

Oracle® Enterprise Manager

Administrator's Guide

Release 2.1

December, 1999

Part No. A75683-01

ORACLE

Oracle Enterprise Manager Administrator's Guide, Release 2.1

Part No. A75683-01

Copyright © 1996, 2000, Oracle Corporation. All rights reserved.

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.

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 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.

Oracle is a registered trademark, and Oracle Trace, SQL*Plus, and SQL*Net, Net8, Oracle Expert, Oracle7, Oracle8, and PL/SQL are trademarks or registered trademarks of Oracle Corporation. All other company or product names mentioned are used for identification purposes only and may be trademarks of their respective owners.

Contents

Send Us Your Comments	xi
Preface.....	xiii
Part I Oracle Enterprise Manager Framework	
1 Console	
Console Framework.....	1-2
Components	1-2
Navigator	1-2
Group	1-2
Jobs	1-3
Events	1-3
Console Menus.....	1-4
Context-Sensitive Menus.....	1-4
File Menu	1-5
View Menu	1-5
System Menu.....	1-5
Tools Menu.....	1-6
Database Applications	1-6
Service Management	1-6
Help Menu.....	1-6
Managing Enterprise Manager Administrators.....	1-7
Managing Administrators.....	1-7

Creating or Editing an Administrator Account	1-8
Reassigning Object Ownership.....	1-9
Administrator Preferences.....	1-10
General	1-11
Notification	1-12
Paging Information	1-12
Email Parameters	1-13
Schedule	1-15
Permissions.....	1-16
Preferred Credentials	1-18
Setting Administrator Credentials	1-19
Saving Logon Information as a Preferred Credential.....	1-20
Saving the Contents of a List.....	1-20

2 Navigator

Navigator Pane	2-2
Populating the Navigator Tree	2-2
Expanding Objects in the Navigator.....	2-2
Viewing Specific Tree Objects.....	2-4
Launching Tools.....	2-4
Groups.....	2-5
Navigator Menu	2-5
Discovering Services	2-7
Adding a Service	2-8
Refreshing a Service	2-8
Removing a Service	2-9
Ping Agent	2-9
Manipulating Objects in the Navigator	2-9
Administering Objects	2-10
Copying Navigator Objects.....	2-10
Paging/Email Blackout	2-10
Specifying Total Paging/Email Blackout.....	2-10
Defining a Paging/Email Blackout	2-11
Making a Copy of an Existing Blackout Schedule	2-12
Turning Blackout Schedules On and Off	2-12

Deleting Paging/Email Blackout Periods.....	2-13
Viewing Blackout Periods	2-13

3 Groups

Group Pane	3-2
Group Menu	3-3
Managing Groups	3-4
Creating a Group.....	3-4
Group General Page	3-5
Group Permissions Page	3-6
Manipulating Group Views.....	3-7
Monitoring Status.....	3-7
Expanding Objects.....	3-8
Groups	3-8
Nodes	3-8
Databases and Other Discovered Services	3-8
Launching Tools from a Group.....	3-8
Adding Objects to a Group	3-9
Deleting Objects in a Group.....	3-9
Modifying a Group.....	3-10
Removing Groups	3-10
Example: Creating a Group.....	3-11

4 Jobs

Job Process	4-2
Job Tasks	4-2
Writing SQL*Plus Scripts	4-3
Job Credentials.....	4-4
Submitting Jobs.....	4-5
Cancelling Jobs.....	4-5
Job Pane	4-6
Active Page.....	4-7
Context-sensitive Menu options	4-8
History Page.....	4-8
Job Menu	4-10

Job Library	4-11
Creating, Modifying, or Viewing a Job	4-11
Creating a New Job.....	4-12
Modifying a Saved Job.....	4-12
Viewing Job Details	4-13
General Page.....	4-13
Job Task Page.....	4-15
Job Parameters Page	4-17
Job Schedule Page.....	4-17
Job Permissions Page.....	4-18
Job Progress Page.....	4-19
Output Dialog Box	4-20
Example: Creating a Job.....	4-21
Oracle Job Tasks	4-23
Oracle Database Tasks	4-24
Enterprise Manager Wizard Database Tasks	4-24
Operating System or Host Tasks	4-27
Tcl Script Examples	4-29
Listener Tasks.....	4-30

5 Events

Event System Overview.....	5-2
Using Events.....	5-2
Creating Events	5-3
Registering Events	5-3
Event Occurrences	5-3
Event Notifications	5-4
Notifying Administrators	5-4
Interpreting Events	5-5
Correcting Problems.....	5-5
Event Categories and Types	5-6
Fault Management Event Tests.....	5-6
Space Management Event Tests	5-7
Resource Management Event Tests	5-7
Performance Management Event Tests.....	5-7

