to replicate data amongst a group of databases (Multi Master Replication) or to replicate a snapshot of data between databases (Snapshot Replication). You can to monitor and administer your Advanced Replication environment.

Warehouse Management

A data warehouse is a relational database that is designed for query and analysis rather than for transaction processing. Oracle supports complex analysis of warehouse data by on-line analytical processing (OLAP) applications. Oracle also provides mechanism to improve performance of your warehouse by using summaries.

Figure 10–15 Warehouse Management

OLAP Management

OLAP Management enables an administrator to create and edit OLAP catalog metadata for an existing star or snowflake schema. OLAP catalog metadata is required by business intelligence applications that use Oracle's Java-based OLAP API.

Within the OLAP catalog repository, owned by OLAPSYS, there are two sets of metadata tables, each with its own write APIs (PL/SQL packages): CWMLite Release 1 and CWMLite Release 2.

The OLAP Management interface within Enterprise Manager uses CWMLite Release 1 only. CWMLite Release 2 metadata must be created programmatically using the CWM2_OLAP supplied PL/SQL packages.

Both CWMLite 1 and CWMLite 2 are supported by the OLAP API.

For more information, see *Oracle9i OLAP User's Guide*.

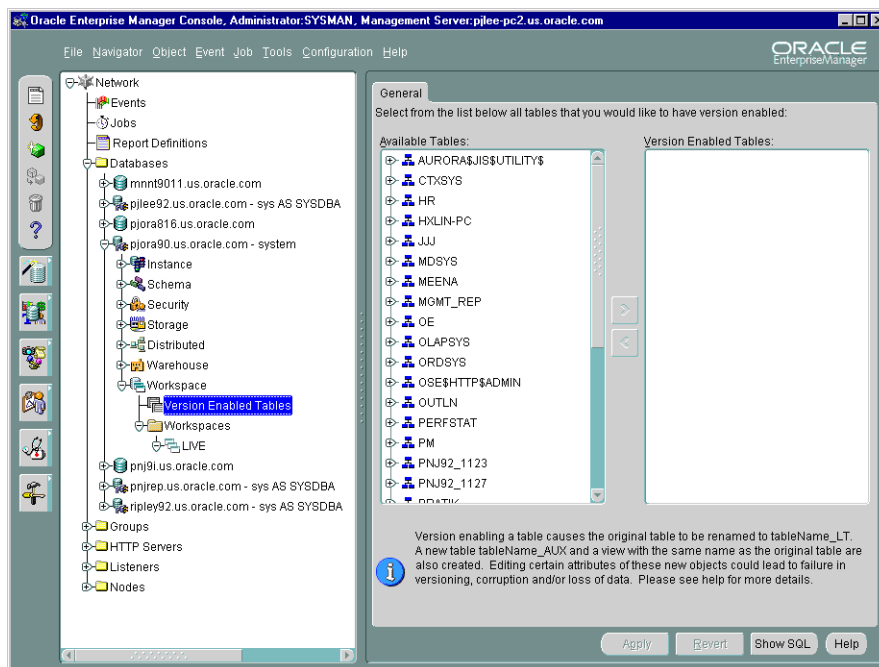
Summary Management

Oracle supports Materialized Views, also known as Summaries, to improve the performance of a data warehouse. Summaries pre-calculate expensive joins and aggregation operations and store them in a table. Oracle speeds-up query execution by transparently re-writing queries to use summaries. Dimensions aid query re-write operation. MV Logs track changes to master table for incremental refresh of Materialized Views.

Workspace Management

Workspace Management allows you to version-enable tables and create, modify, refresh, and merge workspaces.

Figure 10–16 Workspace Management



Oracle Workspace Manager provides a long transaction framework, in which multiple data versions are stored in the database as different workspaces. You can create new versions of data to update, while maintaining a copy of the old data.

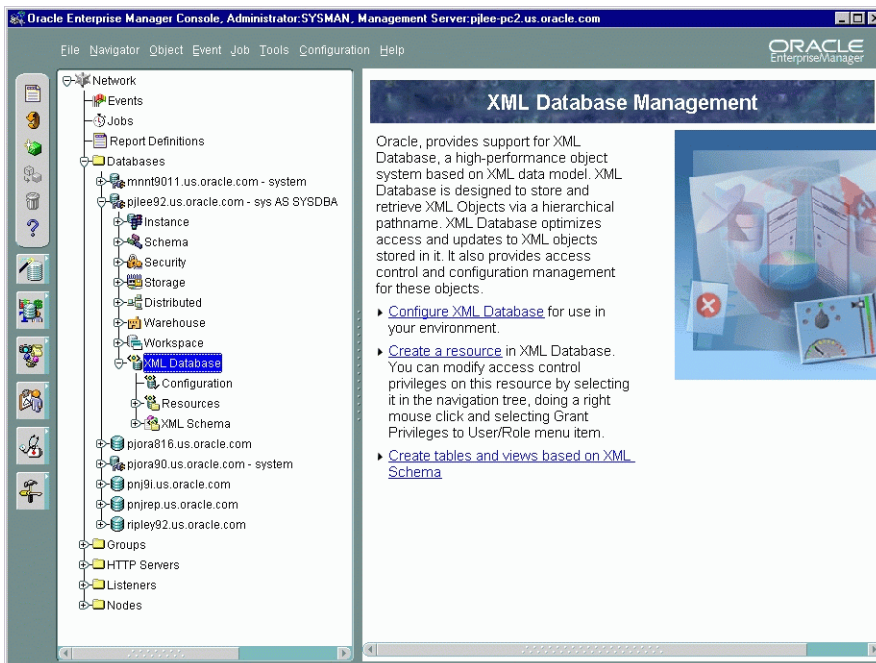
Versioning improves concurrent access of data in the database and allows multiple what-if analyses to be run against the data simultaneously.

For detailed information about Workspace Manager concepts and the application programming interface (API), see *Oracle9i Application Developer's Guide - Workspace Manager*.

XML Database

Oracle provides support for an XML Database, a high-performance object system based on the XML data model. An XML Database is designed to store and retrieve XML Objects via a hierarchical pathname. It also optimizes access and updates to XML objects stored in it and provides access control and configuration management for these objects.

Figure 10–17 XML Database Management



From the XML Database container, you can perform the following tasks:

- Configure XML Database for use in your environment.
- Create a resource in XML Database. You can modify access control privileges on this resource by selecting it in the navigation tree, doing a right mouse click and selecting Grant Privileges to User/Role menu item.
- Create tables and views based on XML Schema

SQL*Plus Worksheet

When you need to administer your database environment with SQL, PL/SQL, or SQL*Plus commands, use Oracle SQL*Plus Worksheet. With Oracle SQL*Plus Worksheet, you can enter SQL and PL/SQL code and DBA commands dynamically and run scripts which are stored as files.

The SQL*Plus Worksheet window consists of an Input pane (top) where commands are entered, and an Output pane (bottom) where the results of your commands appear after you click the Execute button.

SQL*Plus Worksheet maintains a history of the commands you have entered, allowing you to edit and re-execute an earlier command without having to retype it. The last 50 command executions can be displayed by clicking the Command History button. Selections from the Command History dialog box can then be copied and inserted into the Input pane.

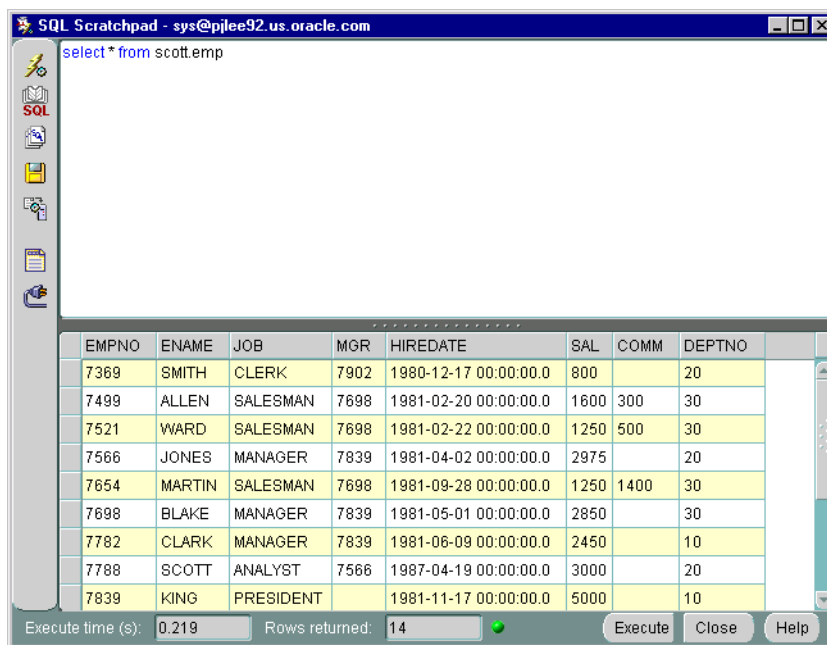
With SQL*Plus Worksheet, you can have multiple copies of the worksheet open at a time, each of which is separate from the others; so work can be committed or rolled back in each worksheet independently.

Note: Additional information on the database administration features and wizards in the Oracle Enterprise Manager can be found in the Oracle Enterprise Manager Online Help.

SQL Scratchpad

Oracle Enterprise Manager now features the SQL Scratchpad, which provides a user interface for you to enter, edit, and execute SQL quickly and easily.

Figure 10–18 SQL Scratchpad

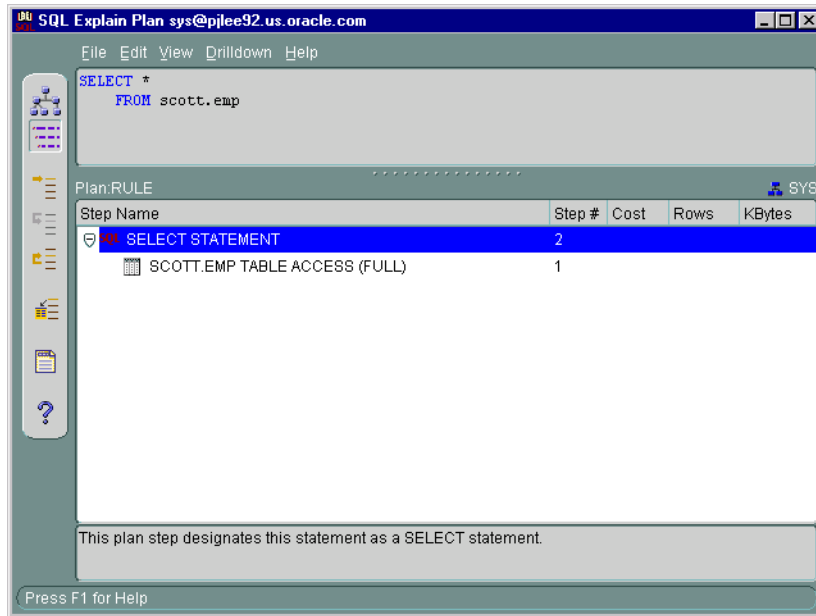


Features of SQL Scratchpad are listed below:

- Low overhead
- Standard editor features: cut, copy, paste
- Full color syntax highlighting
- Load and save SQL

- Graphical Explain Plan which describes the steps chosen by the optimizer for executing the SQL statement.

Figure 10–19 Explain Plan



- View and sort results of executed SQL command as data in a read-only spread table
- Displays the time taken to execute the queries
- Displays the number of rows for select statement
- Cancel SQL execution, which is especially useful for long running queries
- Connect to the database as a different user, including connection with SYSDBA role

Wizards

For help with database administration tasks, Oracle Enterprise Manager offers a variety of wizards:

- **Analyze Wizard:** The Analyze wizard allows you to collect statistics about the objects and store them in the data dictionary, delete statistics about the objects from the data dictionary, validate the structure of the objects, and identify migrated and continued rows of tables or clusters.
- **Backup and Recovery Management Wizards:** The Backup and Recovery wizards are also available to help you back up or restore and recover various objects such as the tablespaces, datafiles, or archivelogs. With the Backup wizard you can also make an image copy of the datafiles and the current controlfile. Beginning with Oracle Enterprise Manager 9.2, the Backup Wizard allows the setting of additional options, such as backup retention policy, deleting obsolete backups and specifying the archivelog deletion policy. With Oracle 9.2, recovery includes Block Media Recovery which improves the speed of recovery significantly in the case of block corruptions.
- **Create Table Wizard:** The Create Table Wizard facilitates the creation of a table.
- **Cube Wizard:** The Create Cube Wizard helps you build a cube object. Cubes represent multidimensional data stored in your data warehouse fact tables.
- **Data Management Wizards (Import/Export/Load):** Oracle Data management wizards automate the transfer of data to and from an Oracle database.
- **Dimension Creation Wizard:** The Create Dimension Wizard will help you build a dimension object. Dimensions represent columns in your data warehouse dimension tables as levels and attributes. Dimensions typically define hierarchical relationships between their levels.
- **Resource Plan Wizard:** The Resource Plan Wizard helps you group user sessions that have similar processing and resource usage requirements and allocate resources among the consumer groups.
- **Summary Advisor Wizard:** The Summary Advisor Wizard provides advice as to which materialized views should be created, dropped, or retained.
- **View Wizard:** The View Wizard facilitates the creation of a view, a tailored presentation of the data contained in one or more tables (or other views).

For more information on wizards, see the Oracle Enterprise Manager Quick Tour or the Oracle Enterprise Manager Online Help.

Managing Backup and Recovery

Introduction

Recovery Manager (RMAN) is an Oracle utility that can back up, restore, and recover database files. It is a feature of the Oracle database server and does not require separate installation. The Oracle Enterprise Manager Backup Management wizards and property sheets provide a graphical user interface to Recovery Manager. This chapter describes how to use Oracle Enterprise Manager to administer your database backup and recovery environment.

Recovery Manager (RMAN)

Recovery Manager uses database server sessions to perform the work of backup and recovery. It stores metadata about its operations in the control file of the target database and, optionally, in a recovery catalog schema in an Oracle database.

You can invoke RMAN as a command-line executable from the operating system prompt or use RMAN features through the Enterprise Manager GUI. The Oracle Enterprise Manager Backup Management wizards and property sheets provide a graphical user interface to Recovery Manager.


The Recovery Manager environment consists of the various applications and databases that play a role in a backup and recovery strategy. The RMAN environment can be as simple as an RMAN executable connecting to a target database, or as complex as an RMAN executable connecting to multiple media managers and multiple target, recovery catalog, and auxiliary databases, all accessed through Oracle Enterprise Manager.

Possible components of the RMAN environment are listed below.

- RMAN executable
- Target database
- Oracle Enterprise Manager
- Recovery catalog database
- Recovery catalog schema
- Standby database
- Media management application
- Media management catalog

Of these components, only the RMAN executable and target database are required. RMAN automatically stores its metadata in the target database control file, so the recovery catalog database is optional. Nevertheless, maintaining a recovery catalog is strongly encouraged. If you create a catalog on a separate machine, and if the production machine fails completely, then you have all the restore and recovery information you need in the catalog.

Figure 11-1 Example RMAN Environment



The figure above depicts an example of a realistic RMAN environment. In this environment, five nodes are networked together, with each machine serving a different purpose.

The five nodes share duties as follows:

- One client node runs the RMAN executable
- One server node hosts the target database and media management subsystem
- One server node hosts the duplicate or standby database
- One server node hosts the recovery catalog database
- One client node runs the Oracle Enterprise Manager application, which provides a GUI interface to the databases in the system

In this scenario, you can run the RMAN executable from a client machine, and then connect to the target, catalog, and auxiliary databases. You can then run backup and recovery jobs. You can also connect to the client hosting Oracle Enterprise Manager and use the Oracle Enterprise Manager to access RMAN.

About the RMAN Executable

RMAN is the client application that manages backup and recovery operations for a target database. The RMAN executable is automatically included with the Oracle software installation. The RMAN client uses Oracle Net to connect to a target database, so it can be located on any host that is connected to the target host through Oracle Net.

About the Target Database

The target database is made up of the control files, datafiles, and optional archived redo logs that RMAN is in charge of backing up, restoring, or recovering. RMAN uses the target database control file to gather information about the database and to store information about its own operations. The actual work of the backup and recovery jobs is performed by server sessions on the target database. You can use a recovery catalog to manage the metadata of the database.

About Oracle Enterprise Manager

You can use Oracle Enterprise Manager as an interface to RMAN. Access the wizards and property sheets using one of the following methods:

- From the Console Navigator, select the database you want to administer; then, choose the tool from the context-sensitive Backup Management menu.

- From the Storage container of the Console Navigator, select the tablespace or datafile you want to administer; then, choose the tool from the context-sensitive Backup Management menu.
- From the Console's Object->Backup Management menu.
- From the Console's Tools>Database Tools->Backup Management menu.
- From the Storage and Instance HTML pages, you can launch the Backup, Recovery and Maintenance wizards.

Note: The Backup Management wizards and property sheets are only available when you are connected to a Management Server.

The Backup Management wizards and property sheets consist of:

- **Backup**
The Backup Wizard helps you to back up various objects such as the database, datafiles, tablespaces, and archivelog, or to make an image copy of the datafiles and the current controlfile.
- **Recovery**
The Recovery wizard helps you restore and recover various objects like databases, datafiles, and tablespaces. It guides you through the process of specifying what you want to restore and recover and submits a recovery job through the Enterprise Manager to complete the operation.
- **Maintenance**
The Maintenance Wizard helps you perform maintenance operations on the target databases and on the recovery catalog. Using this wizard, you can set up backup and retention policies in a target database, or register, reset or resynchronize the target database with the recovery catalog.
- **Create Backup Configuration**
Enterprise Manager creates a default backup configuration for each target database, but you can use the Create Backup Configuration property sheets to create other backup configurations for backup and recovery. A configuration can be used for one database or many databases depending if the systems are the same.
- **Backup Configuration Library**
The Backup Configuration Library page displays backup configurations.

About the Recovery Catalog Database

A recovery catalog database is a database containing the recovery catalog schema, which contains the metadata that RMAN uses to perform its backup and recovery operations.

About the Recovery Catalog Schema

The Recovery Catalog Schema is the user within the recovery catalog database that owns the metadata tables maintained by RMAN. RMAN uses these metadata tables to store information about the target database and its backup and recovery operations. Among other things, RMAN stores information about:

- Backup sets and pieces
- Image copies
- Proxy copies
- Archived redo logs
- The target database schema
- Persistent configuration settings

You can either use a recovery catalog in which to store the repository, or let RMAN store the repository exclusively in the target database control file.

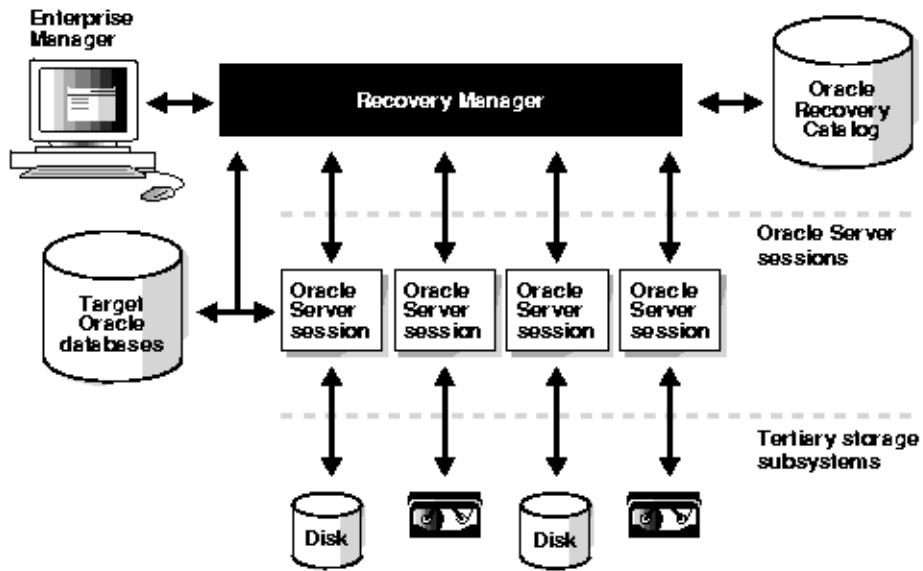
Although RMAN can conduct all major backup and recovery operations using just the control file, note these advantages of using the catalog:

- Some RMAN commands and operations function only with a catalog.
- The recovery catalog retains historical backup information that can get overwritten in the control file.
- The recovery catalog stores information about backups from different incarnations of the database.

The recovery catalog is maintained solely by RMAN; the target database never accesses it directly. RMAN automatically propagates information about the database structure, archived redo logs, backup sets, and datafile copies into the recovery catalog from the target database's control file.

For Oracle9i and later, a recovery catalog is created if you specify for the Enterprise Manager repository to be located in a local database. The recovery catalog will be created in the `CATTS` tablespace for you by default with the recovery catalog user and password of `rman/rman`.

Figure 11–2 Recovery Catalog



Important: The recovery catalog and the Oracle Enterprise Manager repository should not reside in the target database (database to be backed up). The recovery catalog can reside in the same database as your Oracle Enterprise Manager repository. Oracle recommends placing the recovery catalog in a separate tablespace. As with any important data, you should back up your recovery catalog regularly.

To use Recovery Manager with a recovery catalog, you must register your database with the recovery catalog. Refer to "Registering the Recovery Catalog" for more information. No setup is required if you are using the control file.

About the Standby Database

A standby database is a copy of the primary database that is updated using archived logs created by the primary database. RMAN can create or back up a standby database.

About the RMAN Media Management Interface

To store backups on tape, RMAN requires a media manager, which is a vendor-specific application. A media manager is a software program that loads, labels, and unloads sequential media such as tape drives used to back up and recover data. For information on configuring RMAN to make backups to a media manager, refer to the *Oracle9i Recovery Manager User's Guide*.

When doing backups or restores, the RMAN client connects to the target instance and directs the instance to talk to its media manager. No direct communication occurs between the RMAN client and the media manager: all communication occurs on the target instance.

About the Media Management Catalog

A media management catalog is a vendor-specific repository of information about a media management application.

RMAN Backup

A backup is a copy of data. This copy can include important parts of the database such as the control file and datafiles. A backup is a safeguard against unexpected data loss and application errors. If you lose the original data, then you can reconstruct it by using a backup.

This section contains the following topics:

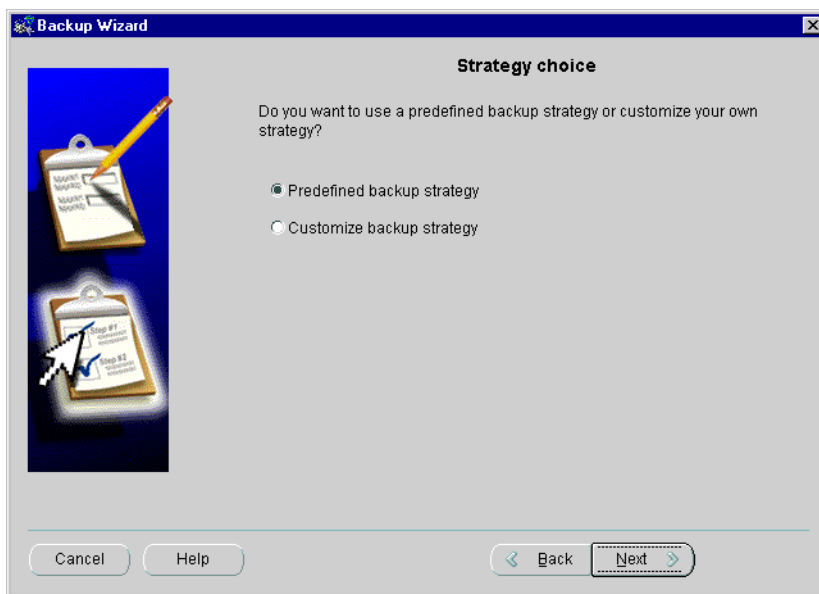
- Backing Up a Database
 - Backing Up a Database Using a Predefined Strategy
 - Backing Up the Database with a Customized Strategy
 - Deleting Obsolete Backups and Copies
 - Selecting a Full or Incremental Backup
 - Choosing an Online or Offline Mode to Back Up Your Database
- Backing Up Individual Files
- Backing Up and Deleting Archived Logs
- Copying a Datafile
- Overriding a Retention Policy in Backup for Special Cases

Backing Up a Database

To back up a database

1. From the Backup Management menu, choose **Backup** to access the Backup Wizard.
2. On the Strategy Choice page, select either **Predefined backup strategy** or **Customize backup strategy**.

Figure 11-3 Predefined Strategy in Strategy Choice Page

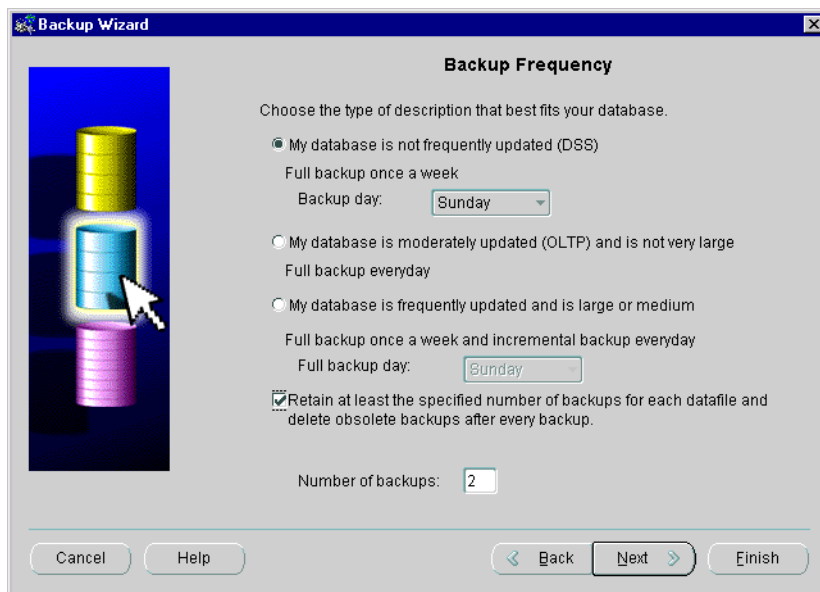


For more information on using a customized backup strategy, see "Backing Up the Database with a Customized Strategy" on page 11-11.

Backing Up a Database Using a Predefined Strategy

Select **Predefined backup strategy** on the Strategy choice page of the Backup Wizard if you want to back up your entire database without having to make too many decisions. The Backup Frequency page appears with general descriptions from which you can choose.

Figure 11–4 Backup Frequency



Picking a Description that Fits Your Database. Pick the description that fit your database in the Backup Frequency page, and RMAN will decide how often to perform a backup based on the general description that you pick.

Specifying the Number of Backups to Retain. If the selected (target) database is Oracle 9i and later, the **Retain at least the specified number of backups for each datafile and delete obsolete backups after every backup** checkbox and the **Number of backups** field are enabled in the Backup Frequency page.

The checkbox is selected by default. The default value is 2 for the number of backups to retain in the **Number of backups** field.

With this default selection, the retention policy of the target database will be set to redundancy 2. At least the 2 most recent full backups will be retained for each datafile. Older backups will be deleted after each new backup is successfully performed.

Later Steps in the Predefined Strategy. Later, in the process, wizard pages will appear in which you will have the option to perform the following tasks:

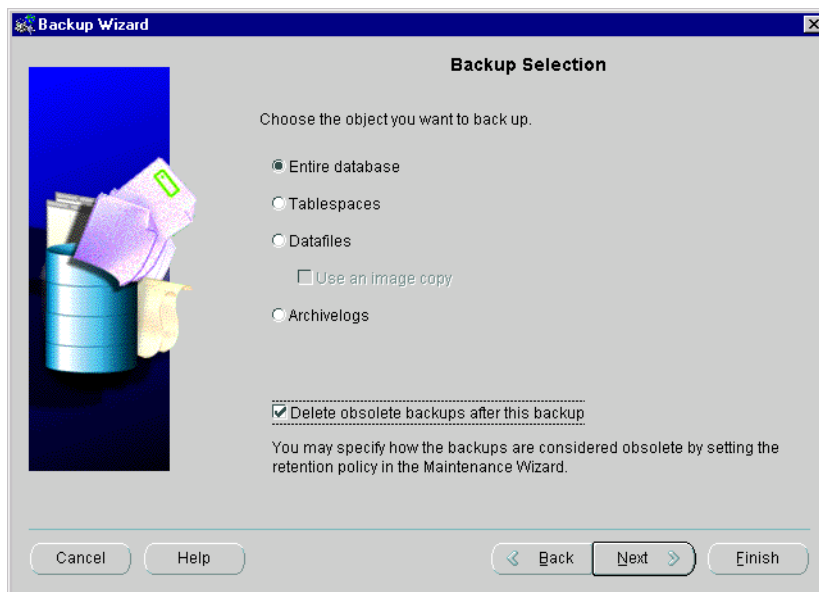
- Specify the start time for performing the backup.
- Select the default configuration or any of the user created configurations for the backup. A default configuration is created by Enterprise Manager for each target database.
- Specify multiple targets to submit the job. Note: To submit the same backup to multiple targets, the databases must have the same structure and the same disk/tape configuration.