Unsolicited Event Tests.....	5-8
Registering Interest in an Unsolicited Event	5-8
Setting the Parameters Property Sheet for Unsolicited Events	5-8
Raising an Unsolicited Event Test	5-10
Raising Unsolicited Events Via Enterprise Manager Jobs	5-10
Unsolicited Event Caveats	5-17
Event Pane	5-17
Alerts Page.....	5-18
Viewing Alerts	5-19
History Page	5-19
Registered Page.....	5-19
Event Menu	5-20
Context-Sensitive Menus.....	5-21
Event Library Dialog.....	5-21
Editing an Event	5-22
Oracle Event Tests	5-22
Event Viewer	5-22
General.....	5-23
Log	5-23
Notification Details.....	5-24
Creating, Modifying, and Registering an Event	5-24
Event General Page	5-26
Event Tests Page	5-28
Event Parameters Page	5-29
Parameters	5-29
Event Permissions Page.....	5-30
Event Progress Page	5-31
Example: Creating and Registering an Event.....	5-32
Administrator Event Notification	5-34
Oracle Event Tests	5-34
Database Fault Management Event Tests	5-35
Node Fault Management Events.....	5-35
Listener Fault Management Events	5-36
Data Gatherer Fault Management Events.....	5-36
Numeric Pager Job/Event Ids	5-36

General Event System Troubleshooting Procedures	5-37
Obtaining Enterprise Manager Repository Status	5-37
Removing Old Job and Events	5-38

6 Management Server and Intelligent Agent

Oracle Intelligent Agent	6-2
Characteristics	6-3
Event and Job Support	6-3
Network Encryption.....	6-3
SNMP Support	6-4
Management Server	6-4
Message Queues.....	6-5
Failed Queue.....	6-5
Notification Queue	6-5
Connection Cache	6-5

Part II The Database Administration Tools

7 DBA Management Pack

DBA Studio	7-2
Common Features of DBA Management Pack Functionality	7-2
Tree Views and Property Sheets.....	7-2
Multi-Column Lists	7-3
Database Version Awareness.....	7-3
Show SQL.....	7-3
Show Dependencies	7-3
Right-Mouse Commands.....	7-3
Advanced Mode.....	7-4
DBA Management Pack Functionality and Wizards	7-4
Instance Management	7-4
Database Operations	7-5
Stored Configurations	7-6
Sessions	7-6
Resource Consumer Groups	7-6

Resource Plans	7-7
Schema Management	7-8
Tree List by Schema	7-9
Editing an Object	7-10
Creating Objects	7-10
Security Management	7-10
User Operations	7-11
Role Operations	7-12
Profile Operations	7-12
Storage Management	7-13
Controlfile Operations	7-14
Tablespace Operations	7-14
Datafile Operations	7-14
Rollback Segment Operations	7-14
Redo Log Group Operations	7-15
Archive Log Operations	7-15
SQL*Plus Worksheet	7-15
Wizards	7-15
Backup Management Tools	7-16
Data Management Tools	7-16

8 Managing Backup and Recovery

Accessing the Backup Management Wizards and Property Sheets	8-2
Backup Wizard	8-3
Predefined Backup Strategy	8-5
Customize Backup Strategy	8-8
Recovery Wizard	8-15
Create Backup Configuration	8-20
Catalog Maintenance Wizard	8-25
Backup Configuration Library	8-30
Configuring a Remote Database for Backup or SYSDBA Administration	8-31
Server Side (NT)	8-31
Server Side (UNIX)	8-32
Recovery Catalog Setup and Registration for Oracle8i	8-33
Create a Tablespace	8-34

Create a User	8-34
Run the catrman Script (for pre-8i only)	8-35
Create the Recovery Catalog (for 8i and later) on the Target Machine	8-35
Registering the Recovery Catalog	8-36
Dropping the Recovery Catalog	8-37
Upgrading the Recovery Catalog	8-38

A Enterprise Manager Event Handler

How the Event Handler Works	A-2
How to Set Up the Event Handler	A-3
Defining Event Filters	A-4
Filter Syntax.....	A-4
The Event Logger Templates	A-6
The Command Executor Templates.....	A-8
Starting up and Running the Event Handler.....	A-9
Starting the Event Handler	A-10
Stopping the Event Handler	A-11
Troubleshooting Tips	A-12
Tracing.....	A-13
Sample Filters and Templates:	A-14
Filters:	A-14
Templates	A-15
Known Issues	A-15
Differences between Solaris and Windows NT	A-16

Index

Send Us Your Comments

Oracle Enterprise Manager Administrator's Guide, Release 2.1

Part No. A75683-01

Oracle Corporation welcomes your comments and suggestions on the quality and usefulness of this publication. 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 about this manual?

If you find any errors or have any other suggestions for improvement, please indicate 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
Redwood Shores, CA 94065 U.S.A.

If you would like a reply, please give your name, address, and telephone number below.

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
- Product Documentation
- Related Publications

Purpose of this Guide

This guide describes Oracle® Enterprise Manager, Release 2.1: Oracle's system management Console, common services, and integrated platform graphical tools. Oracle Enterprise Manager is bundled with the DBA Management Pack, which is 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 services into logical administrative groups

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 Solaris system. See the documentation for your Windows or Solaris system, if necessary.