Backing Up the Database with a Customized Strategy

Select **Customize backup strategy** on the Strategy choice page of the Backup Wizard if you want to select the information you want to back up and the schedule for the execution of the backup.

In order to back up the whole database you must select **Entire database** on the Backup Selection page.

Figure 11-5 Backup Selection



If the target database is 9i or above and the retention policy has been set in the target database, you can also choose to delete obsolete backups after the backup.

For more information, see "Deleting Obsolete Backups and Copies" on page 11-12.

Later in the process, wizard pages will appear in which you will have the option to perform the following tasks:

- Choose whether to include archive logs in the backup and if they should be deleted after each backup.
- Select a full or an incremental backup.
- Choose the online backup or offline backup mode.
- Select the default configuration or any of the user created configurations for the backup. A default configuration is created by Enterprise Manager for each target database.
- Choose to override the backup and retention policy.
- Schedule the execution of a backup.
- Select when to submit the job and whether to add it to the job library.
- Specify multiple targets to submit the job. Note: To submit the same backup to multiple targets, the databases must have the same structure and the same disk/tape configuration.

Deleting Obsolete Backups and Copies

If you are using a customized backup strategy and if the target database is 9i or above and the retention policy has been set in the target database, you can choose to have obsolete backups and copies deleted after the backup.

The retention policy determines which backups and image copies are obsolete. The current retention policy setting may be viewed through the Maintenance Wizard on the "Retention Policy" page. The retention policy may be modified through the Maintenance Wizard by submitting an Enterprise Manager job.

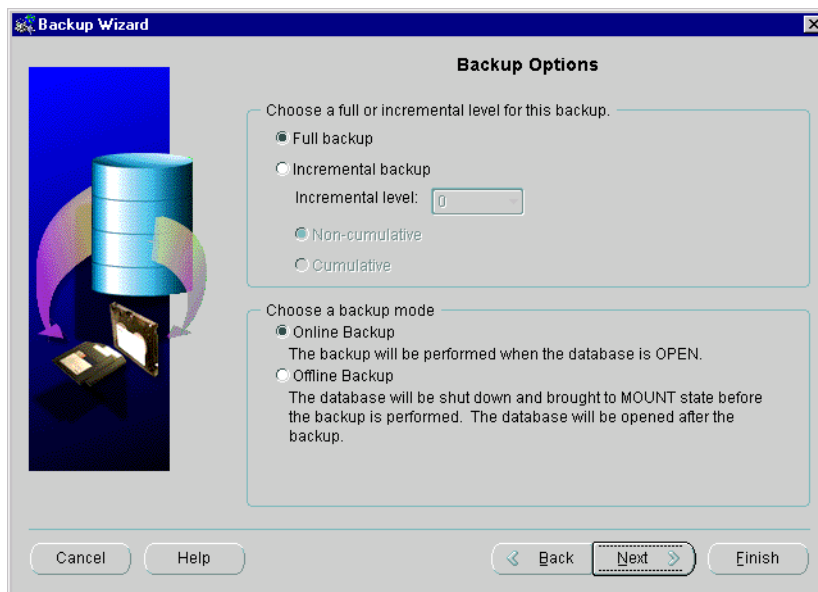
Select **Delete obsolete backups after this backup** on the Backup Selection page of the Backup Wizard. The retention policy determines which backups and image copies are obsolete. If selected, the obsolete backups and copies will be deleted when the backup is finished.

Selecting a Full or Incremental Backup

If you are using a customized strategy for your backup, you can choose to perform a full or an incremental level backup.

Select **Full backup** or **Incremental Backup** in the Backup Options page of the Backup Wizard.

Figure 11–6 Backup Options



A full backup A full backup backs up all blocks into the backup set, skipping only datafile blocks that have never been used. The server process does not skip blocks when backing up archived redo logs or control files. A full backup has no effect on subsequent incremental backups, which is why it is not considered part of the incremental strategy. In other words, a full backup does not affect which blocks are included in subsequent incremental backups.

An incremental backup Incremental backups are a convenient way to conserve storage space because they back up only database blocks that have changed.

The primary reasons for making an incremental backup are

- To save tape when using a media manager or disk space when making disk backups
- To save network bandwidth when backing up over a network

- When the aggregate tape bandwidth available for tape write I/Os is much less than the aggregate disk bandwidth for disk read I/Os
- To be able to recover changes to objects created with the NOLOGGING option (direct load inserts do not log redo, although they do change data blocks and so are captured by incremental backups)
- To reduce backup sizes for NOARCHIVELOG databases. Instead of making a whole database backup every time, you can make incremental backups. Note that incremental backups of a NOARCHIVELOG database are only legal after a consistent shutdown.

Incremental backups are a method by which you only backup modified blocks. An incremental level 0 backup performs the same function as a full backup in that they both backup all blocks that have ever been used except a level 0 will affect what blocks are copied out by subsequent incremental backups. Incremental backups of levels greater than 0 backup only blocks that have changed since previous incremental backups. Blocks which have not changed will not be backed up.

When you choose to make an incremental backup, you can choose a non-cumulative or a cumulative backup.

A non-cumulative backup is a type of incremental backup in which you back up all blocks that have changed since the most recent backup at level n or lower. For example, in a differential level 2 backup you back up all blocks modified since the last level 2, level 1, or level 0 backup. A non-cumulative backup copies less data and therefore takes a shorter time than the cumulative backup, but recovery time may be longer based on the number of incremental backups that must be applied.

A cumulative backup is a type of incremental backup that allows you to back up all the blocks used since the most recent backup at level n-1 or lower. For example, in a cumulative level 2 backup you back up all blocks used since the most recent level 1 or level 0 backup. A cumulative backup copies more data and therefore takes longer than the non-cumulative backup, but recovery time is shorter.

Incremental Backup Strategy Choose a backup scheme according to an acceptable MTTR (mean time to recover). For example, you can implement a three-level backup scheme so that a full or level 0 backup is taken monthly, a cumulative level 1 backup is taken weekly, and a cumulative level 2 is taken daily. In this scheme, you never have to apply more than a day's worth of redo for complete recovery.

When deciding how often to take full or level 0 backups, a good rule of thumb is to take a new level 0 whenever 50% or more of the data has changed. If the rate of change to your database is predictable, then you can observe the size of your incremental backups to determine when a new level 0 is appropriate.

Choosing an Online or Offline Mode to Back Up Your Database

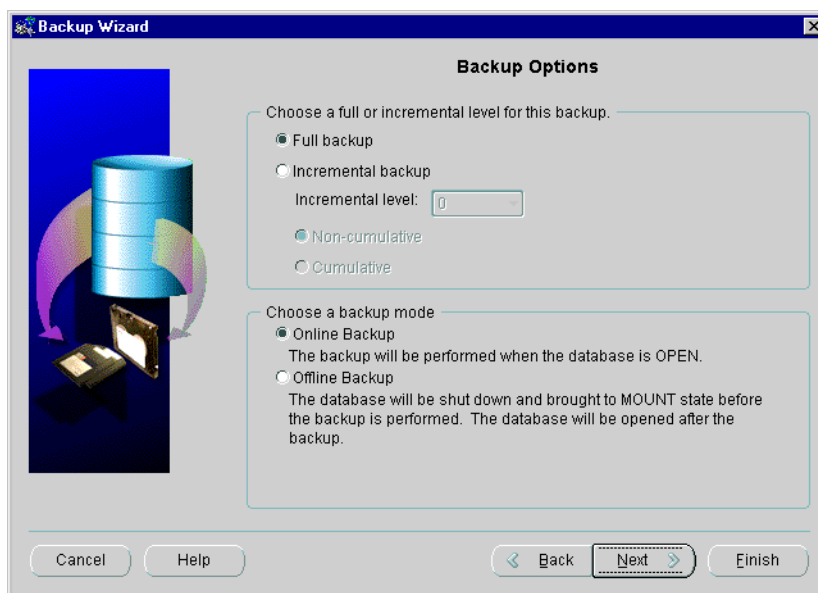
If you are making a database backup using a customized strategy and the target database is in ARCHIVELOG mode, you can choose to make an online or an offline backup.

Select **Online Backup** or **Offline Backup** in the Backup Options page of the Backup Wizard.

An online backup is a backup of one or more datafiles taken while a database is open and the datafiles are online.

An offline backup is a backup when the database is not open.

Figure 11–7 Backup Options



Online Backup is the default selection. If the database is in the OPEN state, the database backup will be performed while the database is OPEN.

If you choose **Offline Backup**, the database will be backed up in the MOUNT state. If the database is in the OPEN state, it will be shut down and brought to the MOUNT state before the backup is performed. When the backup is finished, the database will be brought back to OPEN state.

Backing Up Individual Files

Your database contains a wide variety of types of data. When developing your backup strategy, you must decide what information you want to back up. The basic principle you should use when deciding what to back up is to prioritize data depending on its importance and the degree to which it changes.

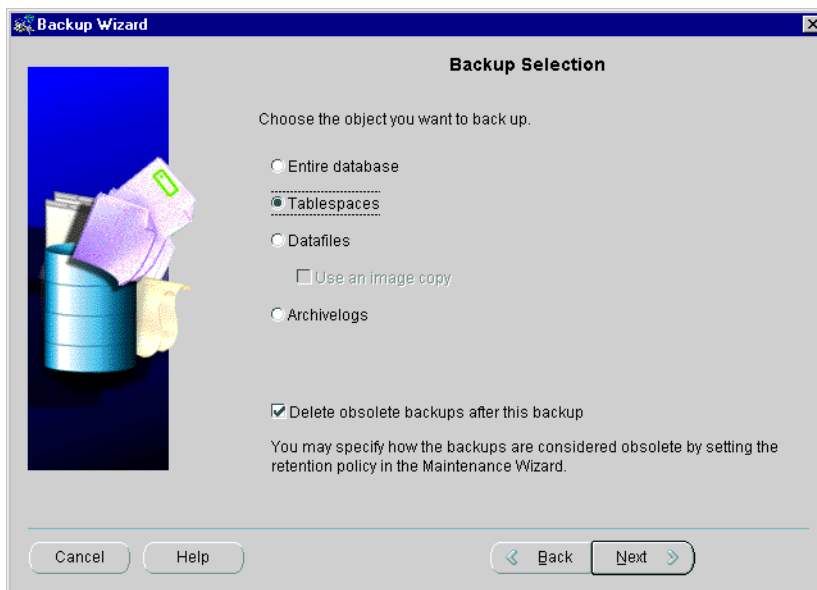
You can backup up individual files with various options.

1. From the Backup Management menu, choose **Backup** to access the Backup Wizard.
2. On the Strategy Choice page, select **Customize backup strategy**.
3. On the Backup Selection page, select **Datafile** or **Tablespace**.

To back up datafiles or tablespaces, the database must be in ARCHIVELOG mode and the Mount State.

See "Starting Up the Database" and "Setting the Database in ARCHIVELOG Mode" for information on starting the database in Mount mode or putting the database in ARCHIVELOG mode.

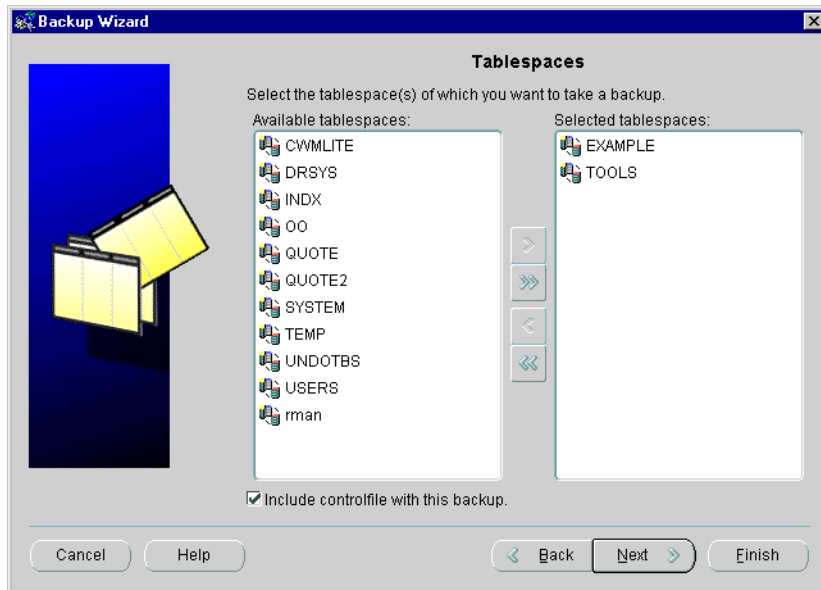
Figure 11–8 Select Tablespace in Backup Selection



If the target database is 9i or above and the retention policy has been set in the target database, you can also choose to delete obsolete backups after the backup.

For more information, see "Deleting Obsolete Backups and Copies" on page 11-12.

4. Choose tablespaces from the Tablespace page or datafiles from the Datafile page. You may choose to include the Controlfile with the backup.

Figure 11–9 Choosing Tablespaces in Tablespace Page

During the process, a few wizard pages will appear in which you will have the option to perform the following tasks:

- Choose to whether to include archive logs in the backup and if they should be deleted after each backup.
- Select a full or an incremental backup.
- Choose an online or offline mode of backup.
- Select the default configuration or any of the user created configurations for the backup. A default configuration is created by Enterprise Manager for each target database.
- Choose to override the backup and retention policy.
- Schedule the execution of a backup.
- Select when to submit the job and whether to add it to the job library.
- Specify multiple targets to submit the job. Note: To submit the same backup to multiple targets, the databases must have the same structure and the same disk/tape configuration.

Backing Up and Deleting Archived Logs

An archived redo log is an online redo log that Oracle has filled with redo entries, rendered inactive, and copied to one or more log archive targets. You should maintain multiple copies if possible.

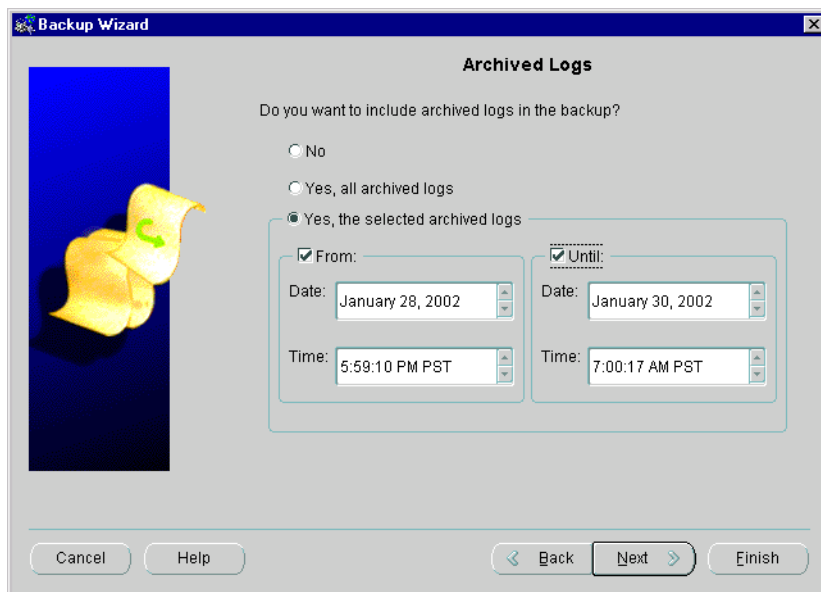
Archived redo logs are the key to successful media recovery. Back them up regularly. You can back up logs by issuing selecting Archive Logs from a customized backup strategy or by backing up datafiles and control files and specifying to include archive logs in the backup.

Typically, database administrators back up archived logs on disk to a third-party storage medium such as tape. You can also back up archived logs to disk.

Backing Up Archived Logs

To select to back up archive logs and the date and time for the first and last archive logs to be backed up

1. From the Backup Management menu, choose **Backup** to access the Backup Wizard.
2. On the Strategy Choice page, select **Customize backup strategy**.
3. On the Backup Selection page, select **Archive Logs**. This option is available only in ARCHIVELOG mode. Refer to "Setting the Database in ARCHIVELOG Mode" for more information.

Figure 11–10 Archived Logs

If you select to include all or selected archived logs in this backup, the Archived Log Deletion page appears.

Deleting Archived Logs

From the Archived Log Deletion page, you can delete the input logs (from the primary archiving destination only) automatically after the backup completes.

Figure 11–11 Archived Log Deletion

Archived Log Deletion

Do you want the archived logs to be deleted after every backup? You may select one of the following options based on your available disk space and desired recovery time.

No, I will delete the archived logs later.
You must have enough disk space for your archived logs. Your database will stop running if you run out of disk space.

Yes, delete archived logs that are older than the specified days and have the specified number of backups.
Days: Number of backups:
You must have enough disk space to keep the logs until they are deleted.

Yes, delete all archived logs after every backup.
This option may result in less disk space usage than the above options. However, the recovery may be slower since the latest archived logs may need to be restored from backups.

Cancel Help < Back Next > Finish

Select one of the following options based on your available disk space and desired recovery time.

- No, I will delete the archived logs later. This choice is available for databases of all supported versions. You must have enough disk space for your archived logs.
- Yes, delete archived logs that are older than the specified days and have the specified number of backups. This option can override backup optimization. With backup optimization on, each archived log will be backed up only once. If you select this option, you can back archive logs exactly the number of times you specify. The archive logs will be deleted after they have enough backups and are older than the specified days. The default number of days is 7. This choice is available for 9.2 target databases.

The default value for number of backups is 2. You may modify the value. With this selection, only archived logs with enough number of backups will be deleted after this backup.

- Yes, delete all archived logs after every backup. This choice is available for databases of all supported versions. This is the default.

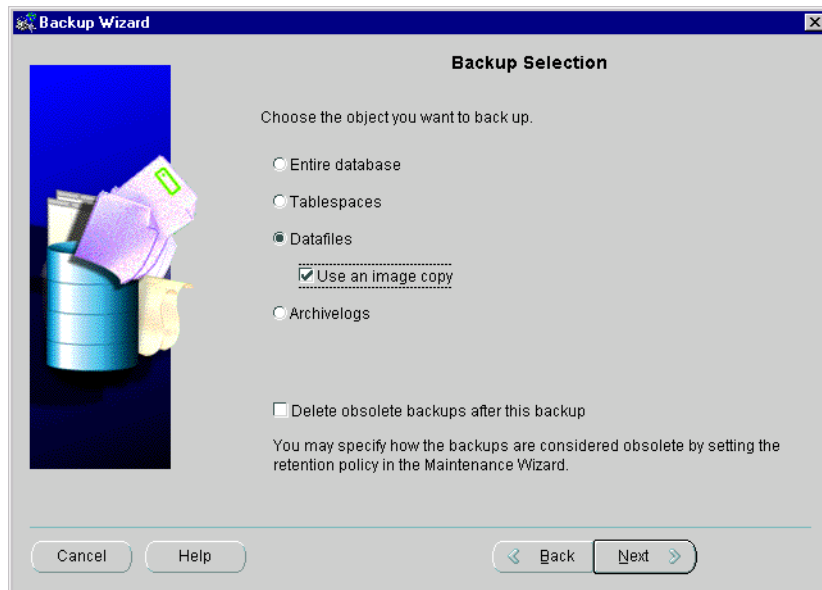
Deleting archive logs saves space. Each archived log will be backed up once before it is deleted.

Copying a Datafile

An image copy contains a single datafile, archived redo log file, or control file that you can use as-is to perform recovery. RMAN only writes image copies to disk.

1. From the Backup Management menu, choose **Backup** to access the Backup Wizard.
2. On the Strategy Choice page, select **Customize backup strategy**.
3. On the Backup Selection page, select **Datafile** and check the **Use an image copy** box.

Figure 11–12 Choosing Datafile and Use Image Copy

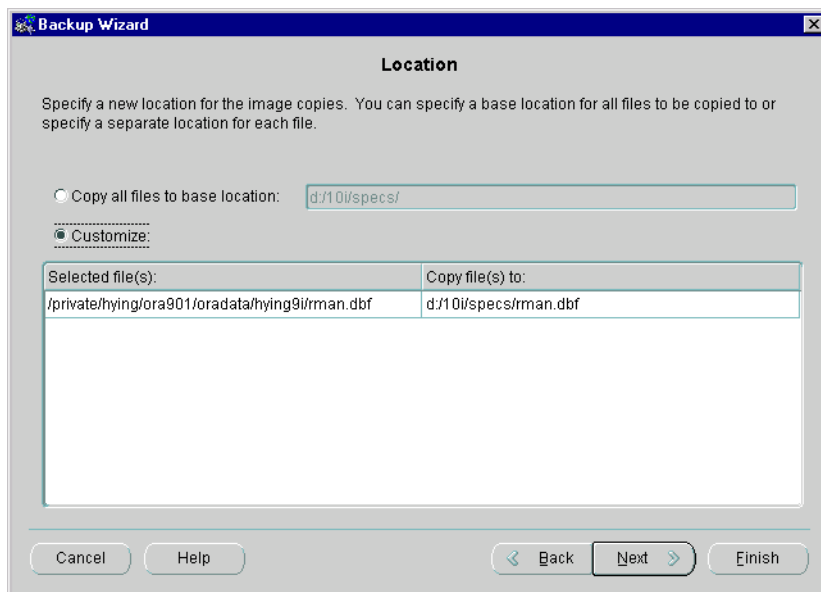


Check the **Use an image copy** box if you want to back up a datafile using an image copy. Oracle supports performing a backup using an image copy of datafiles, or controlfiles. You can only perform an image copy of a controlfile with an image copy of a datafile. Using an image copy of only the controlfile is

directly supported from the Backup Wizard. You may submit a "Run Rman Script" job separately from the Console to perform the controlfile image copy. For information on the Rman script, see "RMAN Job Script" on page 11-75.

4. On the Configuration page, select the default configuration or any of the user created configurations for the backup. A default configuration is created by Enterprise Manager for each target database.
5. On the Datafiles page, select the datafiles of which you want to make an image copy.
6. On the Locations page, choose the file locations from
 - **Copy all files to the base location**, which is the default selection. You can select the base location to place the files. The locations are computed by appending the file names from the Selected file(s) column to the base location given in the **Copy all files to base location** field. The values of the locations change as you edit the base location.
 - **Customize**, which allows you to modify the location of each file.

Figure 11–13 Location page



During the process, a few wizard pages will appear in which you will have the option to perform the following tasks:

- Choose to override the backup and retention policy.
- Schedule the execution of a backup.
- Select when to submit the job and whether to add it to the job library.
- Specify multiple targets to submit the job if you have selected to submit the job to multiple targets. Note: To submit the same backup to multiple targets, the databases must have the same structure and the same disk/tape configuration.

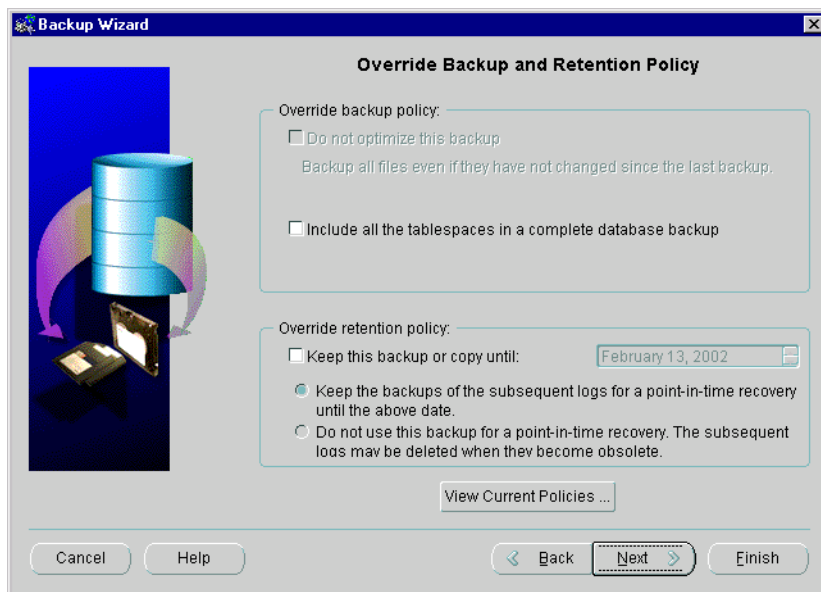
Overriding a Retention Policy in Backup for Special Cases

Backup and retention policies are a set of parameters set in the database which define how to perform a backup and how backups are retained. Once configured, these parameters apply to all the subsequent backups, but you can choose to override these backup and retention policies.

The Override Backup and Retention Policy page enables you to make a special backup occasionally that is different from what you have defined for your policy. For example, you may want to override the retention policy and keep one backup for a relatively longer period or you may want to perform a complete whole database backup by choosing **Do not optimize this backup** and **Include all the tablespaces in a complete database backup**.

The Override Backup and Retention Policy page is available if the target database is a 9i release or above, and the backup and retention policies have been set in the target database.

Figure 11–14 *Override Backup and Retention Policy*



Overriding the Backup Policy

The **Do not optimize this backup** option is enabled when the database has been configured to use backup optimization (from the Maintenance Wizard) and you have chosen to perform a database or archivelog backup. If checked, the Backup Wizard generates the "force" option in the RMAN command. The files are backed up even if they have not changed since the last backup. Choose this option if you want to override the backup optimization configuration.

The **Include all the tablespaces in the complete database backup** option is available if the target database has been configured to exclude some tablespaces from a database backup and you are planning to make a database backup. This allows you to override the "exclude for tablespace" configuration and include all tablespaces in the database backup. This corresponds to the "noexclude" option in RMAN.

Overriding the Retention Policy

The **Keep this backup or copy until <specified time>** option allows you to override the retention policy and keep the backup or copy until a specified time is available if a retention policy has been set in the target database. If the database is in ARCHIVELOG mode and you are planning to perform an online backup, the necessary archived logs will be kept to allow recovery of the database.

If the database is in ARCHIVELOG mode and you are planning to perform an offline backup, you have the option of whether to keep the subsequent archived logs or their backups.

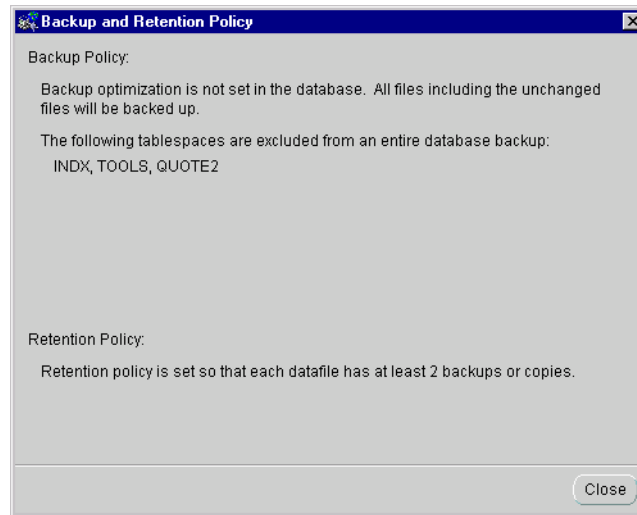
Refer to the choices below:

- Keep the backups of the subsequent logs for a point-in-time recovery
If the subsequent logs are kept, a point-in-time recovery to any time between the backup time and today.
- Do not use the backup for a point-in-time recovery. The subsequent logs may be deleted when they become obsolete.
If the logs are not kept, the backup can only be used for vaulting purpose and can be used only for recovery to the point of time when backup was taken.

Viewing Current Policies

Click the **View Current Policies** button to view the current backup and retention policy setting in the target database. The dialog is read-only.

Figure 11–15 Backup and Retention Policy Read-Only Dialog



Restore and Recover

Typically, you restore and recover a database or subset of a database in the following cases:

- A media failure has damaged some or all control files or datafiles.
- You want to recover the database to a point before a user error, such as a dropped table, occurred.

This section contains the following topics:

- Recovering the Entire Database
- Restoring the Entire Database
- Recovering and/or Restoring Tablespaces or Datafiles
- Restoring the Control File
- Restoring Archive Log Files
- Recovering Datablocks

The basic procedure for performing restore and recovery with RMAN is as follows:

1. Determine which database files require recovery.
2. From the Backup Management menu, choose **Recover** to access the Recovery Wizard to recover the restored files.
3. On the Operation Selection page, select your choice. The choices which appear on this page depends on your database's state.

The default operation is to perform a restore and recover.

Restore only, recover only, and recover data blocks only are advanced options.

You may choose to perform a restore only in one of the following situations:

- if DBVERIFY will be run against the restored files before a recovery.
- if only archived logs need to be restored.
- if datafiles or tablespaces need to be restored from an older backup set.

You may choose to perform a recovery only if you have restored the files before or you know there is no need to restore files.

You may choose to recover data blocks only in a 9i database if you know the corruption is limited to a few data blocks.

4. If the Operation Selection page does not display the choice you want, place the database in the state appropriate for the type of recovery that you want to perform.
5. In the Object Selection page, select a restore and/or recovery operation for the **Entire Database, Tablespaces, Datafiles, or Archivelogs**. Depending on the status of the target database, some options will be disabled or do not appear.

Recovering the Entire Database

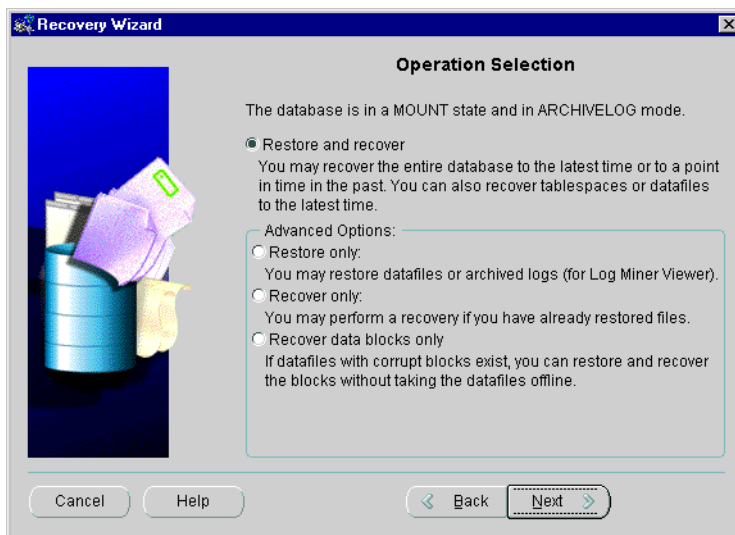
A recovery of an entire database is a recovery of all database files that belong to a database. RMAN uses the backups and copies that you made earlier and restores the files to their correct locations. Then, it uses archived redo logs (if needed) to recover the database.

You can recover the entire database to the latest time or to a point in time if the database is in ARCHIVELOG mode and MOUNT state.

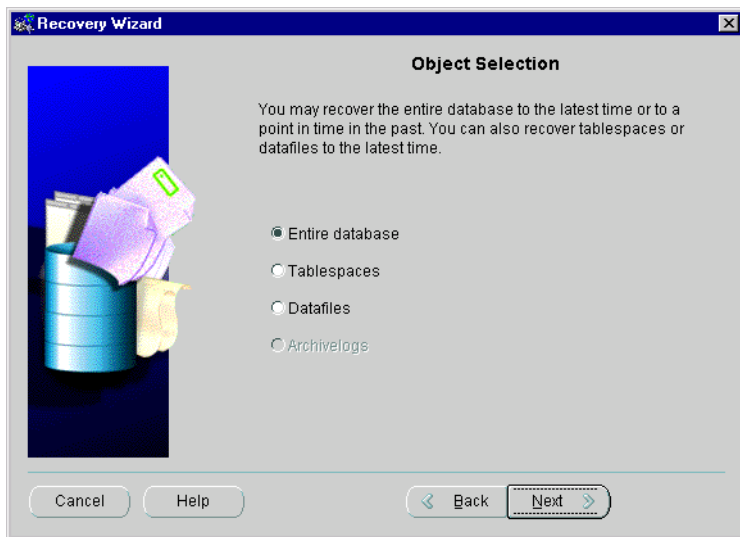
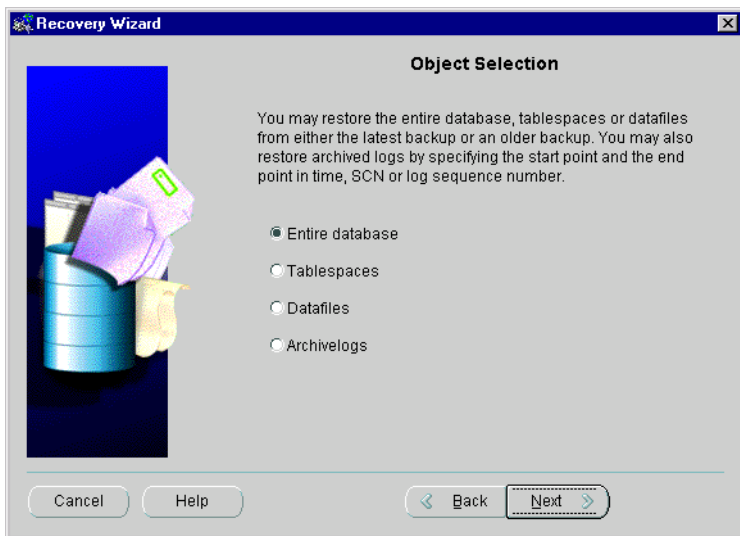
To restore and recover the database using the default disk channel, perform the following steps:

1. From the Backup Management menu, choose **Recover** to access the Recovery Wizard.
2. On the Operation Selection page, select either **Restore and recover**, **Restore only**, or **Recover only**.

Figure 11–16 Database in ARCHIVELOG mode and Mount state



3. On the Object Selection page, select **Entire Database**.

Figure 11–17 Recovery Selection Mount**Figure 11–18 Restore Selection Mount**

- Specify a range (recovery until time) on the Range Selection page of the Recovery Wizard.

You can select to recover the database to a point-in-time in the past if you have selected **Restore and recover** or **Recover only** on the Operation Selection page, and the database is in the MOUNT state and ARCHIVELOG mode.

Figure 11–19 Range Selection for Recover Database

Range Selection

You may recover the entire database to the latest time. If you want to recovery the entire database to a point-in-time in the past, then enter a date and time, or an SCN, or a log sequence number below.

Recover to the latest time

Recover to a point-in-time in the past

Date: December 10, 2001 5:52:08 PM PST

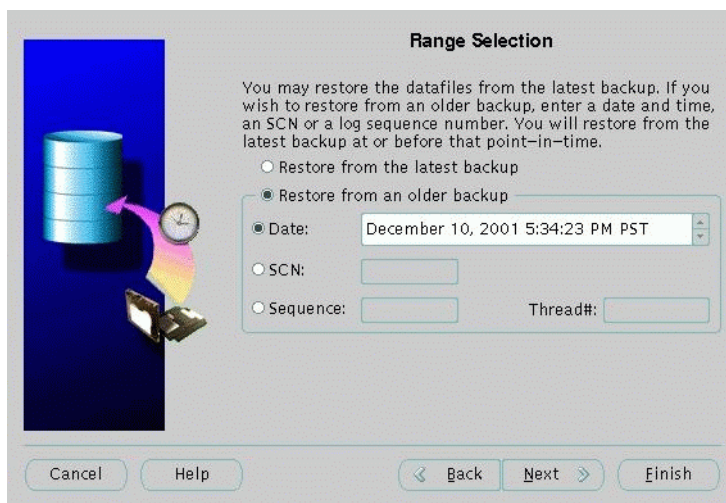
SCN:

Sequence: Thread#:

Cancel Help Back Next

You can specify a point-in-time in the past by giving a time, an SCN or a log sequence number if you had selected **Restore only** in the Operation Selection page.

Figure 11–20 Range Selection in Restore



Later, in the process, wizard pages will appear in which you will have the option to perform the following tasks:

- Specify to restore the files to a different location and rename them; thereby making datafile copies.
- Select the default configuration or any of the user created configurations for the backup. A default configuration is created by Enterprise Manager for each target database.

Restoring the Entire Database

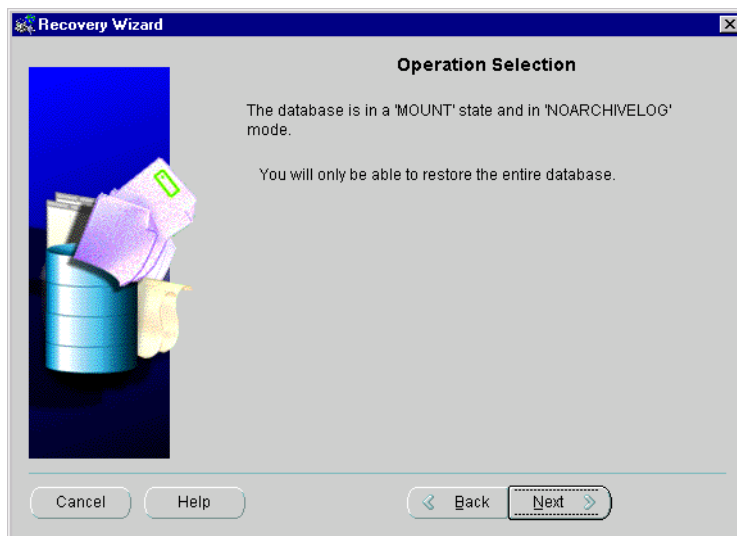
If the database is in NOARCHIVELOG mode and MOUNT state, you can restore only the entire database.

When you run your database in NOARCHIVELOG mode, the filled online redo log files are not archived. If the database's redo log operates in NOARCHIVELOG mode, the database can be completely recovered from instance failure but not from disk failure. Also, the database can be backed up only while it is completely closed. Because no archived redo log is created, no extra work is required by the database administrator.

1. From the Backup Management menu, choose **Recover** to access the Recovery Wizard.

2. On the Operation Selection page, press the **Next** button since you will only have the choice of restoring the entire database.

Figure 11–21 Database in NOARCHIVELOG mode and MOUNT state



3. Select the default configuration or any of the user created configurations for the backup on the Configuration page. A default configuration is created by Enterprise Manager for each target database.
4. Click the **Finish** button.

Recovering and/or Restoring Tablespaces or Datafiles

It is not uncommon for a media failure to affect some but not all files in a database. You can recover tablespaces or datafiles to the latest time if

- the database is in ARCHIVELOG mode and MOUNT state
 - the database is in ARCHIVELOG mode and OPEN state
1. From the Backup Management menu, choose **Recover** to access the Recovery Wizard.
 2. On the Operation Selection page, select **Restore and recover**, **Restore Only**, or **Recover Only**.

Figure 11–22 Database in ARCHIVELOG mode and OPEN state

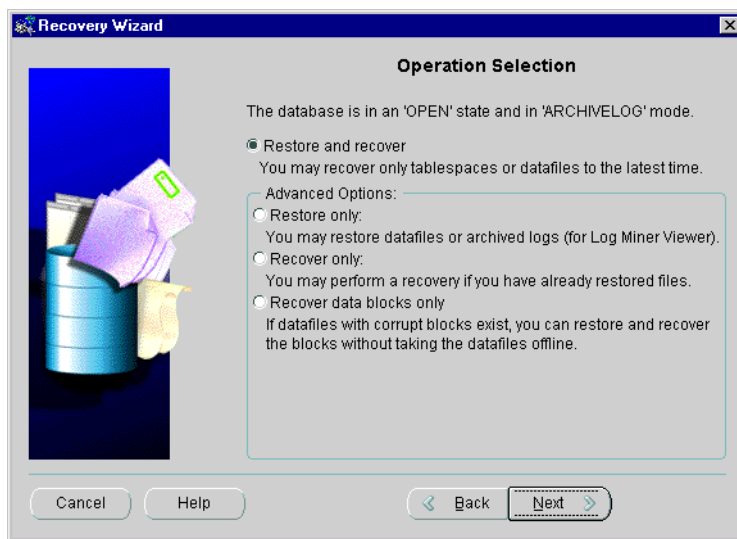
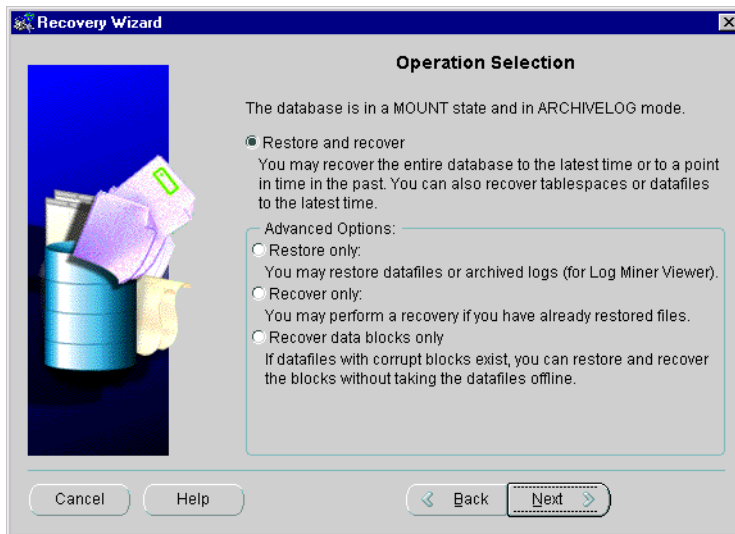


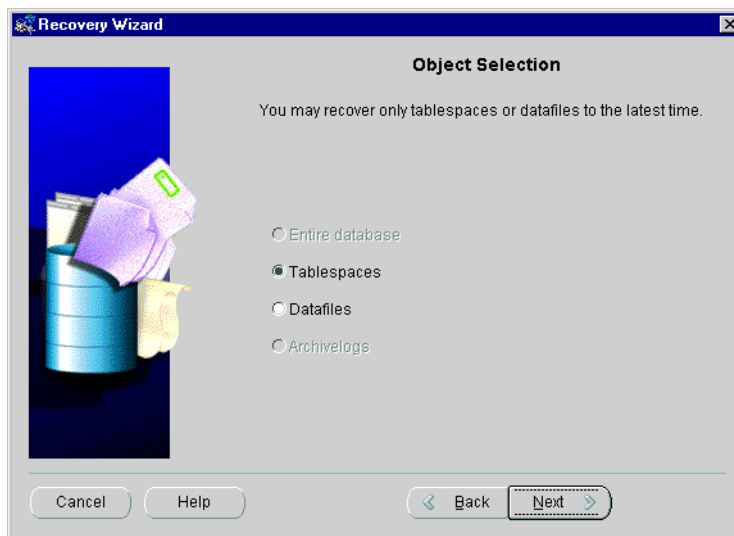
Figure 11–23 Database in ARCHIVELOG mode and Mount state



3. On the Object Selection page, select Tablespaces or Datafiles.

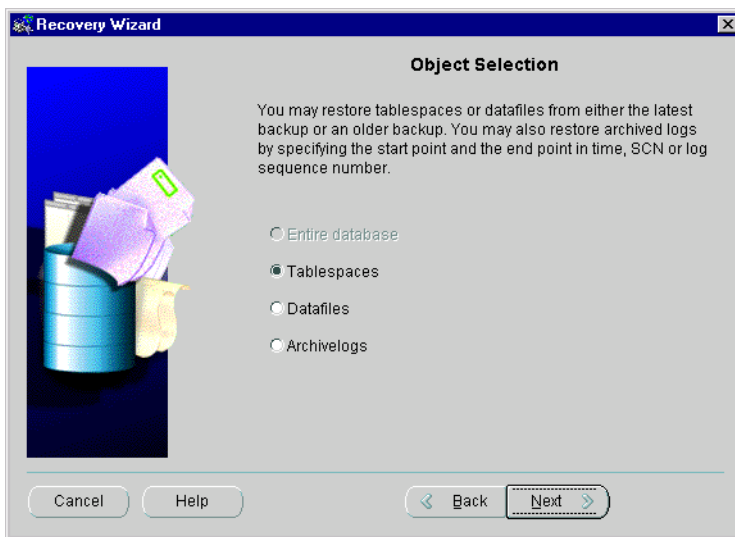
If you have chosen **Restore and recover** or **Recovery only** on the Operation Selection page and the database is in the Open state and ARCHIVELOG mode, you can recover the files to the latest time.

Figure 11–24 Recovery Selection Open

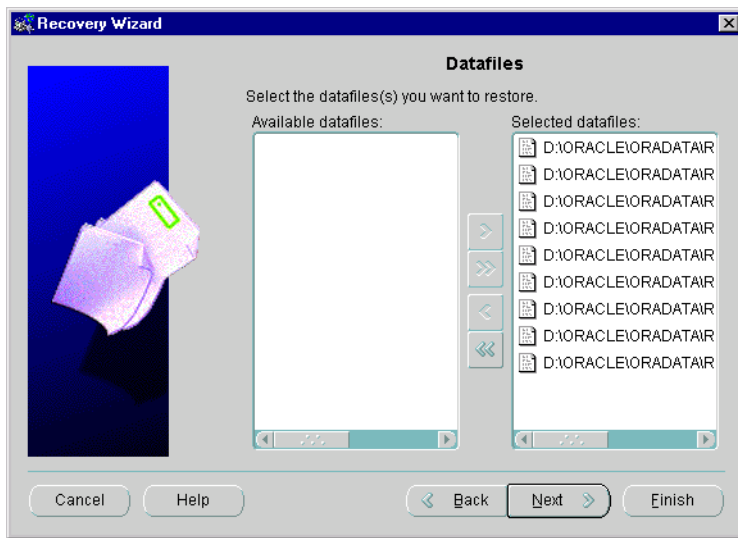


If you had selected **Restore only** on the Operation Selection page, and the database is in the OPEN state and ARCHIVELOG mode, you can restore the files from the latest backup or an older backup.

Figure 11–25 Restore Selection Open



4. Choose tablespaces from the Tablespace page or datafiles from the Datafile page.



Later, in the process, wizard pages will appear in which you will have the option to perform the following tasks:

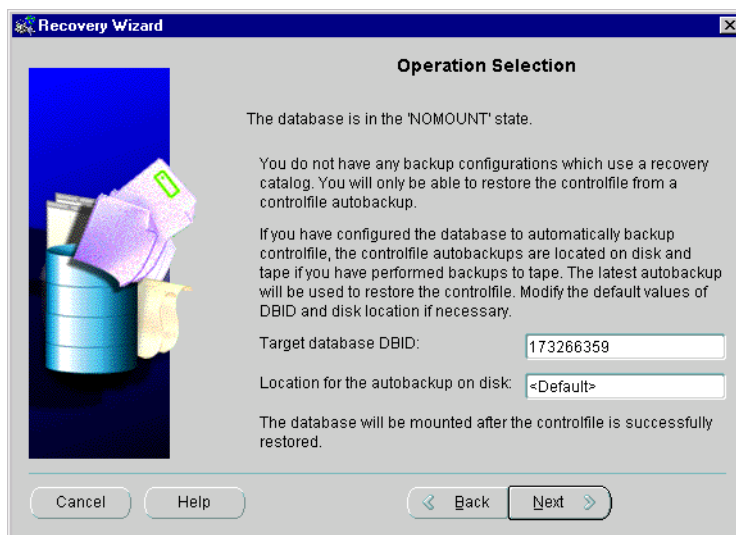
- Specify if you want to restore the files to a different location and rename them.
- Select the default configuration or any of the user created configurations for the backup. A default configuration is created by Enterprise Manager for each target database.

Restoring the Control File

When the database is in the NOMOUNT (Started) state and there are no backup configurations which use a recovery catalog, you can restore a controlfile from a controlfile autobackup for Oracle 9.0.1 target databases. For pre-9.0.1 target databases, an error dialog appears when the database is in NOMOUNT and there are no backup configurations that use a recovery catalog.

1. From the Backup Management menu, choose **Recover** to access the Recovery Wizard.
2. On the Operation Selection page, fill in the Target database DBID and the Location of the autobackup on disk.
3. Select the default configuration or any of the user created configurations for the backup on the Configuration page.
4. Press Finish on the Configuration page.

Figure 11–26 Database in NOMOUNT (Started) state and there is not any backup configuration which uses a recovery catalog



If you have performed backups to tape, your controlfile autobackups are located on both disk and tape. To restore the controlfile from autobackup, select the backup configuration you used when backing up to tape. Both tape and disk will be searched and the latest autobackup will be used to restore the controlfile.

If you have never performed a backup to tape, all your controlfile autobackups are located on disk. To restore the controlfile from autobackup, select the backup configuration you used to backup to disk. The latest autobackup will be used to restore the controlfile.

RMAN requires the DBID to be specified when restoring controlfile from autobackup. If you have performed at least one backup from the Enterprise Manager Backup Wizard, the DBID is stored in the Enterprise Manager repository and the value is filled in the DBID field. Otherwise you will have to enter the DBID manually.

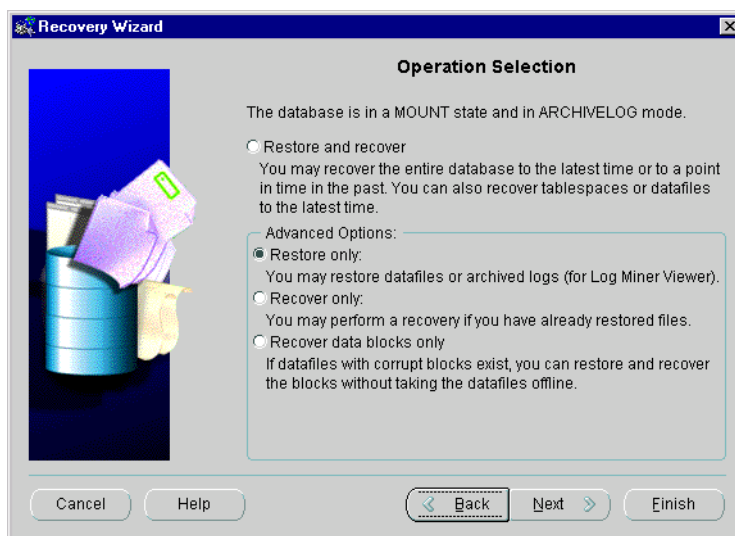
If you have specified a location for the autobackup on disk using the Maintenance Wizard, the location is stored in the Enterprise Manager repository and the value is filled in the location field. Otherwise you have to fill in the value manually. If you do not specify a value, the default location will be searched for autobackups on disk.

Restoring Archive Log Files

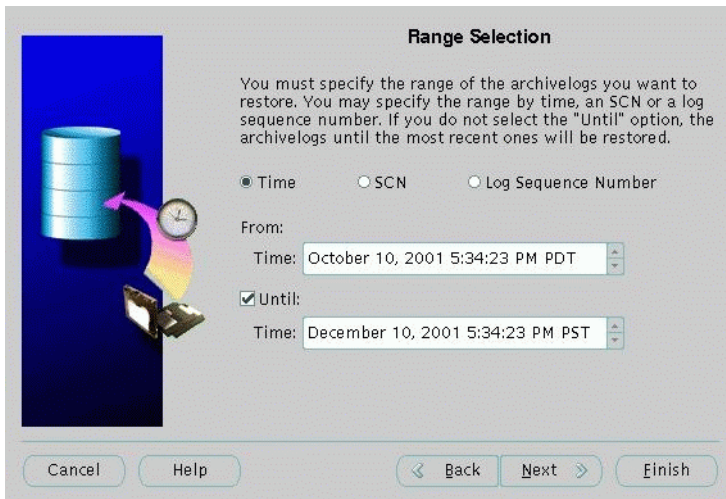
You can restore archive log files for use with LogMiner Viewer if the archive logs need to be restored from backups if they are no longer on disk.

1. From the Backup Management menu, choose **Recover** to access the Recovery Wizard.
2. On the Operation Selection page, select **Restore only**.

Figure 11–27 Restoring Only



3. On the Object Selection page, select **Archivelogs**.
4. To restore archived logs, a range has to be specified using one of the three methods:
 - Specify archived log range by time.
 - Specify archived log range by SCN.
 - Specify archived log range by log sequence.

Figure 11–28 Restore Archivelog by Time

Range Selection

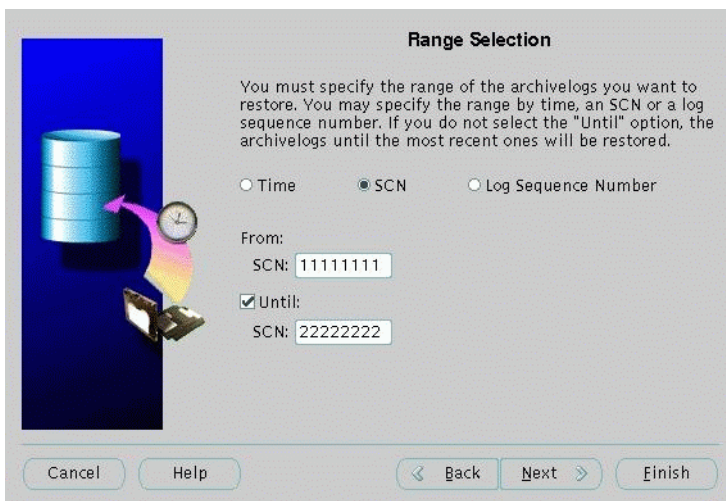
You must specify the range of the archivelogs you want to restore. You may specify the range by time, an SCN or a log sequence number. If you do not select the "Until" option, the archivelogs until the most recent ones will be restored.

Time SCN Log Sequence Number

From:
Time:

Until:
Time:

Cancel Help < Back Next > Finish

Figure 11–29 Restore Archivelog by SCN

Range Selection

You must specify the range of the archivelogs you want to restore. You may specify the range by time, an SCN or a log sequence number. If you do not select the "Until" option, the archivelogs until the most recent ones will be restored.

Time SCN Log Sequence Number

From:
SCN:

Until:
SCN:

Cancel Help < Back Next > Finish

Figure 11–30 Restore Archivelog by Log Sequence

Range Selection

You must specify the range of the archivelogs you want to restore. You may specify the range by time, an SCN or a log sequence number. If you do not select the "Until" option, the archivelogs until the most recent ones will be restored.

Time
 SCN
 Log Sequence Number

From:

Sequence: Thread#:

Until:

Sequence:

5. Select the default configuration or any of the user created configurations for the backup on the Configuration page.
6. Press **Finish** on the Configuration page.

Recovering Datablocks

You can recover datablocks if

- the database is in ARCHIVELOG mode and MOUNT state, or
 - the database is in ARCHIVELOG mode and OPEN state
1. From the Backup Management menu, choose **Recover** to access the Recovery Wizard.
 2. On the Operation Selection page, select **Recover datablocks only**.

Figure 11–31 Recover Datablock Open

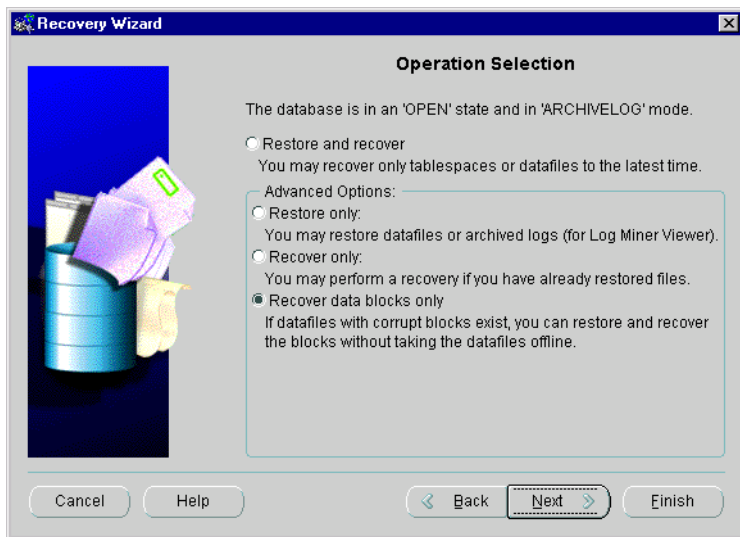
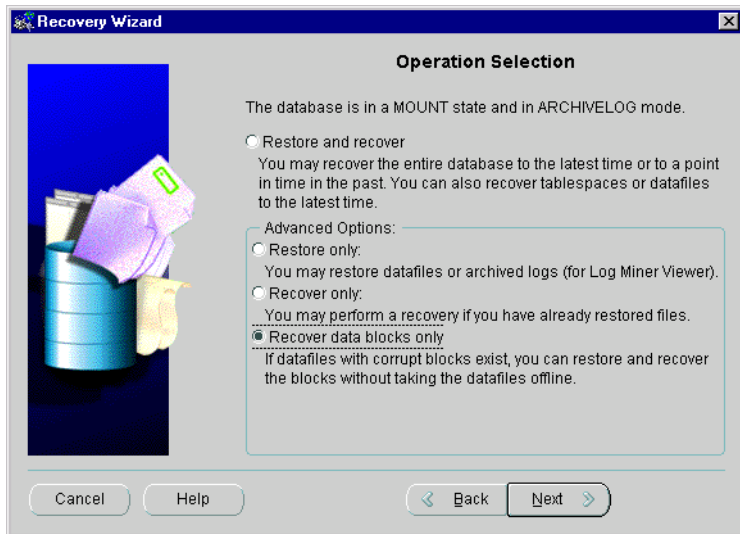


Figure 11–32 Recover Datablock Mount



3. On the Block Media Recovery Method page of the Recovery Wizard, specify the data blocks to recover by selecting one of the options.

Using a Corruption List for Data Block Recovery

The **Corruption List** option is the default and preferred selection for most users. All the data blocks that have been identified corrupted by RMAN during the backup and copy operations can be recovered.

Data block corruption in non-RMAN operations, such as dbverify, is not included in the corruption list. To recover those database blocks, you will have to find the block numbers or data block addresses of the corrupted blocks from Oracle standard output, an alert.log, user trace files, results of the ANALYZE TABLE and ANALYZE INDEX SQL commands, result of the DBVERIFY utility, or third-party media management output. However, the corrupted blocks are always recorded in the alter.log file.

The **View Corruption List** button is visible for 9.2 databases only. By clicking the button, a Corruption List dialog will appear, showing the current data blocks that are labeled corrupted by RMAN.

Using Datafiles for Data Block Recovery

Select the **Datafiles** option on the Block Media Recovery Method page if the block numbers of the corrupted data blocks can be found in one of the places listed above.

When the Data Block Selection by Datafile page appears, enter the block numbers of the corrupted data blocks for each datafile. If you are entering more than one block number for a datafile, separate the entries by a space or comma.

Figure 11–33 Data Block Selection by Datafile

Data Block Selection by Datafile

Specify the data blocks that need media recovery by selecting datafiles and entering the block numbers for each selected datafile. More than one data block can be entered for one datafile. Separate block numbers by space or comma.

Datafile Name	File#	Block Numbers
/private/hying/ora901/oradata/hying9i/system01.d	1	1,2,3
/private/hying/ora901/oradata/hying9i/undotbs01.	2	
/private/hying/ora901/oradata/hying9i/cwmlite01.	3	
/private/hying/ora901/oradata/hying9i/drsys01.dbf	4	
/private/hying/ora901/oradata/hying9i/example01.	5	
/private/hying/ora901/oradata/hying9i/indx01.dbf	6	
/private/hying/ora901/oradata/hying9i/tools01.dbf	7	
/private/hying/ora901/oradata/hying9i/users01.dbf	8	
/private/hying/ora901/oradata/hying9i/QUOTE.dbf	9	
/private/hying/ora901/oradata/hying9i/"dd.dbf"	10	

Cancel Help < Back Next > Finish

Using Tablespaces for Data Block Recovery

Select the **Tablespaces** option on the Block Media Recovery Method page if the data block addresses of the corrupted data blocks can be found in one of the places listed above.

When the Data Block Selection by Tablespace page appears, enter the data block addresses for each tablespace. If you are entering more than one block address for a tablespace, separate the entries by a space or comma.

Figure 11–34 Datablock Selection by Tablespace

Data Block Selection by Tablespace

Specify the data blocks that need media recovery by selecting tablespaces and entering the data block addresses for the each selected tablespace. More than one address can be entered for one tablespace. Separate addresses by space or comma.

Tablespace Name	Data Block Addresses
CWMLITE	12
DRSYS	
EXAMPLE	
INDX	
QUOTE	
QUOTE2	
SYSTEM	
TEMP	
TOOLS	
UNDOTBS	

Maintenance Operations

You will use the Maintenance wizard to help you perform maintenance operations on the target databases and on the recovery catalog.

This section contains the following topics:

- Setting up Backup and Retention Policies in a Target Database
 - Configuring a Backup Policy
 - Configuring a Retention Policy
- Performing Recovery Catalog Maintenance
 - Registering a Database
 - Resynchronizing the Catalog
 - Resetting the Database

Setting up Backup and Retention Policies in a Target Database

To set up backup and retention policies in a target database

1. From the Backup Management menu, choose **Maintenance** to access the Maintenance Wizard.
2. On the Operation Choice page, select **Modify backup and retention policies in the target database**.

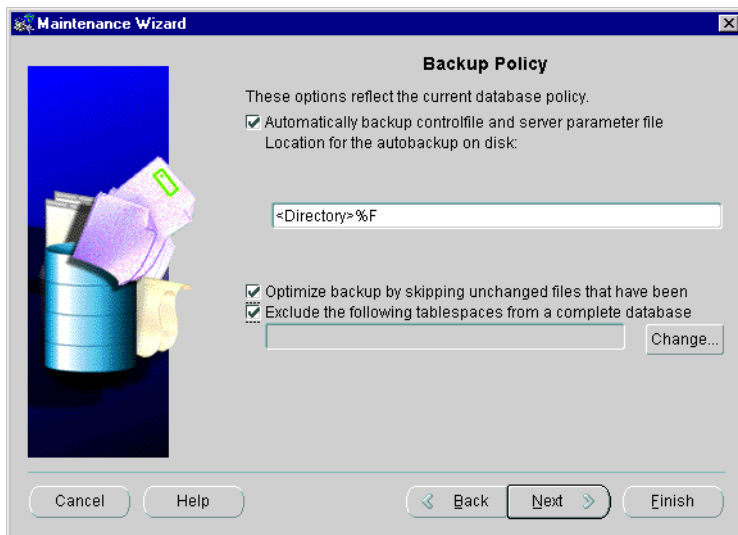
You may modify persistent RMAN configurations in the target database using this option.

The wizard proceeds to the Backup Policy and Retention Policy pages where you can modify backup related RMAN configuration parameters. This option is enabled for 9i and above databases only. A one time job will be submitted to execute the changes.

Configuring a Backup Policy

Choose a backup policy for your database.

Figure 11–35 Backup Policy



Choosing to automatically back up controlfile and server parameter file. By default this option is not selected, but it is highly recommended that you choose the option for your database and set the location for the autobackup on disk.

The controlfile and server parameter file (SPFILE) will be backed up automatically after each backup. Additionally these files are backed up automatically after each structural change in the database. The backups of these files are called controlfile autobackups.

If you are making a backup to disk, the controlfile and SPFILE will be automatically backed up to disk. The database backup location is determined by the format of the allocated disk channel. The controlfile autobackup is located in what is specified in the **Location for the autobackup on disk** field.

If you are making a backup to tape, the controlfile and SPFILE will be automatically backed to tape. The database backup format is determined by the format of the allocated tape channel. The controlfile autobackup format is always %F.

After each structural change (such as adding a tablespace or a datafile) in the database, the controlfile and SPFILE will always be backed up to disk.

The autobackup is located in what is specified in the **Location for the autobackup on disk** field below.

Note: %F is required as part of the format string. %F contains DBID information which is used to uniquely identify a database.

Specify the **Location for the autobackup on disk**. It is the location of the autobackups if you are making a backup to disk and after a database structural change. If you do not specify a location on disk, the default location will be used. The default location is on the same disk as the database. The default location is unlikely to be available when you need to restore the controlfile from autobackup. Therefore, it is highly recommended that you specify a disk location which is different from the database location.

If the autobackup format for disk has not been configured in the database, the location appears as <Directory>%F. You are recommended to change the Directory path.

If the autobackup format for disk has been configured, the configured format will appear in the location field.

Note: %F is required as part of the format string. %F contains DBID information which is used to uniquely identify a database.

Choosing to optimize the backup. Database backup and archivelog backup may be optimized by skipping files that have not changed since the last backup. Typical examples of unchanged files include offline or read-only datafiles as well as archivelogs. Once backup optimization is set, archivelogs can only be backed up once unless you specify **Yes, delete archived logs only after a specified number of**

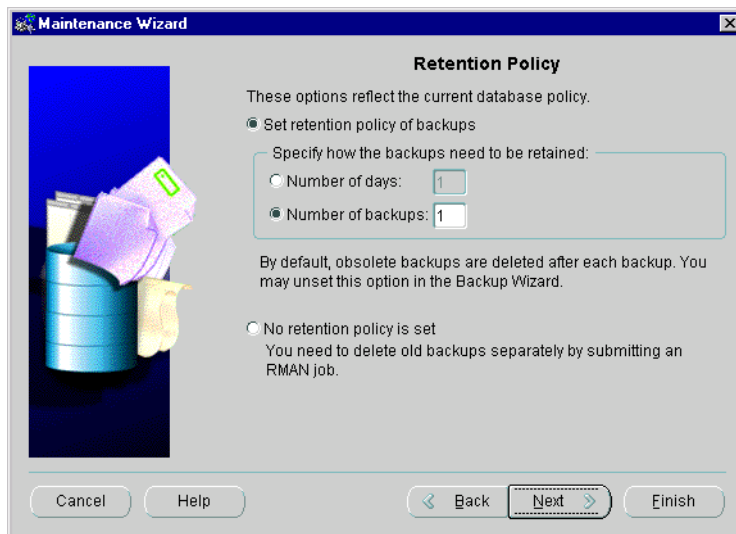
backups option in the Backup Wizard's Archived Log Deletion page (available for 9.2 target databases).

Excluding tablespaces in a backup You can choose to exclude some tablespaces from a complete database backup. These tablespaces can still be backed up separately using a tablespace backup. This option allows you to exclude some extremely large but rarely modified tablespaces from a database backup, and back them up on a different schedule. This may reduce the amount of time required for a regularly scheduled database backup.

Configuring a Retention Policy

A retention policy allows you to define how long to retain your backups before they are marked obsolete. Based on this policy, RMAN will mark backups you do not need as obsolete. You can delete obsolete backups by regularly running delete obsolete operations.

Figure 11–36 Retention Policy



Select the **Set the retention policy of backups** option to use a retention policy. You can set the policy to retain backups so that you are able to recover database until the past "number of days" or to retain a "number of backup copies."

The **Number of days** option corresponds to "configure retention policy to recovery window of x days", where x is the value from the **Number of days** field.

The **Number of backups** option corresponds to "configure retention policy to redundancy y", where y is the value of the **Number of backups** field.

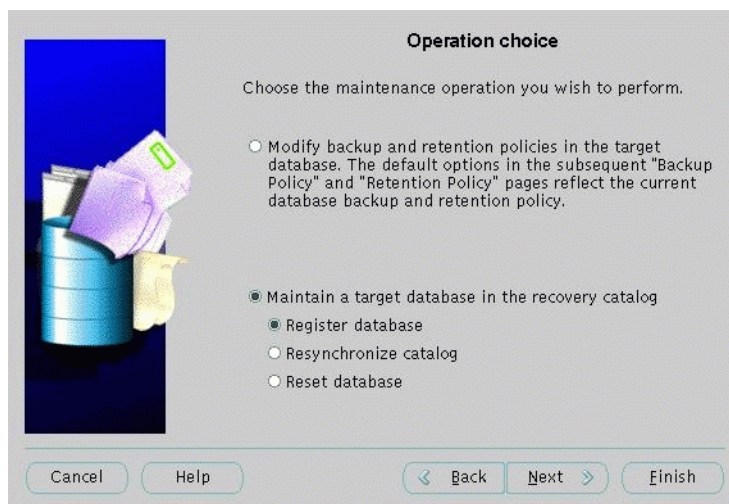
Select the **No retention policy is set** option if you want to manually delete backups. This option corresponds to RMAN's "configure retention policy to none" option.

Performing Recovery Catalog Maintenance

To register, reset, or resynchronize the target database with the recovery catalog

1. From the Backup Management menu, choose **Maintenance** to access the Maintenance Wizard.
2. On the Operation Choice page, select **Maintain a target database in the recovery catalog**.

Figure 11–37 *Maintain target database in the recovery catalog*



Registering a Database

The target database must be registered with the Recovery Catalog before the Backup wizard can use it. You only need to register the database once.

Choose the Register database option in Operation choice page.

Resynchronizing the Catalog

The recovery catalog obtains crucial RMAN metadata from the target database control file. Resynchronization of the recovery catalog ensures that the metadata that RMAN obtains from the control file stays current. Resynchronizations can be full or partial. In a partial resynchronization, RMAN reads the current control file to update changed data, but does not resynchronize metadata about the database physical schema: datafiles, tablespaces, redo threads, rollback segments, and online redo logs. In a full resynchronization, RMAN updates all changed records, including schema records.

RMAN automatically detects when it needs to perform a full or partial resynchronization and executes the operation as needed. You can also force a full resynchronization by issuing a `RESYNC CATALOG` command.

To ensure that the catalog stays current, run the `RESYNC CATALOG` command periodically if you run backups periodically.

You will not need to run the `RESYNC CATALOG` command if you perform at least one backup in *n* days, where *n* is the setting for the initialization parameter `CONTROL_FILE_RECORD_KEEP_TIME`. When you start the RMAN backup (or any other operation) it will automatically perform a "RESYNC CATALOG". In other words, if you perform a backup fairly often (for example, every 2-5 days), then there is no need to do "RESYNC CATALOG" manually.

Because the control file employs a circular reuse system, backup and copy records eventually get overwritten. Resynchronizing the catalog ensures that these records are stored in the catalog and so are not lost.

Resetting the Database

Resetting the database is rarely performed and should only be done if all the information has been lost. You must reset the recovery catalog if the target database had been previously opened with the `RESETLOGS` option. Refer to the *Oracle9i Recovery Manager User's Guide* for information on the `RESETLOGS` option.

Configuring the RMAN Environment

RMAN contains some default configuration settings. These settings apply to all RMAN sessions until you explicitly change them.

An RMAN channel represents one stream of data to a device type and corresponds to one server session. Allocation of one or more RMAN channels is necessary to execute most backup and recovery commands. Each channel establishes a connection from the RMAN executable to a target or auxiliary database instance by starting a server session on the instance. The server session performs the backup, restore, and recovery operations. Only one RMAN session communicates with the allocated server sessions.

RMAN comes preconfigured with a DISK channel that you can use for backups and copies to disk.

This section contains the following topics:

- Creating a Backup Configuration
- Specifying Disk Channel Device for Backup Set
- Specifying Tape Channel Device for Backup Set
- Setting Channel Limits
- Specifying Channel Device for an Image Copy
- Setting Up a Proxy Copy for a Tape Backup Set
- Setting the Storage Parameters for the Current Backup Set
- Registering the Recovery Catalog
- Setting Preferred Credentials for Running Backup Jobs
- Registering Later Databases with the Recovery Catalog
- Resynchronizing the Recovery Catalog with the Target Database
- Setting Up the Recovery Catalog

Creating a Backup Configuration

A configuration is a set of defaults you set up for backup and recovery. Enterprise Manager creates a default backup configuration for each target database, but you can use the Create Backup Configuration property sheets to create other backup configurations for backup and recovery. A configuration can be used for one database or many databases depending if the systems are the same.

To create a backup configuration

1. Start the Oracle Enterprise Manager Console.
2. Expand the Database folder and select the database or an example database from the Navigator.
3. Log in as a user with the DBA role.
4. Choose the **Backup Management** menu from the **Object** menu.
5. Select **Create Backup Configuration** from the **Backup Management** menu.
6. On the General Page, give a name and description to your configuration.
7. On the Channels Page, choose **Backup Set** or **Image Copy**.

The Channels page allows you to specify a channel or channels. A channel establishes a connection from the database to the storage device for backup or restore operations. A channel can be either disk or tape-based. Multiple channels can be created to allow parallel backup/recovery by a single job.

Note: If you are connected to a Certified Configurations database, refer to the *Oracle Certified Configuration Administrator's Reference* for information on features unique to Certified Configurations backups.

Note: At least one channel must exist before performing a backup, restore, or recover operation.

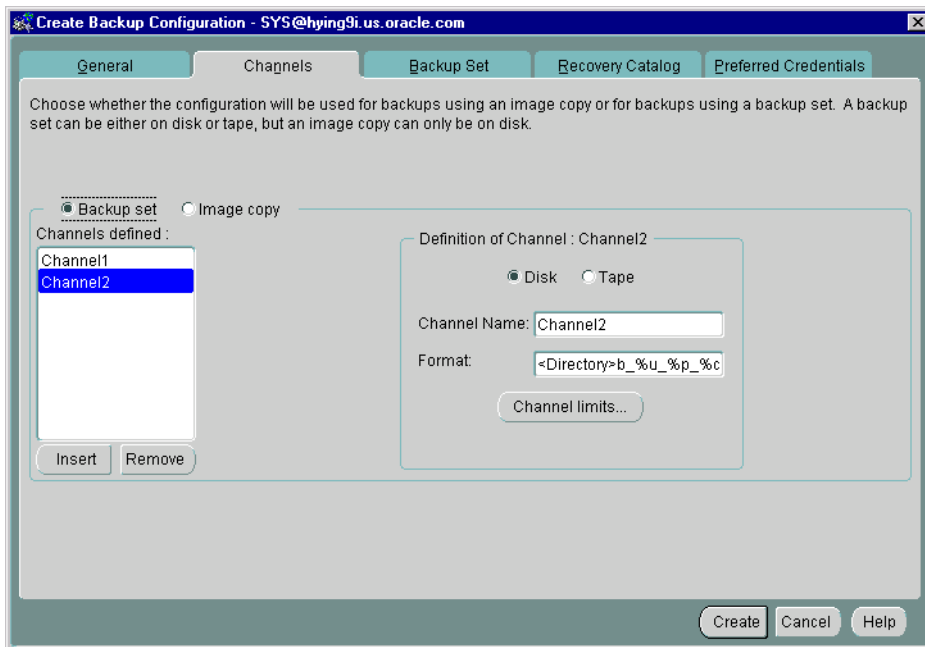
Attention: Oracle9i can only allocate one Recovery Manager channel at a time, thus limiting the parallelism to one stream. The Oracle9i Enterprise Edition allows unlimited parallelism. See *Oracle9i Database New Features* for more information about the features available with Oracle9i and Oracle9i Enterprise Edition.

Specifying Disk Channel Device for Backup Set

To create a configuration that backs up to disk using a backup set, specify the following on the Channels page of the Create Backup Configuration property sheet:

1. Specify Disk for your Backup Set. The Backup set is written to disk via Oracle Recovery Manager.
2. Specify both the directory path and the unique file name format.
<Directory>b_%u_%p_%c

Figure 11–38 Channels Disk



The Directory is the drive and path where backup sets are stored. You must specify a proper directory for the channel. The directory field must end with a proper delimiter, which is OS dependent.

The File Name is the unique backup set name. The following parameters can be used:

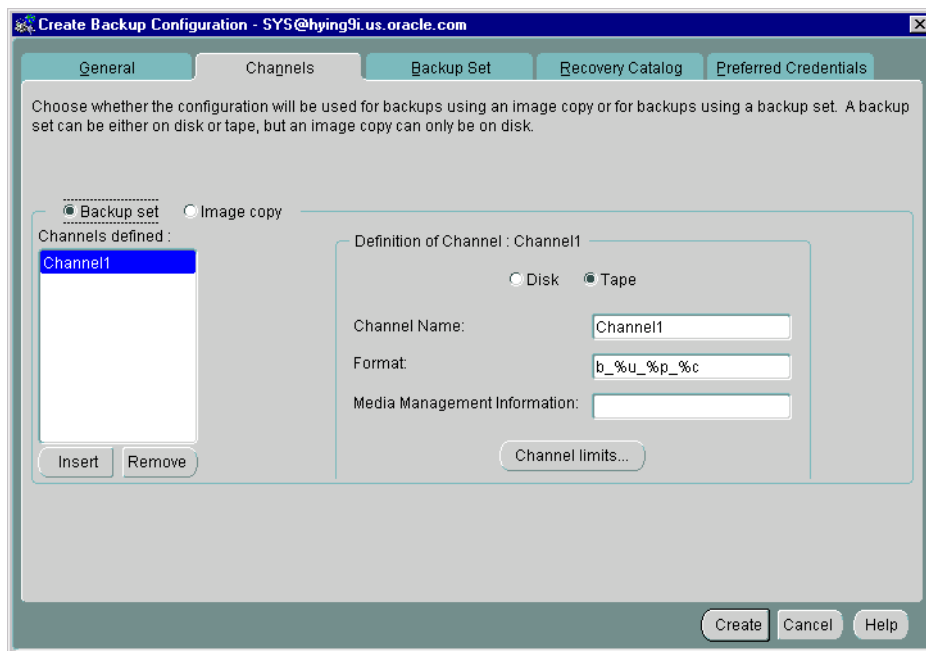
- **b** is the prefix for backup. It is not necessary to have the prefix. You can modify it to anything you want.
- **%p** is the backup piece number within the backup set. This value starts at 1 for each backup set and is incremented by 1 as each backup piece is created.
- **%U = %u_%p_%c**. **%c** has been available since 8.1.x. For information on **%c**, refer to the *Oracle9i Recovery Manager User's Guide*.
- **%u** is the unique name. If you want to backup multiple databases on the same machine using the same configuration, you must use **%u** in the format string in the configuration.
- **%s** is the backup set number. The counter value starts at 1 and is unique for the lifetime of the control file.
- **%t** is the backup set timestamp. Note: The combination of **%s** and **%t** can be used to form a unique name for the backup set.
- **%f** is the unique file format.

Specifying Tape Channel Device for Backup Set

To create a configuration that backs up to tape using a backup set, specify the following on the Channels page of the Create Backup Configuration property sheet:

1. Specify Tape for your Backup Set. Backup is via media management software.
2. Specify the unique file name format.
3. In the Media Management Information field, specify the parameters regarding the device to allocate. For information on `parms` parameter, see the *Oracle9i Recovery Manager User's Guide* and the corresponding Media Management documentation.

Figure 11–39 Channels Tape



The following parameters can be used for specifying the unique file format name:

- `b` is the prefix for backup. It is not necessary to have the prefix. You can modify it to anything you want.

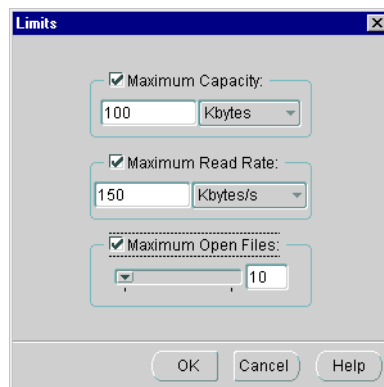
- %p is the backup piece number within the backup set. This value starts at 1 for each backup set and is incremented by 1 as each backup piece is created.
- %U = %u_%p_%c. %c has been available since 8.1.x. For information on %c, refer to the *Oracle9i Recovery Manager User's Guide*.
- %u is the unique name. If you want to backup multiple databases on the same machine using the same configuration, you must use %u in the format string in the configuration.
- %s is the backup set number. The counter value starts at 1 and is unique for the lifetime of the control file.
- %t is the backup set timestamp. Note: The combination of %s and %t can be used to form a unique name for the backup set.
- %f is the unique file format.

Setting Channel Limits

To set the limits for any backup or copy operation, press the Channel Limits button on the Channels page of the Create Backup Configuration property sheet.

For any setting, you move the slider bar to change its value or type in the value. The number in the field changes according to the position of the slider bar.

Figure 11–40 Limits page



Checking **Maximum Capacity** allows you set the maximum number of units that a backup operation can write to a single backup.

Checking **Maximum Read Rate** allows you to control the number of blocks per second read by a backup or copy operation from or to any input datafile. Controlling the read rate ensures that a backup or copy operation does not consume excessive disk bandwidth, which can degrade online performance.

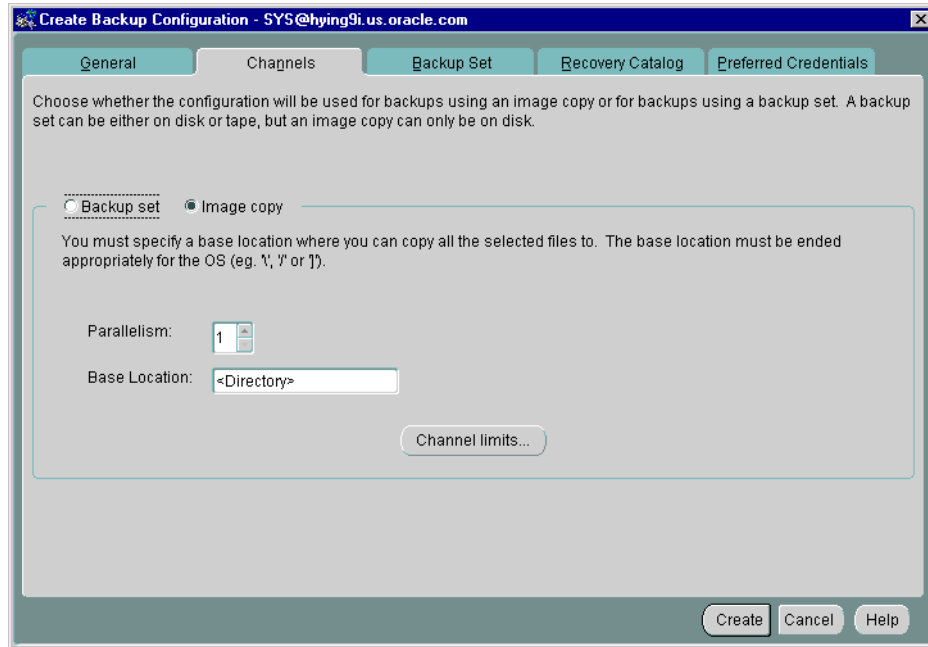
Checking **Maximum Open Files** allows you to control the maximum number of input files that a backup or copy operation can have open simultaneously. Setting maximum number of open files is particularly useful when backing up a large number of archivelogs into a single backup set.

Specifying Channel Device for an Image Copy

To create a configuration that uses an image copy, specify the following on the Channels page of the Create Backup Configuration property sheet:

1. Specify Image Copy.
2. Specify the number of channels required, the base location, and the channel limits for all the channels. You will not have to specify a name for each channel since it is generated automatically.

Image copies can only be written to disk.

Figure 11–41 Channel Device for Image Copy

In the Parallelism field, specify the number of channels for the image copy.

In the Base Location field, specify the location to which you want to copy all the selected files. You must specify a proper directory for the channel. The directory field must end with a proper delimiter, which is OS dependent.

Note: Unlike a backup set, the image copy channel does not need a format because the copy is a one-to-one just like the `cp` command at the OS level.

You can also specify one set of channel limits, which applies to all the channels in an image copy configuration. Assigning different limits for different channels in an image copy is not a necessity.

Note: At least one channel must exist before performing a backup, restore, or recover operation.

Attention: Oracle9i Standard Edition can only allocate one Recovery Manager channel at a time, thus limiting the parallelism to one stream. The Oracle9i Enterprise Edition allows unlimited parallelism. See *Oracle9i Database New Features*

for more information about the features available with Oracle9i and Oracle9i Enterprise Edition.

Setting Up a Proxy Copy for a Tape Backup Set

A proxy copy is a special type of backup in which RMAN turns over control of the data transfer to a media manager that supports this feature.

You can choose to have Media Management Software take over the backup and recovery of your files instead of using RMAN.

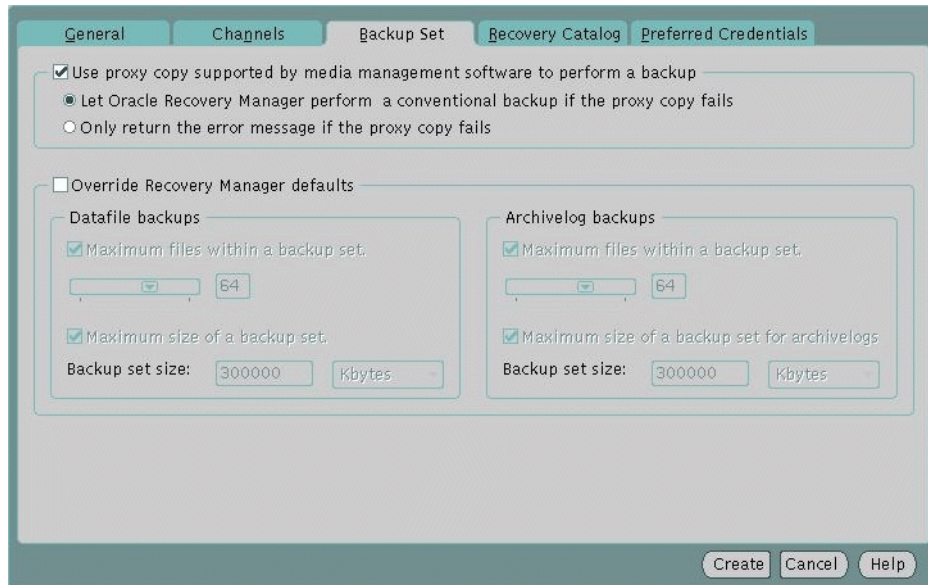
To choose to use a Proxy copy, specify **Use Proxy copy supported by media management software to perform a backup** on the Backup Set page of the Create Backup Configuration property sheet.

RMAN will provide a list of files that need backup or recovery to the media manager. Proxy copy applies only if you are integrated with media manager software that has implemented the proxy feature.

When the **Use Proxy copy supported by media management software to perform a backup** checkbox is selected, you can further specify what to do when the proxy copy fails.

- If the proxy copy fails, continue and let Oracle Recovery Manager perform a conventional backup.
- If the proxy copy fails, halt the backup and return the error message.

Figure 11–42 Proxy Option



The proxy option is enabled only if you have created a Tape channel because having a Tape channel tells the media management software to take over the backup. *Tape* means that data goes to the media management software which in most cases will be backed up to a physical tape.

Setting the Storage Parameters for the Current Backup Set

To set storage parameters for the current backup set and override the Recovery Manager default settings that the Recovery Manager has calculated, specify **Override Recovery Manager defaults** on the Backup Set page of the Create Backup Configuration property sheet.

The **Override Recovery Manager defaults** checkbox is disabled when **Image Copy** on the Channels Page is selected. The checkbox is enabled only when **Backup Set** is selected and the page is not read-only.

Setting Storage Parameters for Datafiles

Checking **Maximum files at a time in a backup set** allows you to set the maximum number of files that can be placed in a single backup set. If the number of files

selected for the current backup exceed this number, multiple backup sets are created. In addition, multiple channels, if defined and available, will also be used.

Checking **Maximum size of a backup set** allows you to set the maximum file size of a backup set. You can specify file size in megabytes or kilobytes. Specifying a set size for backup sets permits better load balancing when performing backups.

Setting Storage Parameters for Archivelogs

Checking **Maximum files at a time in a backup set** allows you to set the maximum number of files that can be placed in a single backup set for archivelogs. If the number of files selected for the current backup exceed this number, multiple backup sets are created. In addition, multiple channels, if defined and available, will also be used.

Checking **Maximum size of a backup set for archivelogs** allows you to set the maximum file size of a backup set for archivelogs. You can specify file size in megabytes or kilobytes. Specifying a set size for backup for archivelogs permits better load balancing when performing backups.

Registering the Recovery Catalog

The enrolling of a database in a recovery catalog is called registration. You can register a database only once in a given catalog schema: for example, you cannot register prod1 in the catowner catalog and then register prod1 again in the catowner catalog.

Figure 11–43 Recovery Catalog page

To register the first database with the recovery catalog, specify the following on the Recovery Catalog page of the Create Backup Configuration property sheet:

1. Select **in a recovery catalog** option as where you want your backup information stored.
2. In the **Username** field, enter the name of the user. For example, `rman`.
3. In the **Password** field, enter a password. For example, `rman`.
4. For the Service Name, choose an existing service or enter a new service.

The service name is the name of the database where the recovery catalog resides.

If you choose to use an existing service, you can choose from a list of all the service names of databases that have been discovered by the Management Server.

If you choose to enter a new service, you must enter a new recovery catalog service by specifying the host name, port and sid.

- *Host Name* is the machine name where the database is located

- *Port* is the database listener port address, usually 1521 or 1526
- *SID* is the database system identifier

A TNS descriptor constructed using the user supplied host name, port and sid will be used to connect to the recovery catalog database during the backup.

5. Press the Create button. The database you selected will automatically be registered.

You only need to register the database once. The Backup Configuration you have set become your default backup configuration if you use the Backup and Recovery wizards and submit a job.

Setting Preferred Credentials for Running Backup Jobs

If the Oracle Enterprise Manager Console's preferred credentials are not set to SYSDBA and you do not want to set the login credentials to SYSDBA in the Console, you can set SYSDBA credentials that will only be used for backup and recovery jobs.

When running backup and recovery jobs, these configurations will take precedence over the preferred credentials set in the Oracle Enterprise Manager Console.

1. Choose the **Backup Management** menu from the **Object** menu.
2. Choose **Create Backup Configuration** from the **Backup Management** menu.
3. On the Preferred Credentials page, set your credentials.

Registering Later Databases with the Recovery Catalog

To use the same configuration on subsequent databases, follow the steps below.

1. Choose the **Backup Management** menu from the **Object** menu.
2. Choose **Maintenance** from the Backup Management menu so that you can register the target database in the recovery catalog.
3. In the Operation Choice Page, choose **Maintain a target database in the recovery catalog** and then **Register Database**. Click Next.
4. In the Configuration Page, ensure that you are using the Backup Configuration you have created/modified. Click Next.
5. In the Multiple Targets Page, choose the target database(s) to submit the job. Click the Finish button. At this point, the database registration is sent as a job to the Oracle Enterprise Manager Job system.

6. When the summary screen appears, click OK to complete the operation.
7. In the Console, click the registration job in the Active Jobs list to see the current state of the job.

Once the job is completed, check the job history in the Jobs property sheet to make sure that the job completed successfully.

Resynchronizing the Recovery Catalog with the Target Database

RMAN automatically resynchronizes the catalog on each backup. So, manual resynchronizing of the catalog is needed only in rare cases.

The recovery catalog is not updated automatically when a log switch occurs or when an log is archived. Also, any structural changes to the target database would require re-synchronization of the Recovery Catalog.

To resynchronizes the Recovery Catalog with the target database so that the recovery catalog is updated with current information from the control file of the target database:

1. Choose the **Backup Management** menu from the **Object** menu.
2. Choose **Maintenance** from the **Backup Management** menu so that you can register the target database in the recovery catalog.
3. In the Operation Choice Page, choose **Maintain a target database in the recovery catalog** and then **Resynchronize catalog**.
4. In the Multiple Targets Page, choose the target database(s) to submit the job. Click the Finish button. At this point, the database registration is sent as a job to the Oracle Enterprise Manager Job system.
5. When the summary screen appears, click OK to complete the operation.
6. In the Console, click the registration job in the Active Jobs list to see the current state of the job.

Once the job is completed, check the job history in the Jobs property sheet to make sure that the job completed successfully.

Setting Up the Recovery Catalog

This section contains the following topics:

- 9i Procedure
- Pre-8i and 8i Procedures

- Create a Tablespace
- Create a User
- Run the catrman Script (for pre-8i only)
- Create the Recovery Catalog (for 8i)

9i Procedure

For Oracle9i, a recovery catalog is created if you specify for the Enterprise Manager repository to be located in a local database.

If you want to create the recovery catalog user and schema with a script, follow the procedure below:

1. Issue the following SQL statements

```
CREATE TABLESPACE "CATTBS"  
  LOGGING  
  DATAFILE 'CATTBS.dbf' SIZE 10M AUTOEXTEND  
  ON MAXSIZE UNLIMITED EXTENT MANAGEMENT LOCAL;  
CREATE USER rman IDENTIFIED BY rman  
  DEFAULT TABLESPACE CATTBS  
  TEMPORARY TABLESPACE TEMP QUOTA UNLIMITED ON CATTBS;  
GRANT RECOVERY_CATALOG_OWNER TO rman;  
GRANT CONNECT, RESOURCE TO rman;
```

2. Start RMAN from command-line using

```
rman CATALOG rman/rman@<database alias>
```

3. After RMAN has started, issue the following command

```
CREATE CATALOG;
```

The creation of the catalog could take several minutes.

Pre-8i and 8i Procedures

To set up a recovery catalog, you must complete the following procedures:

- Create a Tablespace
- Create a User
- Run the catrman Script (for pre-8i only)
- Create the recovery catalog (for 8i)

Create a Tablespace From the Oracle Enterprise Manager Console, perform the following steps to create a tablespace:

1. Expand the Database folder and select the database from the Navigator.
2. Log in as a user with the DBA role.
3. Select Create from the Object menu.
4. Select Tablespace from the Object menu tree and click the Create button.
5. In the Name field of the Create Tablespace property sheet, type the name of the new tablespace. For example, `cattbs`.

A datafile is added by default to the tablespace with a default name and default size, but if you can edit them. In Datafile section, check that the Name field contains the complete path and name of the datafile. For example:

`c:\orant\oradata\cattbs.dbf`. In the Size section, change the size of the new datafile if you want. For example, 10 M.

6. Click the Create button in the Create Tablespace property sheet.

Create a User From the Oracle Enterprise Manager Console, perform the following steps to create a user:

1. Select Create from the Object menu. Note: You must be connected to a target database; otherwise an error message appears.
2. Select User from the Object menu tree and click the Create button. The Create User property sheet appears.
3. In the General Page of Create User property sheet, fill in the following information.
 - a. In the Name field, enter the name of the new user. For example, `rman`.
 - b. In the Password and Confirm Password fields, enter a password. For example, `rman`.
 - c. Choose the default tablespace. For example, `CATTBS`.
 - d. Choose the temporary tablespace. For example, `TEMPORARY_DATA` or `TEMP`.
4. In the Role Page, grant the `RESOURCE` and `RECOVERY_CATALOG_OWNER` roles to the user.
5. In the Quotas Page, specify an unlimited quota for the default tablespace. In this example, the default tablespace is `CATTBS`.

6. Click the Create button after specifying the requisite parameters.

Run the catrman Script (for pre-8i only) Follow the instructions below:

1. Start SQL*Plus on the machine where the database resides.
2. Log in as the recovery catalog user you created. For example, user=RMAN, password=RMAN.
3. Execute `spool create_rman.log` to create a log file that you can use to check for errors.
4. Execute catrman script located in the `Oracle_Home/rdbms/admin` or `Oracle_Home/ora81/rdbms/admin` directory. If you are using SQL*Plus worksheet, type `@<full pathname and name of script>` in the SQL*Plus Worksheet input panel and press the Execute button. For example, type `@Oracle_Home/ora81/rdbms/admin/catrman` and press the Execute button.

Create the Recovery Catalog (for 8i) Follow the instructions below

1. Connect to the recovery catalog from the operating system command line on the target machine:

```
%> rman catalog rman/rman@<service name for database>
```

The correct output is shown below:

```
RMAN-06008: connected to recovery catalog database  
RMAN-06428: recover catalog is not installed
```

2. Issue the create catalog command to create the catalog.
 - For UNIX, type: `RMAN> create catalog tablespace 'cattbs';`
 - For Windows NT, type: `RMAN> create catalog tablespace 'CATTBS';`

Starting Up the Database

The following are description of various database modes.

Started (No Mount)

The instance starts, but does not mount the control file or open the database. This mode is used to recreate a control file or recreate the database from scratch. The database is not open, and therefore access by users is not permitted.

Mounted

An instance that is started and has the control file associated with the database open. You can mount a database without opening it; typically, you put the database in this state for maintenance or for restore and recovery operations. The database is not open, and therefore access by users is not permitted.

Cold backups (closed backups) are taken while the database is closed. Typically, closed backups are also whole database backups. If you closed the database cleanly, then all the files in the backup are consistent. If you shut down the database using a SHUTDOWN ABORT or the instance terminated abnormally, then the backups are inconsistent.

Open

The instance is started, the database mounted and opened. This mode is the default startup mode which allows any valid user to connect to the database and perform typical data access operations. Hot backups (online backups) are performed when the database is opened.

Cold backups (closed backups) are taken while the database is closed. Typically, closed backups are also whole database backups. If you closed the database cleanly, then all the files in the backup are consistent. If you shut down the database using a SHUTDOWN ABORT or the instance terminated abnormally, then the backups are inconsistent.

This section contains the following topics:

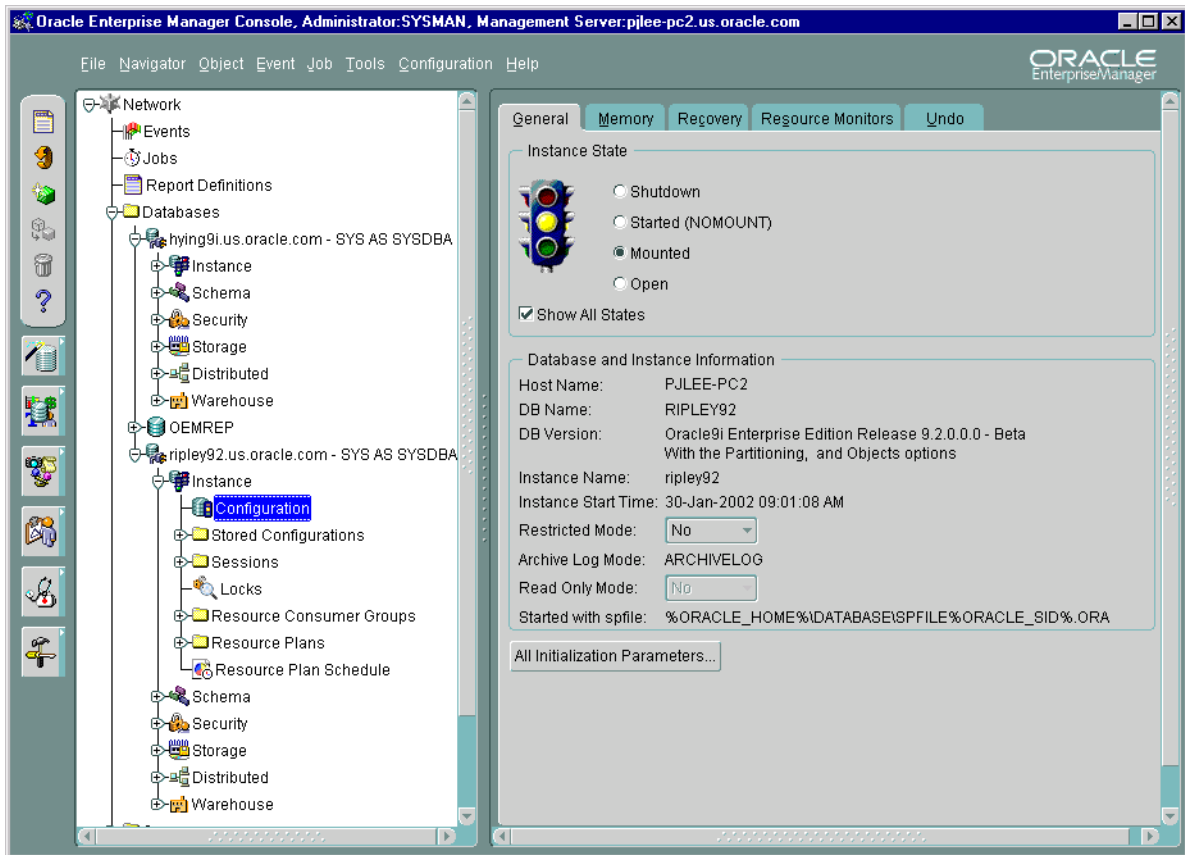
- Starting the Database in Mount Mode
- Starting the Database in Open Mode
- Placing the Database in Mount

Starting the Database in Mount Mode

Startup is the action of placing an Oracle instance in a state that allow users to access the database. To change the archiving status of the database, the database should be in the Mount state (mounted and closed).

Note: To start up a database, you must be connected to the database as SYSDBA.

Figure 11–44 Mounted State



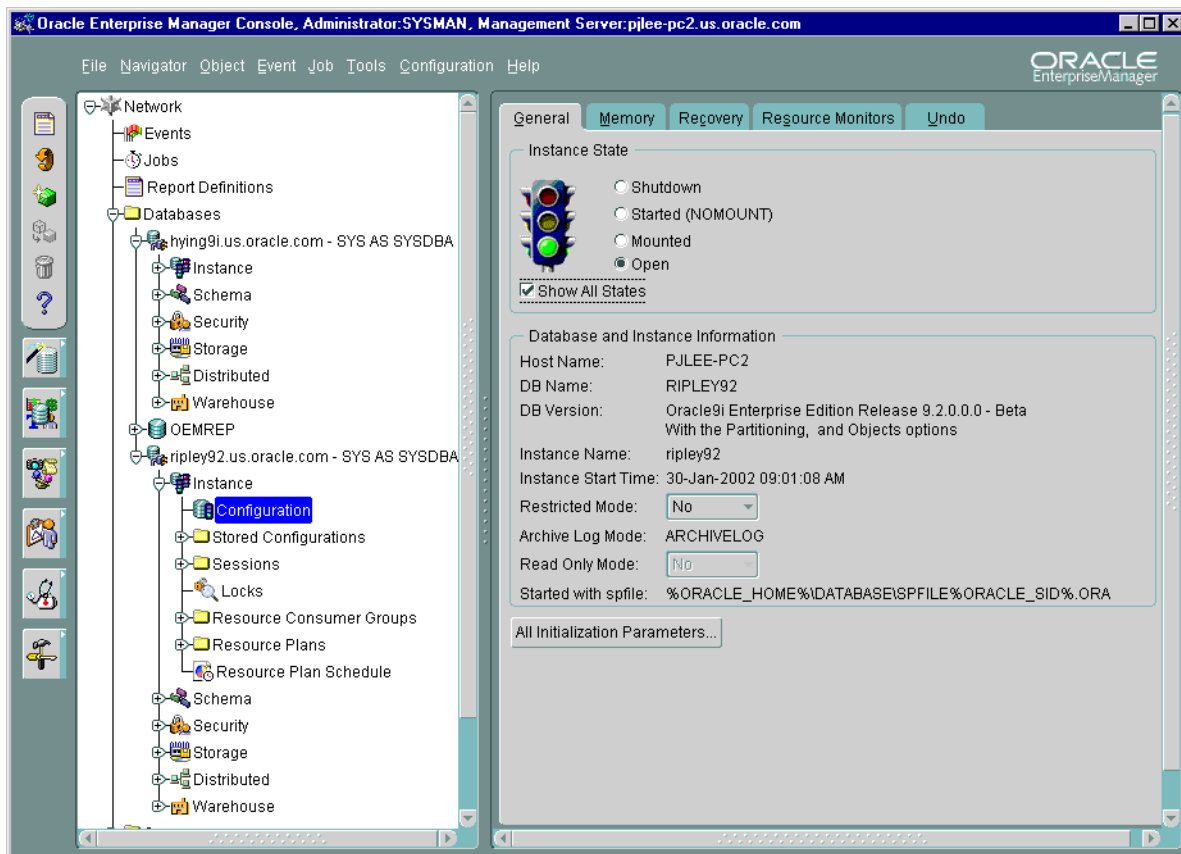
1. Select the Configuration under Instance in the Oracle Enterprise Manager Console.
2. To view all the database states, select **Show All States** box in the General page. Otherwise, only **Shutdown** and **Open** are shown
3. Select the **Mounted** startup option, which will start the instance and mount the database, but leave it closed. This mode is used to change the archiving status of the database, perform recovery, and for datafile recovery. The database is not open, and therefore access by users is not permitted.
4. Click Apply button.

Starting the Database in Open Mode

Startup is the action of placing an Oracle instance in a state that allow users to access the database.

Note: To start up a database, you must be connected to the database as SYSDBA.

Figure 11–45 Open State



1. Select the Configuration under Instance in the Oracle Enterprise Manager Console.
2. Select the **Open** startup option in the General page.
3. Click Apply button.

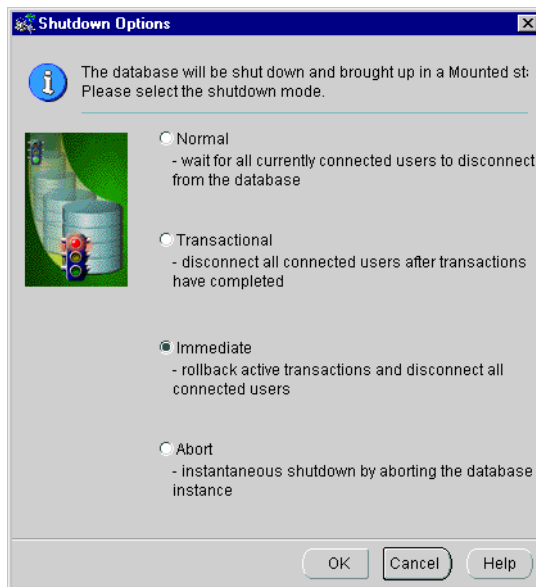
Placing the Database in Mount

Note: To shut down, then start and mount a database, you must be connected to the database as SYSDBA.

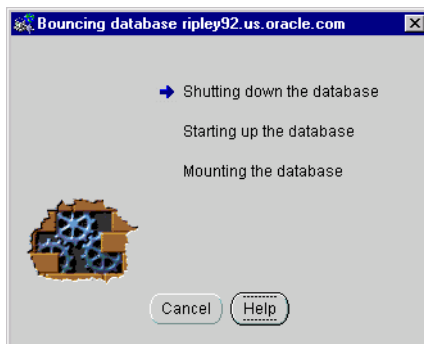
1. Select the Configuration under Instance in the Oracle Enterprise Manager Console.
2. To view all the database states, select **Show All States** box in the General page. Otherwise, only **Shutdown** and **Open** are shown.
3. Select the **Mounted** option.
4. Click Apply button.
5. When the Shutdown Options page appears, select **Immediate** as the Shutdown Mode and click OK.

Shutdown is the action of taking an Oracle instance from a state that allows users to access the database to a dormant state; when the database is shut down, we say that it is closed. Shutdown terminates the processes required for users to access the database, and it releases the portion of your computer's memory within which Oracle was operating.

Current client SQL statements are terminated immediately. Uncommitted transactions are rolled back. Active transactions are rolled back and all users are disconnected.

Figure 11–46 Shutdown Options

A progress dialog appears, giving you the status of the operation.

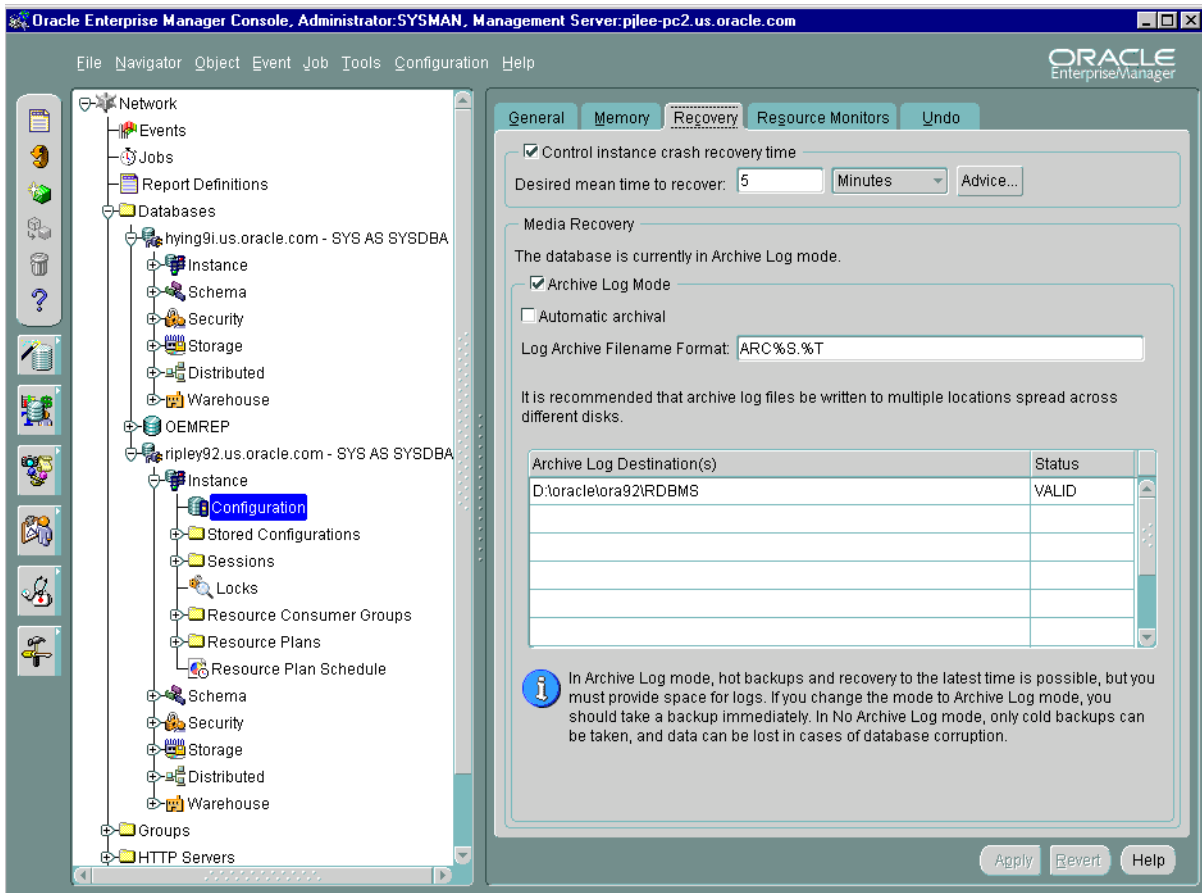
Figure 11–47 Progress Dialog

Press **Close** to dismiss the dialog once the database is mounted.

Setting the Database in ARCHIVELOG Mode

ARCHIVELOG mode permits complete recovery from disk failure as well as instance failure, because all changes made to the database are permanently saved in an archived redo log. The database can also be backed up while it is open and available for use.

Figure 11–48 Archive Log Mode



1. Select the Configuration under Instance in the Oracle Enterprise Manager Console.

2. Select the **Archive Log Mode** in the Recovery page.
3. Click Apply button.

If the database is in the Open state (mounted and open), Oracle Enterprise Manager displays the Shutdown Options dialog box allowing you to shut down the database. You can restart the database in the Mount state before changing the ARCHIVELOG mode.

If the database is in the No Mount or Started state (not mounted), Instance Management asks if you want to open the database in the Mount state.

RMAN Job Script

Use the Run Rman Script feature to issue any command or script within Oracle Enterprise Manager that can also be called from the RMAN command line. The rman script will be run as a job through the Oracle Enterprise Manager Job System when you submit or schedule it.

Oracle Enterprise Manager's Job System enables the automation of standard and administrative tasks. With the Job System, you can create and manage jobs, schedule their execution, and view and share information about defined jobs with other administrators.

Some Examples

This section contains some examples of RMAN Backup. For more examples, refer to the *Oracle9i Recovery Manager User's Guide*.

Backing Up the Recovery Catalog

A single recovery catalog is able to store information for multiple target databases. Consequently, loss of the recovery catalog can be disastrous. Back up the recovery catalog with the same frequency that you back up the target database.

For example, if you make a weekly whole database backup of the target database, then back up the recovery catalog immediately after all target database backups. The backed up catalog contains a record of the target backup preceding it, so if you need to restore the catalog you can use it to restore a target backup.

1. Create backup configurations that do not use the recovery catalog so that the repository for the recovery catalog is the control file in the catalog database.
 - a. Use the Create Backup Configuration property sheet to configure a backup set for tape. On the Recovery Catalog page, choose "In the target database's controlfile" for where you want the backup information to be stored.
 - b. Use the Create Backup Configuration property sheet to configure a backup set for disk. On the Recovery Catalog page, choose "In the target database's controlfile" for where you want the backup information to be stored.
2. Make sure to set the `CONTROL_FILE_RECORD_KEEP_TIME` initialization parameter to a value that is high enough to store an adequate amount of historical backup data for the catalog.
3. Run the recovery catalog database in ARCHIVELOG mode so that you can do point-in-time recovery if needed.
4. Using the Maintenance Wizard, specify a backup and retention policy for the database with the following settings:
 - The control file autobackup feature is ON and a location for the control file specified.
 - The retention policy is set to a REDUNDANCY value greater than 1.

In the event that the catalog database and control file are destroyed, you can restore a control file from an autobackup and then use it to restore the catalog database.

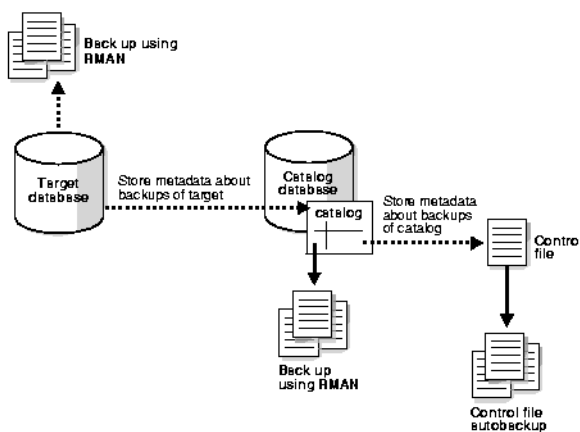
5. Backup the database and archivelogs to disk.

6. Backup the database and archivelogs to tape.

Run the Backup of the database at regular intervals.

With this strategy, the control file autobackup feature ensures that the recovery catalog database can always be recovered.

Figure 11–49 Recovery Catalog Backup



Recovery Without a Catalog

To restore and recover the database without using a recovery catalog, follow the guidelines below:

- Use the Maintenance Wizard to configure a backup policy which has the **Choosing to automatically back up controlfile and server parameter file** option enabled. Enabling the control file autobackup feature causes RMAN to automatically back up the control file and also enables RMAN to restore the control file autobackup without access to a repository (either a recovery catalog or the target database control file)
- Keep all Recovery Manager backup logs

If you lose the current control files, then you can restore a control file autobackup even if you do not use a recovery catalog.

Skipping Tablespaces when Backing Up a Database

Use the Maintenance Wizard to configure a backup policy which has the **Excluding tablespaces in a backup** option selected. For more information, refer to "Excluding tablespaces in a backup" on page 11-49.

The exclusion condition applies to any datafiles that you add to this tablespace in the future.

This tablespace exclusion feature is useful when you do not want to make a specified tablespace part of the regular backup schedule, as in these cases:

- A tablespace is easy to rebuild, so it is more cost-effective to rebuild it than back it up every day.
- A tablespace contains temporary or test data that you do not need to back up.
- A tablespace does not change often and therefore should be backed up on a different schedule from other backups.

These tablespaces can still be backed up separately using a tablespace backup. This option allows you to exclude some extremely large but rarely modified tablespaces from a database backup, and back them up on a different schedule. This may reduce the amount of time required for a regularly scheduled database backup.

Backing Up Often-Used Tablespaces

Many DBAs find that regular whole database backups are not in themselves sufficient for a robust backup strategy. If you run in ARCHIVELOG mode, then you can back up the datafiles of an individual tablespace or even a single datafile. This option is useful if a portion of a database is used more extensively than others, for example, the SYSTEM tablespace and automatic undo tablespaces.

By making more frequent backups of the extensively used datafiles of a database, you avoid a long recovery time.

For example, you may make a whole database backup once every two weeks. If the database experiences heavy traffic during the week, then a media failure on Friday can force you to apply a tremendous amount of redo during recovery. If you had backed up your most frequently accessed tablespaces three times a week, then you could apply a smaller number of changes to roll the restored file forward to the time of the failure.

If you are running in automatic undo management mode, then be sure to regularly back up your undo tablespaces.

If you run in manual undo management mode, then be sure to regularly back up all tablespaces containing rollback segments.

Specifying the Device Type

If you want to change from using an automatic channel tape backup to backing up the database to disk using the default configured DISK channel, refer to "Specifying Disk Channel Device for Backup Set" on page 11-54.

Restarting a Backup

Assume that you back up the database and archived logs every night to tape and you limit each backup set to two datafiles, so it produces multiple backup sets. Assume that the media management device crashes halfway through the backup and is then restarted. The next day you discover that only half the backup sets completed.

In this case, you can run to back up again, making sure that you have a backup policy configured for optimization. Database backups and archivelog backups may be optimized by skipping files that have not changed since the last backup. For more information, refer to "Choosing to optimize the backup." on page 11-48.

Spreading a Backup Across Multiple Disk Drives

When backing up to disk, you can specify a format if you need to spread the backup across several drives for improved performance. In this case, allocate one DISK channel for each disk drive and specify the format string. Specify both the directory path and the unique file name format so that the filenames are on different disks.

1. In the Channels page of the Create Backup Configuration property sheet, specify the following format for Channel1

```
<directory in disk1>%U
```

2. Insert Channel2 and specify the following format for Channel2

```
<directory in disk2>%U
```

3. Insert Channel3 and specify the following format for Channel3

```
<directory in disk3>%U
```

For more information, refer to "Specifying Disk Channel Device for Backup Set" on page 11-54.

Specifying the Size of Backup Sets

When making backups, RMAN divides the total number of files requiring backups by the number of allocated channels to calculate the number of files to place in each backup set.

Set the size of the backup set in the Backup Set page of the Create Backup Configuration Property Sheet.

Performing a Non-Cumulative Incremental Backups

A non-cumulative incremental backup contains only blocks that have been changed since the most recent backup at the same level or lower.

1. Perform an incremental backup at level 0 so that the backup contains all used blocks. The first incremental backup must be a level 0 backup.
2. Perform an incremental backup at level 1 or higher. An incremental backup at level 1 or higher will contain all blocks changed since the most recent level 1 backup. If no previous level 1 backup is available, then RMAN copies all blocks changed since the base level 0 backup.

If you add a new datafile or tablespace to the database, then make a level 0 backup before making another incremental backup. Otherwise, the incremental backup of the tablespace or the database fails because Recovery Manager does not find a parent backup for the new datafiles. Perform a level 0 backup of a single tablespace.

Note that you can perform incremental backups in NOARCHIVELOG mode, but the backups must be consistent. Hence, you cannot take online incremental backups.

For information on incremental backups, refer to "An incremental backup" on page 11-13.

Performing Cumulative Incremental Backups

A cumulative incremental backup at level n contains only blocks that have been changed since the most recent backup at level $n - 1$ or lower. Cumulative backups require more storage space than differential backups, but they are preferable during a restore operation because only one backup for a given level is needed. Note that the first incremental backup must be a level 0 backup that contains all used blocks.

A cumulative backup at level 2 will contain all blocks changed since the most recent level 1 backup, copying all blocks changed since the base level 0 backup only if a previous level 1 is unavailable. In contrast to a cumulative backup, a differential

backup at level 2 will determine which level 1 or level 2 backup occurred most recently and copy all blocks changed since that backup.

For information on incremental backups, refer to "An incremental backup" on page 11-13.

Determining How Channels Distribute a Backup Workload

When you create multiple backup sets and allocate multiple channels, RMAN automatically parallelizes its operation and writes multiple backup sets in parallel. The allocated server sessions share the work of backing up the specified datafiles, control files, and archived redo logs.

RMAN automatically assigns a backup set to a device.

For information on setting channels, refer to "Creating a Backup Configuration" on page 11-52.

Backing Up in NOARCHIVELOG Mode

1. Put the database into the correct mode (MOUNT) for a consistent, whole database backup. Refer to "Placing the Database in Mount" on page 11-72.
2. Back up the database.

Keeping a Long-Term Backup

If you configure a retention policy, then you may want to exclude specified backups from this policy. For example, you may want to archive a consistent backup of the database once a year to serve as a historical record. This long-term backups does not function as a backup that you may perform recovery on, but an archived snapshot of data at a particular time.

To exempt a backup from the retention policy, specify the **The Keep this backup or copy until <specified time>** option. For more information, refer to "Overriding the Retention Policy" on page 11-26.

Displaying Backups that are Exempt from the Retention Policy

Query the v\$backup_set and v\$datafile_copy tables for the KEEP time information.

Optimizing Database Backup

Enable backup optimization. Refer to "Configuring a Backup Policy" on page 11-47. When specific conditions are met, RMAN skips backups of files that are identical to files that are already backed up.

Keyboard Navigation

Oracle Enterprise Manager supports standard keyboard navigation. Standard keyboard navigation includes the use of the tab key, mnemonics (using the Alt key and the underlined character), and accelerators (such as Alt+F4 to exit a window).

The following table contains keyboard actions that are not commonly known or for which there is no firm standard.

Keyboard Action	Result
F10 (release), Space	Drops the system menu for a window or dialog
When focus is on a selected tree item or a selected table item, type Shift + F10	Drops the context menu for the selected item
With focus on an edit field within a table, type in a new value and press Enter	Accepts new value is accepted and the focus is moved to the next row of the table.
With focus on a drop down list within a table, press Space.	Changes the drop down list from opened to closed or from closed to open.

Accelerators for picture buttons are documented in the Help for the dialog or window where the picture buttons appear.

Firewalls and Virtual Private Networks

This appendix discusses how to configure firewall and virtual private networks (VPNs) to allow communication between various components of Oracle Enterprise Manager.

This appendix covers the following topics:

- Firewall Communication for Enterprise Manager
- Virtual Private Network Configuration for Enterprise Manager
- Running the Console in Standalone Mode
- Performance Manager, Capacity Planner, and Firewalls

Firewall Communication for Enterprise Manager


Firewalls protect a company's 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.

The various components of Enterprise Manager 9i (Console, Oracle Management Server, and Intelligent Agents) can be deployed on different nodes, which in turn can be separated by firewalls. This section describes how firewalls can be configured to allow communication between the different components of Enterprise Manager. The three most common deployments are covered:

- *Firewall Between the Console and Management Server*
- *Firewall Between the Management Server and Agent(s) on Monitored Nodes*
- *Firewalls and Network Address Translation (NAT)*

Firewall Between the Console and Management Server

In this configuration, the Enterprise Manager Console and Management Server are separated by a firewall.

Figure B-1 Console and Management Server on Opposite Sides of a Firewall

To enable network communication between the Enterprise Manager Console and Management Server, several network ports must be opened in the firewall to allow TCP traffic. The port range of 7771-7777 covers all these. If the Console is running in a browser, then the firewall needs to allow HTTP traffic over port 3339 from the Enterprise Manager Website HTTP server to any browser client. Functional assignments for these ports are shown in the following table.

Table B-1 Port Usage

Port Number	Usage
3339	Communication between the Enterprise Manger HTTP server and the Enterprise Manager browser client.
7771, 7773, 7776	Communication between the Enterprise Manager Console and the Management Server.
7774	Communication between the Oracle Applications Manager and the Management Server.

Table B-1 Port Usage


Port Number	Usage
7775, 7777	Communication between the paging server and the Management Server.

Special configuration is not required for either the Console or the Management Server in this case.

Firewall Between the Management Server and Agent(s) on Monitored Nodes

In this configuration, the Intelligent Agent that runs on the managed node and the Management Server are on opposite sides of the firewall, as shown in the following illustration.

Figure B-2 Firewall Between the Management Server and Agent



To enable network communication between the Management Server and Intelligent Agents on managed targets, several network ports must be opened in the firewall to allow TCP traffic. Functional assignments for these ports are shown in the following table.

Table B–2 Port Usage

Port Number	Usage
1748, 1754	Management Server communicating with the Agent to discover new targets.
7772	Agent communicating with the Management Server.
7773	Agent communicating with the Management Server via SSL.

No special setup and configuration is required for the Management Server or Intelligent Agents in this situation.

If the Management Server and administered database (or other managed target) are separated by a firewall, then the Management Server acts as a proxy for the Enterprise Manager Console, resulting in the remote database viewing the Management Server as the client. For this reason, there must be a SQL*Net proxy between the Management Server and the administered database. If the Console is launched in Standalone Mode, there must be a SQL*Net proxy between the Console and the Management Server, and between the Management Server and ALL collections services (Data Gatherer) connections.

Firewalls and Network Address Translation (NAT)

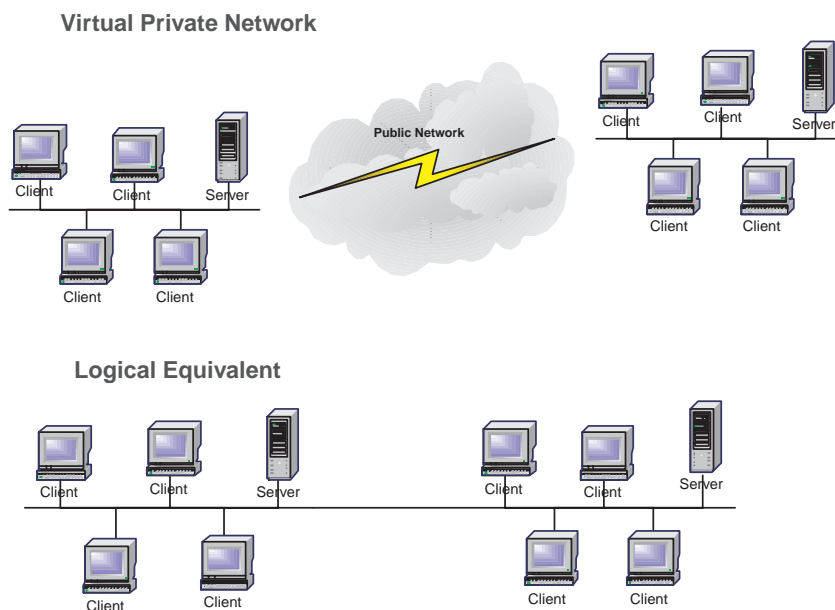
Some firewalls use a feature called Network Address Translation (NAT). This feature masks the true IP address of a client by translating it to a different IP address. Packets sent from a remote client to a server through the firewall will be known to the server by this translated address. As the client and server communicate, the NAT software handles the mapping of the true IP address to its translated address. Of the two Enterprise Manager configurations previously discussed, only an Enterprise Manager Console and Management Server can be separated by firewalls using NAT. No changes are required for Enterprise Manager to support NAT in this configuration.

The Management Server and Intelligent Agent cannot be separated by firewalls using NAT because the Management Server and Agent communication includes the other's host address information, which is stored in the data packet rather than in the IP header. Since NAT only looks for (and translates) addresses in the IP header, NAT will not work with Management Server/Agent communication.

Virtual Private Network Configuration for Enterprise Manager

Virtual Private Networks (VPNs) allow remote employees to connect in a secure fashion to a corporate server located in the corporate Local Area Network (LAN) using the routing infrastructure provided by a public network (such as the Internet). From the user's perspective, the VPN is a point-to-point connection between the user's computer and a corporate server. The nature of the intermediate network is irrelevant to the user because it appears as if the data is being sent over a dedicated private link.

Figure B-3 Virtual Private Network




In order to provide a secure point-to-point channel of communication, VPN software includes services such as user authentication and data encryption. It also implements security standards defined by the IP Security (IPSEC) protocol. IPSEC is a series of guidelines for the protection of Internet Protocol (IP) communications. It specifies standardized ways for securing private information transmitted over public networks. Communication between security systems developed by different vendors is possible if they comply with the IPSEC standards.

To create secure VPNs, VPN software typically operates in IPSEC Tunnel Mode. In this mode, data sent from a client is first encrypted and then encapsulated before being transmitted over an insecure, public network such as the Internet. Upon arriving at its destination, VPN software unpacks, decrypts and authenticates the data received, then forwards it on to its final destination.

Many e-businesses use both VPNs and firewalls as part of their security infrastructure. In these configurations, the firewall must allow IPSEC-compliant traffic to pass through (port 500 is used by default). Application data that is sent via VPN is first encapsulated and tunneled through port 500 in the firewall, unpacked, and sent to its final destination. Targets that have been set up to use VPN thus avoid having to open up additional ports in the firewall. Applications that run on VPN-enabled nodes can also communicate safely and securely across the firewall.

VPN Connections Between the Enterprise Manager Client and Management Server

As previously discussed, VPNs that comply with IPSEC standards allow the secure transfer of information over the internet: Remote clients can connect to a secure server with minimum configuration and maximum security. It is also possible to use VPNs in conjunction with firewalls. The following example shows a VPN environment with the Enterprise Manager Console and the Management Server on opposite sides of the firewall.

Figure B-4 Firewall Configuration In a VPN Environment

In this example, both the Console and Management Server machines have VPN software configured to provide a secure communication channel between the two. Specifically, the machine running Enterprise Manager client must have the VPN client software installed. The machine running the Management Server must have the VPN gateway software installed. Additionally, the firewall must be configured to allow only IPSEC traffic (IPSEC by default uses port 500). In this configuration, all the network traffic between the Console and the Management Server will be tunneled automatically through port 500 by the VPN software.

No additional configuration is required for Enterprise Manager components since the VPN software handles communication tasks automatically.

When the Enterprise Manager Console is launched, the user may be prompted by a VPN client software dialog to enter user security information. Once a valid username and password are provided to the VPN client, subsequent communication between the Console and Management Server across the virtual network will appear seamless.

No additional changes are required for the firewall configuration if IPSEC traffic is already allowed.

VPN Connections Between the Management Server and Intelligent Agents

Some VPN providers may allow server processes on different nodes to communicate. In these configurations, it is possible to deploy the Management Server on one VPN-enabled node and the Agent on another VPN-enabled node. The same principles as described in the previous section apply. It is important to note that communication between the Management Server node and Agent node is bi-directional, so each would need to function as both a VPN client and VPN server. Hence, both the VPN client and server software must be installed on each node.

Note: Sun Solaris version 8 supports VPN server process communication. Refer to the Solaris System Administrator's Guide for more details.

Running the Console in Standalone Mode

If the Enterprise Manager Console is launched in Standalone mode, the Console connects directly to a managed target (e.g. database) in the traditional client-server mode, using SQL*Net to communicate. If a firewall separates the Console and its target, several options are available. These include:

1. Use VPN software on the Enterprise Manager Console node and target node.

In this case, the setup will be similar to the Console - Management Server using VPN setup as described in the previous section.

2. Use Oracle Net features that support firewall configurations. These include:

- Oracle Net Firewall Proxy.

Several firewall vendors (e.g. CheckPoint) include an Oracle Net proxy capability which allows SQL*Net traffic to pass through its firewalls. This functionality is primarily available from the firewall vendors.

- Oracle Net's Connection Manager

Oracle Net's Connection Manager provides database connection pooling capabilities. Client applications connect to Connection Manager which in turn redirects the connection to the database. In this case, firewalls need to allow connections from the client to Connection Manager.

Refer to the Oracle Net documentation for more details.

Performance Manager, Capacity Planner, and Firewalls

The Performance Manager or Capacity Planner applications can be launched separately in order to connect directly to the Intelligent Agent (which incorporates the data collection services) on a target node. If a firewall separates Performance Manager and the Agent, or Capacity Planner and the Agent, then the firewall needs to be configured as follows:

- Port 1808 over TCP: Used for communication between Performance Manager and Agent; or Capacity Planner and the Agent.
- Port 1809 over TCP: Used for communication between the Performance Manager and the Agent over SSL or for Capacity Planner and the Agent over SSL

No additional configuration is required for Performance Manager, Capacity Planner or the Intelligent Agent.

OEMUTIL Utility

OEMUTIL is a command-line utility that performs various job-related and event-related functions, providing a command-line alternative to performing these same operations from within the Enterprise Manager Console. This allows job and event-related operations to be submitted as batch jobs which can run unattended and later checked for completeness.

This chapter provides information on how to setup and use the OEMUTIL utility. The following topics are discussed:

- Starting OEMUTIL
- Using OEMUTIL
- OEMUTIL Commands

Starting OEMUTIL

To enable OEMUTIL, set the environment variable `ORACLE_OEM_CLIENTTRACE` to "true" at the command line.

Windows NT or 2000:

```
>set ORACLE_OEM_CLIENTTRACE=true
```

UNIX:

```
>setenv ORACLE_OEM_CLIENTTRACE=true
```

Using OEMUTIL

OEMUTIL can perform a single function or command (e.g. submit a new Job) or perform multiple commands in succession. For example, you want OEMUTIL to submit job A and B from the job library, and then deregister Event C.

Performing a single command:

To use OEMUTIL to execute a single task, run OEMUTIL at the command line using the following syntax:

```
> oemapp oemutil <username>/<password>@<oms> <command> <parameters>
```

where:

- `<username>` is a valid Enterprise Manager administrator
- `<password>` is the Enterprise Administrator's password
- `<oms>` is the name of the machine running the Oracle Management Server
- `<command><parameters>` is the OEMUTIL command and its associated parameters.

Performing Commands in Succession:

You can use OEMUTIL to execute multiple commands in succession by first creating a single file consisting of the commands you want to execute and run OEMUTIL using the following syntax:

```
> oemapp oemutil -cmdfile <command file name>
```

where:

- `-cmdfile` is the option that instructs OEMUTIL to accept a single file containing the commands to be executed.
- `<command file name>` is the name of the ASCII file containing the various commands you want to execute. The commands in the file will be executed in the order they are listed. The result of each command execution will be displayed sent to standard output. If one command in the sequence fails, then the failure is displayed and OEMUTIL proceeds to execute the next command in sequence.

The command file may contain any number of commands. Each command must be on a separate line along with its corresponding arguments. Command arguments may be quoted if they have white space or special characters in them such as job and event names. Characters can be "escaped" by using the backslash character (`\`). For example, to insert a literal quote, use `\`. A command file may look something like the following example.

Example C-1 Sample Command File

```
omsCredentials sysman/sysman@myomsmachine
submitJob fileListing SYSMAN:dbs|oracle_sysman_group ls "-l"
submitJobFromLibrary backupJob sysman mypc.us.oracle.com
deregisterEvent spaceEvent sysman o817.mypcl
```

OEMUTIL Commands

The following section summarizes all commands and associated parameters for OEMUTIL.

omsCredentials

The *omsCredentials* command specifies the credentials to use when logging into the Oracle Management Server. Logging in to the Management Server is required before any other command can be run. All subsequent commands in the command file will use these same credentials until the next *omsCredentials* command is encountered. There can be multiple *omsCredentials* commands in a single command file, allowing you to use a one batch command file to perform operations using any number of Management Servers.

Syntax

If you are going to use a batch command file with OEMUTIL, then the *omsCredentials* command is required and *must be the first command* in the command file.

If you are issuing commands directly to OEMUTIL (not providing a batch file), then the *omsCredentials* command is not required.

```
omsCredentials <username>/<password>@<oms>
```

where:

■ .

Table C-1 *omsCredentials* Command Parameters

Parameter	Value
<username>	A valid Enterprise Manager administrator. This is the same account used to logon to the Enterprise Manager Console.
<password>	The Enterprise Administrator's password
<oms>	The name of the machine running the Oracle Management Server.

Example:

```
omsCredentials sysman/sysman@dlsun966
```

submitJob

The *submitJob* command allows you to submit a new job against node(s). The job will be submitted with a schedule set to "Immediate".

Syntax

If running OEMUTIL with this single command, the syntax is:

```
oemapp oemutil <username>/<password>@<oms> submitJob <jobName> <nodeName>  
<osCommand> <osParameters>
```

If specifying this command in a command file, the syntax is:

```
submitJob <jobName> <nodeName> <osCommand> <osParameters>
```

Table C-2 *submitJob Command Parameters*

Parameter	Value
username	Name of a valid Enterprise Manager user.
password	Corresponding password of Enterprise Manager user
oms	Name of the machine running the Management Server.
jobName	Name to assign to a job.
nodeName	<p>Name of the target node where the job will run. The name must have been previously discovered in the Console. The name should match the discovered node name in the Navigator tree in the console.</p> <p>You may also specify a group name if you'd like to run the job against the multiple nodes in a group. Specify the group using the syntax:</p> <pre>groupOwner:groupName oracle_sysman_group.</pre> <p>Example:</p> <p>To specify a group called dbms owned by sysman, use the syntax:</p> <pre>SYSMAN:dbms oracle_sysman_group</pre> <p>If you are specifying a group as a command-line parameter to OEMUTIL, then you need to enclose the group name in quotes. For example:</p> <pre>"SYSMAN:dbms oracle_sysman_group"</pre>
osCommand	The OS command to run.
osParameters	Any parameters associated with the OS command.

Success indicated by OEMUTIL means the job has been successfully submitted to the Management Server. Use the Enterprise Manager Console to monitor the actual status of the job.

To submit jobs that use different schedules such as 'on day of week' or 'on day of month' or jobs that run against other target types such as 'databases' or 'http servers', do the following:

1. Define the job using the Enterprise Manager Console.
2. Save the job to the Job Library.

- Using OEMUTIL, submit the job from the library via the `submitJobFromLibrary` command. See the next section for information on the `submitJobFromLibrary` command.

submitJobFromLibrary

The *submitJobFromLibrary* command allows you to submit a job that is defined in the Job Library. The submitted job uses the properties of the job as it is defined in the Job Library (tasks, parameters, schedule, permissions, etc) except for job targets. The job target needs to be specified as a parameter to this command.

Important: The job that is saved in the library must be assigned at least one valid target. However, when submitting this job from the library, the target you provide to the OEMUTIL command will be used.

Syntax

If running OEMUTIL with this single command, the syntax is:

```
oemapp oemutil <username>/<password>@<oms> submitJobFromLibrary <jobName>  
<ownerName> <targetName> <admin to be notified>
```

If specifying this command in a command file, the syntax is:

```
submitJobFromLibrary <jobName> <ownerName> <targetName> <admin to be notified>
```

Table C-3 *submitJobFromLibrary* Command Parameters

Parameter	Value
username	Name of a valid Enterprise Manager administrator.
password	Enterprise Manager administrator's password.
oms	Name of the machine on which the Management Server is running.

Table C-3 *submitJobFromLibrary Command Parameters*

Parameter	Value
jobName	<p>Name of the job to be submitted. The name should match the name of the job as it is defined in the Job Library.</p> <p>If the job name has space, whitespaces or other special characters, then you will need to use the command file option. Put quotes around the name containing spaces, whitespaces or other special characters.</p> <p>Example: "ORCL backup_job"</p>
ownerName	Name of the owner of the job to be submitted.
targetName	<p>Name of the target against which the job is to be submitted. The target must have been discovered by the Intelligent Agent. The target name must match the name displayed in the Console Navigator tree.</p> <p>The target specified in this parameter will be used when submitting the job. Any targets predefined in the job as it exists in the Job Library will be ignored.</p> <p>To submit the job against multiple targets, define a group that contains these targets, and specify the group name as the target. Specify the group using the syntax:</p> <pre>groupOwner:groupName oracle_sysman_group.</pre> <p>Example: To specify a group called "dbs" owned by sysman, use the syntax:</p> <pre>SYSMAN:dbs oracle_sysman_group</pre> <p>If you are specifying a group as a command-line parameter to OEMUTIL, then you need to enclose the group name in quotes.</p> <p>Example: "SYSMAN:dbs oracle_sysman_group"</p>
admin	(optional) Name of the Enterprise Manager Administrator for whom Notify permissions are to be set. This Administrator must have at least View permissions for the job.

Success indicated by OEMUTIL means the job has been successfully submitted to the Management Server. Use the Console to monitor the actual status of the job.

registerEventFromLibrary

The *registerEventFromLibrary* command allows you to register an Event that is defined in the Event Library.

Important: The event that is saved in the library must have at least one valid target. However, when registering this event from the library, the target provided to the OEMUTIL command will be used.

Syntax

If running OEMUTIL with this single command, the syntax is:

```
oemapp oemutil <username>/<password>@<oms> registerEventFromLibrary <eventName>
<ownerName> <targetName> <admin to be notified>
```

If specifying this command in a command file, the syntax is:

```
registerEventFromLibrary <eventName> <ownerName> <targetName> <admin to be notified>
```

Table C-4 *registerEventFromLibrary* Command Parameters

Parameter	Value
username	Name of a valid Enterprise Manager administrator.
password	Enterprise Manager administrator's password.
oms	Name of the machine on which the Management Server is running.
eventName	Name of the event to be submitted. The name should match the name of the event as it is defined in the Event Library. If the event name has space, whitespaces or other special characters, then you will need to use the command file option. Put quotes around the name containing spaces, whitespaces or other special characters. Example: "Check Tablespace usage"
ownerName	Name of the owner of the event to be submitted.

Table C-4 registerEventFromLibrary Command Parameters

Parameter	Value
targetName	<p>Name of the target against which the event is to be registered. The target must have been discovered by the Intelligent Agent. The target name must match the name displayed in the Console Navigator tree.</p> <p>To register the event against multiple targets, define a group that contains these targets, and specify the group name as the target for this event.</p> <p>Specify the group using the syntax:</p> <pre>groupOwner:groupName oracle_sysman_group.</pre> <p>Example:</p> <p>To specify a group called dbs owned by sysman, use the syntax:</p> <pre>SYSMAN:dbs oracle_sysman_group</pre> <p>If you are specifying a group as a command-line parameter to OEMUTIL, then you need to enclose the group name in quotes.</p> <p>Example:</p> <pre>"SYSMAN:dbs oracle_sysman_group"</pre>
admin	<p>(optional) Name of the Enterprise Manager Administrator for whom Notify permissions are to be set. This Administrator must have at least View permissions for event.</p>

Success in using this command means the event has been sent to the Management Server for registration. To confirm that the event has been registered with the Intelligent Agent, use the Enterprise Manager Console to verify that the status of the event is "Registered".

deregisterEvent

The *deregisterEvent* command allows you to deregister an Event.

Syntax

If running OEMUTIL with this single command, the syntax is:

```
oemapp oemutil <username>/<password>@<oms> deregisterEvent <eventName> <owner>  
<targetName> <targettype>
```

If specifying this command in a command file, the syntax is:

```
deregisterEvent <eventName> <owner> <targetName> <targettype>
```

Table C-5 *deregisterEvent Command Parameters*

Parameter	Value
username	Name of a valid Enterprise Manager administrator.
password	Enterprise Manager administrator's password.
oms	Name of the machine on which the Management Server is running.
eventName	Name of the event to be deregistered. If the event name has space, whitespaces or other special characters, then you will need to use the command file option. Put quotes around the name containing spaces, whitespaces or other special characters. Example: "Check Tablespace usage"
owner	Name of the owner of the event. The owner is the Enterprise Manager Administrator who originally registered the event.
targetName	Name of the target against which the event is to be deregistered. The target must have been discovered by the Intelligent Agent. The target name must match the name displayed in the Console Navigator tree.

Table C-5 *deregisterEvent Command Parameters*

Parameter	Value
targettype	The type of target against which the event is to be deregistered. The valid target types are: <ul style="list-style-type: none"> ▪ oracle_sysman_database (Target is a database) ▪ oracle_sysman_node (Target is a node) ▪ oracle_sysman_listener (Target is a net listener) ▪ oracle_sysman_cmanager (Target is a concurrent manager) ▪ oracle_sysman_ops (Target is an Real Application Cluster node) ▪ oracle_sysman_webserver (Target is an apache webserver) ▪ oracle_sysman_hotstandby (Target is a standby database)

changeCredentials

The *changeCredentials* command allows you to change preferred credentials for database targets in the Enterprise Manager repository.

This command does not apply to any other target type. It also does not change the credentials in the databases themselves, nor does it update the credentials in jobs and events that have already been submitted. Previously submitted jobs and registered events must be subsequently de-registered and re-registered in order for them to get the new credentials.

Syntax

If running OEMUTIL with this single command, the syntax is:

```
oemapp oemutil <username>/<password>@<oms> changeCredentials <EM username>
<targetName> <user> <password> <role>
```

If specifying this command in a command file, the syntax is:

```
changeCredentials <EM username> <targetName> <user> <password> <role>
```

Table C-6 *changeCredentials Command Parameters*

Parameter	Value
username	Name of a valid Enterprise Manager administrator.

Table C-6 *changeCredentials Command Parameters*

Parameter	Value
password	Enterprise Manager administrator's password.
oms	Name of the machine on which the Management Server is running.
EM username	Name of the Enterprise Manager Administrator whose database credentials are to be changed.
targetName	Name of the database target against which credentials are to be changed. If you are changing the DEFAULT database credentials, specify <default> as the targetName. If you use <default>, you might have to place quotes around the entry depending on your platform. For example, "<default>"
user	Name of the database user.
password	Password associated with database user.
role	Role associated with the database user: NORMAL, SYSDBA, SYSDBA, SYSDOPER The database user role must be specified.

For any command, if the target node name does not match the name displayed in the Console Navigator, you will receive the following error:

```
oracle.sysman.emSDK.client.omsClient.BadAttributeException: VD-3: Some of the targets specified have been removed or do not have agents.
```

If you specify an invalid database target using the `changeCredentials` command, the command will appear to succeed, but will have no effect.

Repository Views

The Enterprise Manger Configuration Assistant creates repository views whenever a new repository is created. These views are created on top of the repository tables. They enable users to view information from the repository without the danger of manipulating data the wrong way.

Administrator Views

SMP_VIEW_ADMINISTRATORS

This view shows a list of all administrators which have been defined on the machines. Information will be useful to system administrators.

The boolean values are represented as a numerical value so that you can change these values into a string appropriate for the Globalization Support settings which are currently active. These values can then be decoded into proper localized values when the report is run

Table D-1 SMP_VIEW_ADMINISTRATORS

Columns	Descriptions
ADMINISTRATOR_NAME	Name of defined database administrator that is used throughout the repository
SUPERUSER	Shows whether the database administrator is a superuser. Possible values: 1 (YES) and 0 (NO)
JOB_SYSTEM	Shows whether the database administrator can access the JOB system. Possible values: 1 (YES) and 0 (NO)
EVENT_SYSTEM	Shows whether the database administrator can access the EVENT system. Possible values: 1 (YES) and 0 (NO)

SMP_VIEW_ADMIN_CREDENTIALS

List of all the preferred credentials defined per administrator. Every target known in the navigator will have an entry in this view for each administrator who has setup preferred credentials for this target. This view is of interest for system administration.

Table D-2 SMP_VIEW_ADMIN_CREDENTIALS

Columns	Descriptions
ADMINISTRATOR_NAME	Name of defined database administrator used throughout the repository

Table D-2 SMP_VIEW_ADMIN_CREDENTIALS(Cont.)

Columns	Descriptions
TARGET_NAME	Name of the target. The target name uniquely identifies the target. Several targets can exist with the same name, each with a different type. The column TARGET_NAME links with TARGET_NAME in the SMP_VIEW_NODE_SERVICES view.
TARGET_TYPE	Type of the target
TARGET-NLS_TYPE	Type string used for globalization support of the reports.

SMP_VIEW_ADMIN_SETUP

This view shows the notification setup per administrator. This information is important to system administrators.

Table D-3 SMP_VIEW_ADMIN_SETUP

Columns	Descriptions
ADMINISTRATOR_NAME	Name of defined database administrator that is used throughout the repository. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view
EMAIL_ADDRESS	Email address defined for this administrator
EMAIL_TITLE	Title of the email message
PAGING_SERVICE	Name of the paging service to use for this administrator
PAGING_PIN	Paging pin code to use for this administrator

Service Views

SMP_VIEW_AGENTS

Shows topology information/details of all discovered Intelligent Agents on nodes in the Navigator. Manually added nodes are not included in this view.

Table D-4 SMP_VIEW_AGENTS

Columns	Descriptions
NODE_NAME	Name of the node. The NODE_NAME column links with NODE_NAME in the SMP_VIEW_NODES view.
AGENT_ALIAS	Name of the machine as recognized by the Agent
LAST_CHECKED	The last time and date this node was checked by the Management Server to see if it is reachable.
AGENT_TIMEZONE	Offset in hours This value can be subtracted from the all times reported by this Intelligent Agent.
AGENT_STATUS	Current status of the Intelligent Agent. Possible values: 1 (UP) and 0 (DOWN)
AGENT_STATE	Current condition of the Intelligent Agent. Possible values: 1 (GOOD) and 0 (BAD)
OMS_MACHINE	Name of the Management Server machine servicing this node This field is NULL if this is a manually added node. This name is the canonical name of the Management Server machine. The OMS_MACHINE column links with OMS_MACHINE from the SMP_VIEW_OMS_MACHINES view.
SEVERITY	Aggregated alert severity of this target: 0 - No events registered against this target 15 - All events are clear 20 - One or more events are in Warning state 25 - One or more events are in Alert state 115 - Agent down, previous state was Clear 120 - Agent down, previous state was Warning 125 - Agent down, previous state was Alert

SMP_VIEW_TARGETS

Shows topology information, a list of all targets managed by Intelligent Agents, which potential targets are present in the repository for jobs and events. This view is a subset of the V\$SMP_NODE_SERVICES view and contains only these targets

which have an Intelligent Agent running on the node. Manually added targets will not show up in this view.

Table D-5 SMP_VIEW_TARGETS

Columns	Descriptions
TARGET_NAME	Name of the target. The TARGET_NAME combination uniquely identifies the target. Several targets can exist with the same name, each with a different type. The TARGET_NAME column links with TARGET_NAME in the SMP_VIEW_NODE_SERVICES view.
TARGET_TYPE	Type identifier of this target. The type is the type returned by the Intelligent Agent. Only two manually added target types are possible: <ul style="list-style-type: none"> ■ Nodes (oracle_sysman_node) ■ Databases (oracle_sysman_database)
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
NODE_NAME	Name of node on which this target runs. The NODE_NAME column links with NODE_NAME in the SMP_VIEW_NODES view.
USERDATA	Parameters on how to contact this target
SEVERITY	Aggregated alert severity of this target. 0 - No events registered against this target 15 - All events are clear 20 - One or more events are in Warning state 25 - One or more events are in Alert state 115 - Agent down

SMP_VIEW_TARGET_PROPERTIES

Shows topology information, listing all the properties the Intelligent Agent has reported for this target. Only 8.1.7 or higher Agents will transmit target properties during discovery. Manually added nodes will never appear in this list, since they cannot have properties send over by the Intelligent Agent.

Table D-6 SMP_VIEW_TARGET_PROPERTIES

Columns	Descriptions
TARGET_NAME	Name of the target. The TARGET_NAME combination uniquely identifies the target. Several targets can exist with the same name, each with a different type. The TARGET_NAME column links with TARGET_NAME in the SMP_VIEW_TARGETS view.
TARGET_TYPE	Type identifier of this target. The type is the type returned by the Intelligent Agent.
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
NAME	Name of the property
VALUE	Value of the property

SMP_VIEW_NODES

Shows Topology Information, list of all known nodes in the repository. Both manually added nodes as well as nodes managed by an Agent (nodes defined in the navigator) will be represented in this view. For each node, the current state of the Intelligent Agent running on the node will be reported.

The boolean values are represented as a numerical value so that these values can change into a string appropriate for the Globalization Support settings which are currently active.

These values can then be decoded into proper localized values when the report is run

Table D-7 SMP_VIEW_NODES

Columns	Descriptions
NODE_NAME	Name of the node. The NODE_NAME field is the name of the node defined in the navigator. The same name will also appear in the target list as target type, oracle_sysman_node.

Table D-7 SMP_VIEW_NODES(Cont.)

Columns	Descriptions
DISCOVERY	Shows whether the node is discovered or manually added. Possible values: 1 (AUTOMATIC) and 0 (MANUAL) For all nodes with an operational Intelligent Agent, the DISCOVERY will be set to 1, the OMS_MACHINE and LAST_CHECKED fields will contain a value. For all manually added nodes, the fields DISCOVERY and AGENT_STATUS will be set to 0, and the fields OMS_MACHINE and LAST_CHECKED will be set to NULL.
LAST_CHECKED	The last time and date this node was checked by the Management Server to see if it is reachable.
AGENT_STATUS	Current status of the Intelligent Agent. Possible values: 1 (UP) and 0 (DOWN)
OMS_MACHINE	Name of the Management Server machine servicing this node. This field is NULL if this is a manually added node. This name is the canonical name of the Management Server machine. The OMS_MACHINE column links with OMS_MACHINE from the SMP_VIEW_OMS_MACHINES view.

SMP_VIEW_NODE_SERVICES

Shows topology information which lists all targets present in the repository. Both manually added as well as targets managed by an Intelligent Agent will be represented in this view.

Table D-8 SMP_VIEW_NODE_SERVICES

Columns	Descriptions
TARGET_NAME	Name of the target. The TARGET_NAME combination uniquely identifies the target. Several targets can exist with the same name, each with a different type.
TARGET_TYPE	Type identifier of this target. The type is the type returned by the Intelligent Agent.
TARGET-NLS_TYPE	Type string used for globalization support of the reports.

Table D-8 SMP_VIEW_NODE_SERVICES(Cont.)

Columns	Descriptions
NODE_NAME	Name of node on which this target runs. The NODE_NAME column links with NODE_NAME in the SMP_VIEW_NODES view.
USERDATA	Parameters on how to contact this target

Group Views

SMP_VIEW_GROUPS

A System Administration view which lists all groups defined in the repository (Which groups are defined in the navigator).

The boolean values are represented as a numerical value so that these values can change into a string appropriate for the Globalization Support settings which are currently active.

These values can then be decoded into proper localized values when the report is run

Table D-9 SMP_VIEW_GROUPS

Columns	Descriptions
GROUP_NAME	Name of the group
DESCRIPTION	User description of this group
BACKGROUND_IMAGE	Image to use as a background for this group. This is a physical path location to a file on the disc.
ICON_SIZE	Size of the icons. Boolean field: 0 (SMALL) or 1 (LARGE)

SMP_VIEW_GROUP_SERVICES

A System Administration view which lists all targets in a group, showing which targets are placed into which group. This view contains only the targets defined in the group at that level

Table D–10 SMP_VIEW_GROUP_SERVICES

Columns	Descriptions
GROUP_NAME	Name of the group. The GROUP_NAME column links with GROUP_NAME of the SMP_VIEW_GROUPS view.
TARGET_NAME	Name of the targets put in this group and if the target is nested in another group (Nested groups) The TARGET_NAME column links with TARGET_NAME of the SMP_VIEW_NODE_SERVICES view.
TARGET_TYPE	Type of target The type can be any of the registered types within the navigator. Since a group can be a target in another group
TARGET-NLS_TYPE	Type string used for globalization support of the reports.

SMP_VIEW_ALL_GROUP_ALL_SERVICES

A System Administration view which lists all targets in a group, recursively. If a group has a group as its member, all targets of the member group will be shown as well. The member groups themselves will not show up in the view as targets of the group.

Table D–11 SMP_VIEW_ALL_GROUP_ALL_SERVICES

Columns	Descriptions
GROUP_NAME	Name of the group. The GROUP_NAME column links with GROUP_NAME of the SMP_VIEW_GROUPS view.
MEMBER_NAME	Name of the member group. A group is also considered to be a member of itself The MEMBER_NAME column links with GROUP_NAME of the SMP_VIEW_GROUPS view.
TARGET_NAME	Name of the targets put in this group. The TARGET_NAME column links with TARGET_NAME of the SMP_VIEW_NODE_SERVICES view.
TARGET_TYPE	Type of target The type can be any of the registered types within the navigator.
TARGET-NLS_TYPE	Type string used for globalization support of the reports.

SMP_VIEW_TOP_GROUPS

Lists all views which are not part of other groups and have targets in them.

Table D-12 SMP_VIEW_TOP_GROUPS

Columns	Descriptions
GROUP_NAME	Logical name of the group
GROUP_OWNER	Administrator who created this group
DESCRIPTION	Description of the group
BACKGROUND_IMAGE	Background image used in the console
ICON_SIZE	Size of icons. Boolean field: 0 (SMALL) or 1 (LARGE)

Operational System Metrics

SMP_VIEW_OMS_MACHINES

An operational data view which lists all Management Servers currently connected to and working with the repository.

Table D-13 SMP_VIEW_OMS_MACHINES

Columns	Descriptions
OMS_MACHINE	Name of the machine
LAST_CHECKED	Timestamp of the last heartbeat of this Management Server, using the timezone of the Management Server machine.

SMP_VIEW_SESSIONS

An operational data view which lists all administrators logged on to this repository.

Table D-14 SMP_VIEW_SESSIONS

Columns	Descriptions
ADMINISTRATOR_NAME	Name of database administrator currently connected. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.
OMS_MACHINE	Name of Management Server machine the database administrator requested to login to. This name is the canonical name of the Management Server machine. The OMS_MACHINE column links with OMS_MACHINE from the SMP_VIEW_OMS_MACHINES view.
LOGIN_TIME	Time the administrator logged in on the Management Server machine

SMP_VIEW_NOTIFICATION_QUEUE

A view for diagnostics and debugging. It lists all the pending notifications still in the queue. The contents of this view will constantly change. If there are records in this view, the Management Server(s) do not have sufficient time to complete the work

Table D-15 SMP_VIEW_NOTIFICATION_QUEUE

Columns	Descriptions
SEQUENCE_NUM	Internal number to process the notifications. Since the Management Server processes the notifications sorted on SEQUENCE_NUM, this may give an estimate of what is to be processed and when.
SUBSYSTEM	The subsystem who has generated this operation. Can be: VdeEvent - Events VdjJob - Jobs VdmNotificationManager - Notifications
NOTIFICATION_TYPE	The type of notification for the subsystem

Table D–15 SMP_VIEW_NOTIFICATION_QUEUE(Cont.)

Columns	Descriptions
NODE_NAME	Name of the node responsible for this operation. The NODE_NAME column links with NODE_NAME in the SMP_VIEW_NODES view.
TARGET_NAME	Name of the target responsible for this operation. The TARGET_NAME column links with TARGET_NAME in the SMP_VIEW_TARGETS view.
TARGET_TYPE	Type of target
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
ADMINISTRATOR_NAME	Name of the administrator involved. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.
TIMESTAMP	Time the notification was received

SMP_VIEW_NOTIFICATION_HISTORY

A diagnostics and debugging view which lists all the processed notifications. The notification history can be cleaned. The information in this view is only a snapshot of what is still in the repository.

The boolean values are represented as a numerical value so that these values can change into a string appropriate for the Globalization Support settings which are currently active.

These values can then be decoded into proper localized values when the report is run

Table D–16 SMP_VIEW_NOTIFICATION_HISTORY

Columns	Descriptions
OBJECT_TYPE	The subsystem type who generated this notification Currently the notification types are: EVENT - Event subsystem JOB - Job subsystem

Table D–16 SMP_VIEW_NOTIFICATION_HISTORY(Cont.)

Columns	Descriptions
OBJECT_ID	The ID of the object within the sub-system. Events - EVENT_ID Jobs - JOB_ID The OBJECT_ID column links with either JOB_ID in the SMP_VIEW_JOBS view, or EVENT_ID in the SMP_VIEW_EVENTS view, depending on the value of OBJECT_TYPE.
TARGET_NAME	Name of the target responsible for this notification. The TARGET_NAME column links with TARGET_NAME in the SMP_VIEW_TARGETS view.
ADMINISTRATOR_NAME	Name of the administrator involved. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.
METHOD	Notification method. Currently these methods are defined: Email - SMTP Paging - Paging Server
TIMESTAMP	Time the notification was send
STATUS	Shows whether the notification was sent successfully Boolean field: 0 (NO) or 1 (YES)
OBJECT_STATUS	Status of the object

SMP_VIEW_OPERATION_QUEUE

A diagnostics and debugging view which lists all the pending operations still in the queue. The content of this view will constantly change. If there are records in this view, the Management Servers do not have sufficient time to complete the work

Table D–17 SMP_VIEW_OPERATION_QUEUE

Columns	Descriptions
OBJECT_ID	The ID of either a job or event

Table D-17 SMP_VIEW_OPERATION_QUEUE(Cont.)

Columns	Descriptions
SUBSYSTEM	The subsystem who has generated this operation. Can be: VdeEvent - Events VdjJob - Jobs VdmNotificationManager - Notifications
OPERATION_TYPE	The type of operation for the subsystem
NODE_NAME	Name of the node responsible for this operation. The NODE_NAME column links with NODE_NAME in the SMP_VIEW_NODES view.
TARGET_NAME	Name of the target responsible for this operation. The TARGET_NAME column links with TARGET_NAME in the SMP_VIEW_NODE_SERVICES view.
TARGET_TYPE	Type of target
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
ADMINISTRATOR_NAME	Name of the administrator involved. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.
OMS_MACHINE	Management Server currently working on this operation. Value will be NULL if not yet allocated The OMS_MACHINE column links with OMS_MACHINE from the SMP_VIEW_OMS_MACHINES view.

Job Definition Views

SMP_VIEW_JOBS

A view of the Job Subsystem which lists all jobs present in the repository:

- Which jobs are defined in the system.
- Jobs can either be ad-hoc or defined in the library

The boolean values are represented as a numerical value so that these values can change into a string appropriate for the Globalization Support settings which are currently active.

These values can then be decoded into proper localized values when the report is run

Table D-18 SMP_VIEW_JOBS

Columns	Descriptions
JOB_ID	<p>ID of this job.</p> <p>The JOB_ID is an unique identifier for this job.</p> <p>Since a job can be submitted at different times, and because the same job can be created more than once, the JOB_NAME is not enough to uniquely identify a job in the repository.</p> <p>Therefore, all queries used to obtain job information must use the JOB_ID identifier.</p>
JOB_NAME	<p>Name of the group. JOB_NAME contains the name of the job as entered by the administrator. This name is not unique and can appear several times for different jobs.</p>
ADMINISTRATOR_NAME	<p>Name of the administrator owning this job. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.</p>
DESCRIPTION	<p>Description of the job</p>
TARGET_TYPE	<p>Type of the targeted service. A job is always designed to run against a specific set of target targets.</p> <p>This job type is represented by the TARGET_TYPE field which basically determines which different job tasks are possible for this job.</p>
TARGET-NLS_TYPE	<p>Type string used for globalization support of the reports.</p>
FIXIT_JOB	<p>Shows whether it is a fixit job</p> <p>Boolean field: 0 (NO) or 1 (YES)</p>
LIBRARY_JOB	<p>Shows whether it is a job defined only in the library</p> <p>Boolean field: 0 (NO) or 1 (YES)</p>
INTERVAL_JOB	<p>Shows whether it is a job to be executed in an interval</p> <p>Boolean field: 0 (NO) or 1 (YES)</p>

Table D–18 SMP_VIEW_JOBS(Cont.)

Columns	Descriptions
MODIFIED_BY	Name of the administrator who last modified this job. The MODIFIED_BY column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.
MODIFIED_DATE	Time and date this job was last modified. The value of the MODIFIED_DATE field is specified in the time zone of the administrator who has modified it. When the job is created, the first set of modification information is always the owner of the job and the creation date.

SMP_VIEW_JOB_EXECUTIONS

A view of the Job Subsystem which shows the status of each execution of a job. Only the end result of each execution is given in each view.

Table D–19 SMP_VIEW_JOB_EXECUTIONS

Columns	Descriptions
JOB_ID	ID of this job. The JOB_ID is an unique identifier for this job. Since a job can be submitted at different times, and because the same job can be created more than once, the job_name is not enough to uniquely identify an job in the repository. Therefore, all queries used to obtain job information must use the JOB_ID identifier. The JOB_ID column links with JOB_ID of the SMP_VIEW_JOBS view.
JOB_NAME	Name of the job. JOB_NAME contains the name of the job as entered by the administrator. This name is not unique and can appear several times for different jobs.
NODE_NAME	Name of the node on which this job is running. The NODE_NAME column links with NODE_NAME in the SMP_VIEW_NODES view.

Table D-19 SMP_VIEW_JOB_EXECUTIONS(Cont.)

Columns	Descriptions
TARGET_NAME	<p>Target this job is scheduled/submitted for. The TARGET_NAME combination uniquely identifies the target this job was meant for. Several targets can be present for any particular job.</p> <p>The TARGET_NAME column links with TARGET_NAME of the SMP_VIEW_NODE_SERVICES view.</p>
TARGET_TYPE	Type of the targeted service
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
ADMINISTRATOR_NAME	Name of the administrator who owns this job. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.
EXEC_NUM	<p>The execution number. The EXEC_NUM field represents the number of times the job was actually scheduled against a particular target. For each time the job gets to status SCHEDULED, this number is increased by 1.</p> <p>For the SUBMITTED state, the EXEC_NUM value will always be 0, as the Intelligent Agent has not responded for this job yet.</p>
TIMESTAMP	Time the action was recorded. The timestamp is the time of the job occurrence on the Intelligent Agent side.

Table D-19 SMP_VIEW_JOB_EXECUTIONS(Cont.)

Columns	Descriptions
STATUS	<p data-bbox="629 296 929 326">Status of the job notification</p> <ul style="list-style-type: none"> <li data-bbox="629 340 1258 418">■ 1 - Submitted - The Management Server has submitted the request to the Intelligent Agent and is waiting for a confirmation <li data-bbox="629 432 1258 510">■ 2 - Scheduled - The Intelligent Agent has responded to the submit request and is now waiting till the actual execution time for the job to get executed <li data-bbox="629 524 1079 553">■ 4 - Running - The job has been started <li data-bbox="629 567 1243 619">■ 9 - Completed - The job terminated successfully (with an exit code of zero 0) <li data-bbox="629 633 1250 663">■ 11 - Failed - The job failed. The exit code was non-zero <li data-bbox="629 677 1250 755">■ 13 - Pending delete - An administrator requested a delete of a scheduled job. Waiting on Intelligent Agent confirmation <li data-bbox="629 769 1243 821">■ 14 - Deleted - The Intelligent Agent has confirmed the delete of the job <p data-bbox="629 835 1265 913">The status 0 can never appear in this view as library jobs have no active status and do not interact with the Intelligent Agent.</p> <p data-bbox="629 927 1258 979">Also, status 15 cannot appear in this view as this is also not an active status of a job.</p>

SMP_VIEW_JOB_HISTORY

A Job Subsystem view which shows the history of all the job executions. It is only a snapshot of the available information since the job history can be cleanup up and jobs can get manually deleted. It presents information about when and how the job was executed and an overview of the job at this time.

Table D-20 SMP_VIEW_JOB_HISTORY

Columns	Descriptions
JOB_ID	<p>ID of this job. The JOB_ID is a unique identifier for this job.</p> <p>Since a job can be submitted at different times, and because the same job can be created more than once, the job_name is not enough to uniquely identify an job in the repository.</p> <p>Therefore, all queries used to obtain job information must use the JOB_ID identifier.</p> <p>The JOB_ID column links with JOB_ID of the SMP_VIEW_JOBS view.</p>
JOB_NAME	<p>Name of the job. JOB_NAME contains the name of the job as entered by the administrator.</p> <p>This name is not unique and can appear several times for different jobs.</p>
TARGET_NAME	<p>Target this job is scheduled/submitted for. The TARGET_NAME combination uniquely identifies the target this job was meant for. Several targets can be present for any particular job.</p> <p>The TARGET_NAME column links with TARGET_NAME of the SMP_VIEW_NODE_SERVICES view.</p>
TARGET_TYPE	Type of the targeted service
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
EXEC_NUM	<p>The execution number. The EXEC_NUM field represents the number of times the job was actually scheduled against a particular target. For each time the job gets to status SCHEDULED, this number is increased by 1.</p> <p>For the SUBMITTED state, the EXEC_NUM value will always be 0 as the Intelligent Agent has not responded for this job yet.</p>
TIMESTAMP	Time the action was recorded. The timestamp is the time of the job occurrence on the Intelligent Agent side.

Table D–20 SMP_VIEW_JOB_HISTORY(Cont.)

Columns	Descriptions
STATUS	<p>Status of the job notification</p> <ul style="list-style-type: none"> ■ 1 - Submitted - The Management Server has submitted the request to the Intelligent Agent, and is waiting for a confirmation ■ 2 - Scheduled - The Intelligent Agent has responded to the submit request, and is now waiting till the actual execution time for the job to get executed ■ 4 - Running - The job has been started ■ 9 - Completed - The job terminated successfully (with an exit code of zero 0) ■ 11 - Failed - The job failed. The exit code was non-zero ■ 13 - Pending delete - An administrator requested a delete of a scheduled job. Waiting on Intelligent Agent confirmation ■ 14 - Deleted - The Intelligent Agent has confirmed the delete of the job

SMP_VIEW_JOB_STATUS

A job subsystem view which shows the current status of all active jobs. One line will appear in this view for each target this job is submitted against. For library jobs, the NODE_NAME field is NULL, and the STATUS and EXEC_NUM fields are forced to 0 since no actual executions have taken place. Only active jobs are represented in this view.

Table D–21 SMP_VIEW_JOB_STATUS

Columns	Descriptions
JOB_ID	<p>ID of this job.</p> <p>The JOB_ID is a unique identifier for this job. Since a job can be submitted at different times, and because the same job can be created more than once, the job_name is not enough to uniquely identify a job in the repository. Therefore, all queries used to obtain job information must use the JOB_ID identifier.</p> <p>The JOB_ID column links with JOB_ID in the SMP_VIEW_JOBS view.</p>

Table D-21 SMP_VIEW_JOB_STATUS(Cont.)

Columns	Descriptions
JOB_NAME	Name of the job. JOB_NAME contains the name of the job as entered by the administrator. This name is not unique and can appear several times for different jobs.
NODE_NAME	Name of the node on which this job is running The name of the node will be NULL for library jobs The NODE_NAME column links with NODE_NAME in the SMP_VIEW_NODES view.
TARGET_NAME	The target this job is scheduled/submitted for. The TARGET_NAME, TARGET_TYPE combination uniquely identifies the target this job was meant for. Several targets can be present for any particular job. The TARGET_NAME column links with TARGET_NAME in the SMP_VIEW_TARGETS view.
TARGET_TYPE	Type of the targeted service
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
ADMINISTRATOR_NAME	Name of the administrator who owns this job. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.
EXEC_NUM	The execution number. The EXEC_NUM field represents the number of times the job was actually scheduled against a particular target. For each time the job gets to status SCHEDULED, this number is increased by 1. For the SUBMITTED state, as well as library jobs, the EXEC_NUM value will always be 0, as the Intelligent Agent has not responded for this job yet.
TIMESTAMP	Time the job notification was recorded

Table D-21 SMP_VIEW_JOB_STATUS(Cont.)

Columns	Descriptions
STATUS	Status of the job: <ul style="list-style-type: none">■ 1 - Submitted - The Management Server has submitted the request to the Intelligent Agent and is waiting for a confirmation■ 2 - Scheduled - The Intelligent Agent has responded to the submit request, and is now waiting till the actual execution time for the job to get executed■ 4 - Running - The job has been started■ 9 - Completed - The job terminated successfully (with an exit code of zero 0)■ 11 - Failed - The job failed. The exit code was non-zero■ 13 - Pending delete - An administrator requested a delete of a scheduled job. Waiting on Intelligent Agent confirmation■ 14 - Deleted - The Intelligent Agent has confirmed the delete of the job■ 15 - Expired - The finish time for an internal job has passed

SMP_VIEW_AGENT_JOBS

View of job subsystem with overview of all active jobs with the Intelligent Agent specific information.

The boolean values are represented as a numerical value so that these values can change into a string appropriate for the Globalization Support settings which are currently active.

These values can then be decoded into proper localized values when the report is run

Table D-22 SMP_VIEW_AGENT_JOBS

Columns	Descriptions
JOB_ID	<p>ID of this job.</p> <p>The JOB_ID is an unique identifier for this job.</p> <p>Since a job can be submitted at different times, and because the same job can be created more than once, the job_name is not enough to uniquely identify a job in the repository.</p> <p>Therefore, all queries used to obtain job information must use the JOB_ID identifier.</p>
AGENT_ID	<p>ID attributed by the Intelligent Agent for this job.</p> <p>The AGENT_ID is the unique Intelligent Agent ID for this job. This ID is unique for this Intelligent Agent, although others Intelligent Agents can attribute the same ID to a job or event.</p>
JOB_NAME	<p>Name of the group. JOB_NAME contains the name of the job as entered by the administrator.</p> <p>This name is not unique, and can appear several times for different jobs.</p>
ADMINISTRATOR_NAME	<p>Name of the administrator owning this job. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.</p>
DESCRIPTION	<p>Description of the job</p>
NODE_NAME	<p>Node on which this job is running. The NODE_NAME column links with NODE_NAME in the SMP_VIEW_NODES view.</p>
TARGET_NAME	<p>Name of the targeted service. The TARGET_NAME column links with TARGET_NAME in the SMP_VIEW_TARGETS view.</p>
TARGET_TYPE	<p>Type of the targeted service. A job is always designed to run against a specific set of target targets.</p> <p>This job type is represented by the TARGET_TYPE field, which basically determines which different job tasks are possible for this job.</p>
TARGET-NLS_TYPE	<p>Type string used for globalization support of the reports.</p>

Table D-22 SMP_VIEW_AGENT_JOBS(Cont.)

Columns	Descriptions
FIXIT_JOB	Shows whether it is a fixit job Boolean field: 0 (NO) or 1 (YES)
INTERVAL_JOB	Shows whether the job is to be executed in an interval Boolean field: 0 (NO) or 1 (YES)
EXEC_NUM	The execution number. The EXEC_NUM field represents the number of times the job was actually scheduled against a particular target. For each time the job gets to status SCHEDULED, this number is increased by 1. For the SUBMITTED state, the EXEC_NUM value will always be 0, as the Intelligent Agent has not responded for this job yet.
MODIFIED_BY	Name of the administrator who last modified this job. The MODIFIED_BY column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.
MODIFIED_DATE	Time and date this job was last modified. The value of the MODIFIED_DATE field is specified in the timezone of the administrator who has modified it. When the jobs is created, the first set of modification information is always the owner of the job and the creation date.
STATUS	Status of the job condition 1 - Submitted 2 - Scheduled 4 - Running 9 - Completed successfully 11 - Job failed 13 - Pending delete 14 - Job deleted 15 - Job expired
STARTED	Time and date the job started
FINISHED	Time and date the job finished

Table D–22 SMP_VIEW_AGENT_JOBS(Cont.)

Columns	Descriptions
NEXT_EXEC	Time and date for the next execution. Only valuable for interval jobs
AGENT_STATUS	Shows whether an Intelligent Agent is currently running Boolean field: 0 (DOWN) or 1 (UP)

Event Definition Views

SMP_VIEW_EVENTS

A view of the Event Subsystem which lists all events defined in the repository, only 1 line for each event defined in the repository.

The boolean values are represented as a numerical value so that these values can change into a string appropriate for the Globalization Support settings which are currently active.

These values can then be decoded into proper localized values when the report is run

Table D–23 SMP_VIEW_EVENTS

Columns	Descriptions
EVENT_ID	Uniquely distinguishes events from each other. The EVENT_ID is an unique identifier for this job. Since a event can be registered and de-registered at different times, and because the same event can be registered against more than one target, the EVENT_NAME is not enough to uniquely identify an event in the repository. Therefore, all queries used to obtain event information must use the EVENT_ID identifier. The EVENT_ID column links with EVENT_ID from the SMP_VIEW_EVENTS view.
EVENT_NAME	Name of this event

Table D-23 SMP_VIEW_EVENTS(Cont.)

Columns	Descriptions
ADMINISTRATOR_NAME	Name of the administrator who created this event. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.
DESCRIPTION	Description of the event
NODE_NAME	Node on which this event is running. The NODE_NAME column links with NODE_NAME from the SMP_VIEW_NODES view.
TARGET_TYPE	Type of target this event is meant for
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
LIBRARY_EVENT	Shows whether an event is saved in the event library Boolean field: 0 (NO) or 1 (YES)
FIXIT_JOB_ID	ID of the fixit job to execute when this event triggers Value is ZERO (0) if no fixit associated with this event.
UNSOLICITED	Shows whether it is a 3rd party event Boolean field: 0 (NO) or 1 (YES)
SNMP_TRAP	Shows whether an SNMP trap is sent when event triggers Boolean field: 0 (NO) or 1 (YES)

SMP_VIEW_EVENT_HISTORY

View of event subsystem with historical overview of all event occurrences currently in the repository.

The boolean values are represented as a numerical value so that these values can change into a string appropriate for the Globalization Support settings which are currently active.

These values can then be decoded into proper localized values when the report is run

Table D-24 SMP_VIEW_EVENT_HISTORY

Columns	Descriptions
EVENT_ID	<p>Needed to uniquely distinguish events from each other. The EVENT_ID is an unique identifier for this job.</p> <p>Since a event can be registered and de-registered at different times, and because the same event can be registered against more than one target, the EVENT_NAME is not enough to uniquely identify an event in the repository.</p> <p>Therefore, all queries used to obtain event information must use the EVENT_ID identifier.</p> <p>The EVENT_ID column links with EVENT_ID from the SMP_VIEW_EVENTS view.</p>
EVENT_NAME	Name of this event.
OCCURRENCE_ID	Sequence number of event occurrences. These numbers are not necessarily consecutive
NODE_NAME	Name of the node on which this event is running. The NODE_NAME column links with NODE_NAME from the SMP_VIEW_NODES view.
TARGET_NAME	<p>Name of the target that caused the event to trigger. An event is always registered against a specific set of target services.</p> <p>This event type is represented by the TARGET_TYPE field, which basically determines which different event tests are possible for this event.</p> <p>The TARGET_NAME column links with TARGET_NAME from the SMP_VIEW_NODE_SERVICES view.</p>
TARGET_TYPE	Type of target this event is meant for
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
OWNER	<p>Administrator who is currently assigned to check this occurrence.</p> <p>The OWNER column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.</p>
TIMESTAMP	Date/Time this occurrence happened
MESSAGE	Specific occurrence text
AGENT_STATUS	<p>Current status of the Intelligent Agent</p> <p>Boolean field: 0 (DOWN) or 1 (UP)</p>

Table D–24 SMP_VIEW_EVENT_HISTORY(Cont.)

Columns	Descriptions
SEVERITY	Current severity level of this occurrence
ACTIVE	Shows whether this event occurrence is still outstanding in the console Boolean field: 0 (NO) or 1 (YES)

SMP_VIEW_EVENT_TEST_HISTORY

View of event subsystem with historic overview of the notifications of all tests per events present in the repository.

The boolean values are represented as a numerical value so that these values can change into a string appropriate for the Globalization Support settings which are currently active.

These values can then be decoded into proper localized values when the report is run

Table D–25 SMP_VIEW_EVENT_TEST_HISTORY

Columns	Descriptions
EVENT_ID	Needed to uniquely distinguish events from each other. The EVENT_ID is an unique identifier for this job. Since a event can be registered and de-registered at different times, and because the same event can be registered against more than one target, the EVENT_NAME is not enough to uniquely identify an event in the repository. Therefore, all queries used to obtain event information must use the EVENT_ID identifier. The EVENT_ID column links with EVENT_ID from the SMP_VIEW_EVENTS view.
EVENT_NAME	Name of this event
TEST_ID	Needed to uniquely distinguish this test from other tests within the event
EVENT_TEST	Name of the test within the event
OCCURRENCE_ID	Sequence number of event occurrences. These numbers are not necessarily consecutive

Table D-25 SMP_VIEW_EVENT_TEST_HISTORY(Cont.)

Columns	Descriptions
NODE_NAME	Name of the node on which this event is running. The NODE_NAME column links with NODE_NAME from the SMP_VIEW_NODES view.
TARGET_NAME	Name of the target that caused the event to trigger. The TARGET_NAME column links with TARGET_NAME from the SMP_VIEW_NODE_SERVICES view.
TARGET_TYPE	Type of target this event is meant for. An event is always registered against a specific set of target targets. This event type is represented by the TARGET_TYPE field, which basically determines which different event tests are possible for this event.
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
OWNER	Administrator who is currently assigned to check this occurrence. The OWNER column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.
TIMESTAMP	Date/Time this occurrence happened
MESSAGE	Specific occurrence text
EVENT_SEVERITY	Overall severity level of the event
TEST_SEVERITY	Severity level of the test within the event
ACTIVE	Shows whether this event occurrence is still outstanding in the console. Boolean field: 0 (NO) or 1 (YES)

SMP_VIEW_EVENT_STATUS

View of Event Subsystem showing current status of all events present in the repository.

The boolean values are represented as a numerical value so that these values can change into a string appropriate for the Globalization Support settings which are currently active.

These values can then be decoded into proper localized values when the report is run

Table D-26 SMP_VIEW_EVENT_STATUS

Columns	Descriptions
EVENT_ID	<p>Needed to uniquely distinguish events from each other. The EVENT_ID is an unique identifier for this job.</p> <p>Since a event can be registered and de-registered at different times, and because the same event can be registered against more than one target, the EVENT_NAME is not enough to uniquely identify an event in the repository.</p> <p>Therefore, all queries used to obtain event information must use the EVENT_ID identifier.</p> <p>The EVENT_ID column links with EVENT_ID from the SMP_VIEW_EVENTS view.</p>
EVENT_NAME	Name of this event
NODE_NAME	Name of the node on which this event is running. The NODE_NAME column links with NODE_NAME from the SMP_VIEW_NODES view.
TARGET_NAME	Name of the target that caused the event to trigger. The TARGET_NAME column links with TARGET_NAME from the SMP_VIEW_NODE_SERVICES view.
TARGET_TYPE	<p>Type of target this event is meant for</p> <p>An event is always registered against a specific set of target targets.</p> <p>This event type is represented by the TARGET_TYPE field, which basically determines which different event tests are possible for this event.</p>
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
ADMINISTRATOR_NAME	Name of the administrator who owns this event. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.
TIMESTAMP	Date/Time this occurrence happened

Table D-26 SMP_VIEW_EVENT_STATUS(Cont.)

Columns	Descriptions
EVENT_STATUS	Current status of this event 204: Pending registration 205: Registered 206: Registration failed 207: Pending deregistration 208: Deregistered
AGENT_STATUS	Current status of the Intelligent Agent Boolean field: 0 (DOWN) or 1 (UP) If the Intelligent Agent is DOWN or BAD, this boolean field will be marked as DOWN
SEVERITY	Current severity level of this event 15 - clear 20 - warning 25 - alert

SMP_VIEW_EVENT_TESTS

View of event subsystem with list of all tests defined per event.

Table D-27 SMP_VIEW_EVENT_TESTS

Columns	Descriptions
EVENT_ID	Needed to uniquely distinguish events from each other. The EVENT_ID is an unique identifier for this job. Since a event can be registered and de-registered at different times, and because the same event can be registered against more than one target, the EVENT_NAME is not enough to uniquely identify an event in the repository. Therefore, all queries used to obtain event information must use the EVENT_ID identifier. The EVENT_ID column links with EVENT_ID from the SMP_VIEW_EVENTS view.
EVENT_NAME	Name of this event

Table D–27 SMP_VIEW_EVENT_TESTS(Cont.)

Columns	Descriptions
EVENT_TEST	Internal name of the event test

SMP_VIEW_AGENT_EVENTS

View of Event Subsystem showing Intelligent Agent specific information about all registered events.

The boolean values are represented as a numerical value so that these values can change into a string appropriate for the Globalization Support settings which are currently active.

These values can then be decoded into proper localized values when the report is run

Table D–28 SMP_VIEW_AGENT_EVENTS

Columns	Descriptions
EVENT_ID	<p>Needed to uniquely distinguish events from each other. The EVENT_ID is an unique identifier for this job.</p> <p>Since a event can be registered and de-registered at different times, and because the same event can be registered against more than one target, the EVENT_NAME is not enough to uniquely identify an event in the repository.</p> <p>Therefore, all queries used to obtain event information must use the EVENT_ID identifier.</p> <p>The EVENT_ID column links with EVENT_ID from the SMP_VIEW_EVENTS view.</p>
AGENT_ID	ID attributed by the Intelligent Agent for this event.
EVENT_NAME	Name of this event
EVENT_TEST	Specific test the Intelligent Agent has to evaluate
ADMINISTRATOR_NAME	<p>Name of the administrator who created this event. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.</p>
DESCRIPTION	Description of the event

Table D-28 SMP_VIEW_AGENT_EVENTS(Cont.)

Columns	Descriptions
NODE_NAME	Node on which this event is running. The NODE_NAME column links with NODE_NAME from the SMP_VIEW_NODES view.
TARGET_NAME	Name of the target this event is meant for. The TARGET_NAME column links with TARGET_NAME in the SMP_VIEW_TARGETS view.
TARGET_TYPE	Type of target this event is meant for
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
FIXIT_JOB_ID	ID of the fixit job to execute when this event triggers Value is ZERO (0) if no fixit associated with this event.
UNSOLICITED	Shows whether this is a 3rd party event Boolean field: 0 (NO) or 1 (YES)
SNMP_TRAP	Shows whether to send an SNMP trap when event triggers Boolean field: 0 (NO) or 1 (YES)
SEVERITY	Current severity level of this event 15 - clear 20 - warning 25 - alert
AGENT_STATUS	Shows whether the Intelligent Agent is currently running Boolean field: 0 (DOWN) or 1 (UP)

Capacity Planner and Data Gatherer Collections

SMP_VIEW_COLLECTIONS

Capacity Planner view showing list of defined Capacity Planner collections in the repository.

The boolean values are represented as a numerical value so that these values can change into a string appropriate for the Globalization Support settings which are currently active.

These values can then be decoded into proper localized values when the report is run

Table D-29 SMP_VIEW_COLLECTIONS

Columns	Descriptions
DG_NAME	Name of node on which the Data Gatherer is running which is collecting the information.
NODE_NAME	Name of node on which this collection runs. The NODE_NAME column links with NODE_NAME from the SMP_VIEW_NODES view.
TARGET_NAME	Name of the target for which this collection is defined. The TARGET_NAME column links with TARGET_NAME from the SMP_VIEW_NODE_SERVICES view.
TARGET_TYPE	Type identifier of this target.
TARGET-NLS_TYPE	Type string used for globalization support of the reports.
ADMINISTRATOR_NAME	Name of the administrator owning this collection. The ADMINISTRATOR_NAME column links with ADMINISTRATOR_NAME from the SMP_VIEW_ADMINISTRATORS view.
ACTIVE	Shows whether the collection is still active Boolean field: 0 (NO) or 1 (YES)
LOCAL	Shows whether the collection is stored locally in the repository Boolean field: 0 (NO) or 1 (YES)

Index

Symbols

, 1-26

A

acknowledging events

Event menu, 6-40

Event window, 6-38

Active Jobs page, 5-7

administration task

Navigator objects, 3-9

administrator

adding, editing, and deleting, 1-9

managing, 1-8

preferences, 1-13

reassigning objects when deleting, 1-13

super, 1-8

Administrator Notification

email and paging, 6-56

Advanced Mode

View menu, 1-5

Analyze Wizard, 10-32

Apache Web Server, 1-32

Apache Web Server, discovering, 1-32

Archive Log Operations, Storage Management, 10-24

audience for this guide, xxi

B

Background images

adding to groups, 4-6

Backup and Recovery Management Wizard, 10-32

backup and recovery, managing, 11-1

blackouts, paging/email, 1-34

Broadcast Message task, 5-33

C

clone (of a database object), 10-18

communication daemon

populating Navigator, 3-2

Components, 1-2

Navigator, 1-2

composite job, 5-3

configuration operations, 10-8

Console

File menu, 1-4

Help menu, 1-7

menus, 1-3

using menus, 1-3

console

list of tasks that can be performed from, xx

Console in standalone mode, 2-1

adding databases to the tree, 2-4

connecting to a database, 2-6

editing local preferred credentials, 2-9

starting, 2-2

Console panes

Groups, 4-1

Navigator, 3-1

context menu, 3-5

context-sensitive menus

Navigator, 1-4

using in the Navigator, 1-4

Controlfile Operations, Storage Management, 10-23

- copying tree objects
 - Navigator, 3-9
- create
 - event set, 6-40
 - job, 5-11
 - or modify a job, 5-12
 - or modify an event set, 6-31
- Create Like, 10-18
- Create Table Wizard, 10-32
- Cube Wizard, 10-32
- CWMLite, 10-26

D

- Data Management Wizards, 10-32
- database
 - connection made in tree, 3-3
 - recovery, 11-29
- database administration, 10-1
- database schema objects, 10-14
- database security, managing, 9-28
- Databases folder
 - Navigator tree, 3-3
- datafile operations, Storage Management, 10-23
- DBA management functionality
 - comprehensive overview page, 10-3
 - database reports, 10-4
 - database version awareness, 10-3
 - DB Search Capabilities, 10-5
 - general information about databases, 10-3
 - logging of database changes, 10-4
 - multi-column lists, 10-3
 - property sheets, 10-3
 - Show Dependencies, 10-4
 - Show SQL, 10-4
 - showing Object DDL, 10-4
 - tree views, 10-3
- Detail, 1-2
- detail pane, 1-2
- Dimension Creation Wizard, 10-32
- directory base, 9-10
- Directory servers, 9-3
- discovering services
 - Navigator, 3-6
 - SQL*Net network, 3-2

- discovery
 - problems, 3-8
- discovery status, 3-8
- domain administrators, 9-39
- domain schema mapping, 9-40
- drag and drop
 - Navigator, 3-9
 - Navigator objects, 3-2

E

- EMS
 - Event Management system, 6-1
- enterprise domain, membership, 9-36
- enterprise roles, administering, 9-41
- Enterprise security, installing, 9-5
- event
 - set, 6-2
 - unsolicited, 6-11
- Event Filters, 7-7
- Event Handler, 7-2
- Event Handler templates, 7-10
- event handler, configuration commands, 7-17
- event handler, configuration file, 7-18
- Event Handler, configuration parameters, 7-7
- Event Handler, customizing, 7-5
- event handler, disabling, 7-17
- event handler, enabling, 7-17
- event handler, multi-OMS, 7-23
- Event Handler, setup, 7-4
- event handler, viewing configuration, 7-17
- Event History page
 - Event window, 6-38
- Event Logger templates, 7-10
- Event Management, 6-1
 - categories of events, 6-8
 - creating an event set, 6-31
 - fault management, 6-9
 - introduction, 6-1
 - menu, 6-40
 - modifying an event set, 6-31
 - Oracle events, 6-56
 - parameter settings, 6-49
 - performance management, 6-10
 - process, 6-2

- resource management, 6-10
- space management, 6-9
- types of events, 6-8
- viewing an event set, 6-41
- window, 6-36

Event Set Library page

- Event window, 6-41
- predefined event sets, 6-41

Event Set Management property sheet

- creating an event, 6-40
- modifying a new event, 6-40

example

- Job Scheduling, 5-26

expanding

- map objects, 4-9

F

- Fault Management events, 6-8
- Firewall, VPN Connections, B-7
- Firewalls, B-2
- Firewalls, Capacity Planner, B-10
- Firewalls, Performance Manager, B-10
- Firewalls, Port Usage-Console/OMS, B-3
- Firewalls, Port Usage-OMS/Managed Target, B-4
- fixit job, 5-3, 6-1
 - with events, 6-1

G

General page

- job, 5-13

Group

- pane, 4-2

group, 4-3

- add or delete, 4-10
- adding background images, 4-6
- modify, 4-11
- modifying properties, 4-11

Groups folder

- Navigator tree, 3-3

H

headings

H1 Head1, C-2

I

- In-Doubt Transactions, 10-12
- Instance Management, 10-7

J

job

- alter, 5-11
- composite, 5-3
- create, 5-11
- manage, 5-11
- modify, 5-11
- view, 5-11

Job History page, 5-9

Job Library, 5-12

job output

- Progress Page, 5-23

Job Scheduling, 5-1

- credential preference, 5-4, 5-18
- destination requirements, 5-15
- example, 5-26
- General page, 5-13
- job destination, 5-13
- job name, 5-13
- job tasks, 5-15
- job type, 5-13
- managing jobs, 5-5
- menu, 5-11
- scheduling jobs, 5-5
- scheduling process, 5-2
- task parameters, 5-18
- window, 5-6

Job Tasks

- wizards, 5-29

jobs, cancelling, 5-6

L

launching a database tool

- from a map, 4-10
- from the Console, 3-4
- from the Navigator, 1-2

- Navigator, 3-4
 - with the right-mouse button, 3-4
- levels of permission, 1-24
- Listeners folder
 - Navigator tree, 3-3
- Lists, Saving, 1-28
- Locks list, 10-11
- LRS
 - See Log roll-forward server (LRS), A-1, D-1

M

- Maintenance Wizard, 11-5
- management region, 1-29
- manipulation of the objects
 - in the Navigator, 3-9
- Map, 4-1
 - bitmap background, 4-7
 - create, 4-8
 - creating with dragging and dropping, 4-8
 - customized views, 4-1
 - expanding objects, 4-9
 - status of objects, 4-8
 - user-defined views, 4-3
- menus
 - Console, 1-3
 - File menu, 1-4
 - Help menu, 1-7
 - Job, 5-11
- modify
 - event set, 6-40
- multiple tasks
 - submitting a job, 5-13

N

- NAT, B-5
- Navigator, 3-1
 - Enterprise Manager, 1-2
 - manipulating objects, 3-9
 - menu, 3-4
 - objects, 1-2
 - window, 3-2
- network
 - container, 3-3

- Network Address Translation, B-5
- Node Properties, 3-9
- Nodes folder
 - Navigator tree, 3-3

O

- object
 - creating, 10-18
 - editing, 10-18
- oemapp console oem.loginmode=standalone
 - command, 2-3
- OEMUTIL, C-1
- OEMUTIL Commands
 - changeCredentials, C-11
 - deregisterEvent, C-9
 - omsCredentials, C-3
 - registerEventFromLibrary, C-8
 - submitJob, C-4
 - submitJobFromLibrary, C-6
- OEMUTIL, multiple commands, C-2
- OEMUTIL, single command, C-2
- OEMUTIL, starting, C-2
- OLAP Management, 10-26
- Oracle Contexts, administering, 9-20
- Oracle DB UpDown
 - event profile, 6-42
- Oracle events, 6-56
- Oracle Host UpDown
 - event profile, 6-42
- Oracle Listener UpDown
 - event profile, 6-42
- Oracle Wallet, 9-14
- OraTcl
 - job scripts, 5-2
- Output Dialog box
 - Job Scheduling, 5-23
- Outstanding Events page
 - Event system, 6-37

P

- paging status codes for numeric pages, 1-20
- Paging/Email Blackout, 1-34
 - Total, 1-34

- parameters
 - events, 6-49
 - jobs, 5-17
- Parameters Page
 - job, 5-17
- password management, 10-20
- Performance Management events, 6-8
- Ping Agent, 3-5
- ping agent menu item, 3-8
- Port Usage
 - firewall, B-2
- Pre-defined Profiles
 - Event system, 6-42
- Preferred credentials
 - , 1-26
- preferred credentials
 - connecting in a map, 4-10
- preferred credentials, setting local, 2-9
- problems
 - Navigator discovery, 3-8
- profile operations, Security Management, 10-20
- Progress page
 - job, 5-23
- purpose of this guide, xx

R

- recovery
 - database, 11-29
- redo log group operations, Storage Management, 10-24
- Refreshing Discovery
 - Refresh Topology Menu Choice, 3-8
- registered destination
 - map objects, 4-8
- registering events, 6-31
- Registrations page
 - Event window, 6-39
- related publications, xxvi
- removing
 - job with Remove menu item, 5-11
 - problems with a job, 5-11
 - problems with an event, 6-41
 - registered events, 6-41
- Report Definition, 8-2

- Report element, 8-4
- report elements, 8-15
- report generation, applications, 8-10
- Reporting, 8-2
- Reporting website, 8-6
- reports, creating, 8-9
- reports, editing, 8-10
- reports, user-defined, 8-11
- Resource Consumer Groups, 10-12
- Resource Management events, 6-8
- Resource Plan Schedule, 10-13
- Resource Plan Wizard, 10-32
- Resource Plans, 10-12
- right mouse button
 - used in the Navigator, 1-4
- right-mouse button
 - Related Tools menu, 3-4
- role operations, Security Management, 10-20
- rollback segment operations, Storage Management, 10-23
- rollback segments, 10-21
- Run DBA Script task, 5-30
- Run OS Command task, 5-33
- Run SQL*Plus task, 5-30
- Run Tcl Script task, 5-34

S

- sample images, 4-6
- saving
 - job history, 5-10, 5-11
- schedule
 - and manage jobs, 5-5
 - execution of a job, 5-18
 - job, 5-5
- Schema Management, 10-14
- Security Management, 10-19
- Service Discovery
 - Navigator menu option, 3-5
- Sessions Folder, 10-11
- Sessions List, 10-11
- setting a bitmap background
 - Map menu, 4-7
- short-cut menu, 1-4
- Shutdown Database task, 5-31

- Shutdown Listener task, 5-36
- signal flag
 - on map objects, 4-8
- Space Management event, 6-8
- specifying service names, 1-27
- SPFILE, 10-10
- SQL error
 - job task, 5-30
- SQL*Net network services, 3-2
- SQL*Plus Worksheet, 10-29
- Standalone Console procedures
 - changing from connected to Management Server mode to standalone mode, 2-8
 - changing from standalone mode to connected to Management Server mode, 2-11
 - connecting to the database as a different user, 2-7
 - removing a database from tree, 2-7
 - viewing which role you are connected as, 2-7
- Startup Database task, 5-31
- Startup Listener task, 5-36
- status
 - of a map object, 4-8
 - of a registered event set, 6-4
- Stored Configurations, 10-10
- Summary Advisor Wizard, 10-32
- summary of the pending jobs
 - Active Jobs page, 5-7
- Super Administrator, 1-8

T

- tablespace operations, Storage Management, 10-23
- tablespaces, administering, 10-21
- target access, 1-11
- Targets, report, 8-5
- task parameters
 - Job system, 5-17
- Tcl
 - job scripts, 5-2
 - script examples, 5-34
- third-party events
 - creating, 6-10
 - with job task, 5-35

- tnsnames.ora, 3-8
- Tools menu, 1-6

U

- unsolicited event, 6-11
- unsolicited events, 6-8
- user name, Enterprise Security, 9-9
- user operations, Security Management, 10-19
- user preferences
 - jobs, 5-4, 5-18
 - when running jobs, 5-4, 5-18
- user search base, specifying, 9-22
- user security, 9-8
- user-defined map views
 - creating, 4-3

V

- View Wizard, 10-32
- Virtual Private Networks (VPNs), B-6
- VPNs and Standalone Consoles, B-9

W

- web server
 - discovering, 1-33
- Wizard Job Tasks, 5-29
- worksheet, Oracle SQL*Plus Worksheet, 10-29
- Workspace Management, 10-27
- Workspace Manager, 10-27