How this Guide Is Organized

This guide is divided into chapters, as described below.

Part I: *Oracle Enterprise Manager Framework*

Chapter 1, "Console"

This chapter describes the Enterprise Manager Console.

Chapter 2, "Navigator"

This chapter explains how to use the Navigator component of Enterprise Manager.

Chapter 3, "Groups"

This chapter explains how to use the Map component of Enterprise Manager.

Chapter 4, "Jobs"

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

Chapter 5, "Events"

This chapter explains how to use the Event Management component of Enterprise Manager.

Chapter 6, "Management Server and Intelligent Agent"

This chapter discusses the Intelligent Agent and primary services of the Management Server.

Part II: *The Database Administration Tools*

Chapter 7, "DBA Management Pack"

This chapter introduces the DBA Management Pack, 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 8, "Managing Backup and Recovery"

This chapter describes 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 7, "DBA Management Pack". 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.

Product Documentation

The Oracle Enterprise Manager product documentation includes the following:

- The *Oracle Enterprise Manager Readme* provides important notes regarding the updates to the software, documentation, and late-breaking news.
- The *Oracle Enterprise Manager Configuration Guide* provides information about configuring Oracle Enterprise Manager components.
- The *Oracle Enterprise Manager Concepts Guide* provides an overview of the Enterprise Manager system.
- The *Oracle Enterprise Manager Administrator's Guide* describes the components and features of the Oracle Enterprise Manager system.
- The *Oracle Enterprise Manager Messages Manual* describes the Oracle Enterprise Manager error messages and methods for diagnosing the messages.
- The *Oracle Intelligent Agent User's Guide* provides information on configuring and troubleshooting the Version 1 Intelligent Agent.

The guides are available on the Oracle Enterprise Manager CD in HTML and Adobe Acrobat (PDF) formats for viewing with a web browser. In addition to the Enterprise Manager documentation, extensive online help is provided for Oracle Enterprise Manager components.

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 Oracle7 or Oracle8i documentation. The related books referred to in this guide are listed below:

- For general information about the Oracle8i Release 2 (8.1.6) and how it works, see *Getting to Know Oracle8i* and *Oracle8i Concepts*.
- For information about administering Oracle8i Release 2 (8.1.6), see the *Oracle8i Administrator's Guide*.
- For information on Oracle's SQL commands and functions, see the *Oracle8i SQL Reference*.
- For information about Oracle messages and codes, refer to *Oracle8i Error Messages*.
- For information about the Oracle networking system, see your network-specific documentation.
- For information about the Oracle8i Release 2 (8.1.6) on Windows NT, see the *Oracle8i Installation Guide for Windows NT*, the *Oracle8i Administrator's Guide for Windows NT*, and the *Release Notes*.
- For information on Oracle Parallel Server, refer to the *Oracle8i Parallel Server Setup and Configuration Guide*, which provides essential information for preparing Oracle8i Release 2 (8.1.6) for use with Oracle Parallel Server and Oracle Enterprise Manager. The "Introducing Oracle Parallel Server" chapter provides a conceptual and component overview of Oracle Parallel Server. This information will help you prepare and plan your Oracle Parallel Server installation and configuration in an Oracle8i Release 2 (8.1.6) environment.

Oracle Corporation also publishes several files that are available on your distribution media. These files are usually named README, RELNOTE, BUGHST, and RESTRICT and have extensions such as .WRI, .DOC, and .TXT. Read these files to learn about changes to the software or documentation that has not been described in the guides. The exact names and locations of the files mentioned above may vary, depending on your operating system.

Part I

Oracle Enterprise Manager Framework

- Chapter 1, "Console"
- Chapter 2, "Navigator"
- Chapter 3, "Groups"
- Chapter 4, "Jobs"
- Chapter 5, "Events"
- Chapter 6, "Management Server and Intelligent Agent"

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

- Console Framework
- Console Menus
- Managing Enterprise Manager Administrators

Console Framework

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 available through other vendors. The graphical configuration of the Console and the tools available are determined by the optional products purchased and user preferences. See Figure 1-2, "Enterprise Manager Console" for an illustration of a Console screen.

Components

The Console includes viewing panes for the Navigator, Group, Jobs, and Events systems. Other integral Enterprise Manager components are the repository, the Management Server, and Oracle Intelligent Agent.

Objects that administrators create, such as events and jobs, can be shared among Enterprise Manager administrators by assigning permissions. See "Permissions" on page 1-16 for more information.

Navigator

The Navigator displays a hierarchical list of all the services in a network, providing a direct view of services such as databases, groups, listeners, and nodes, plus the objects that they contain. The Navigator shows all the network services with their relationships to other objects.

Through the Navigator, you can also launch administration tools for managing objects such as administrators and tablespaces. You can also perform many administration tasks in the Navigator, such as creating, editing, or dropping database users. See Chapter 2, "Navigator" for information.

Group

The Group system allows you to organize managed services into logical groups that you can then manage collectively with a graphical display. With the Group system, you can create, save, modify, and view groups. To create a group of services, drag and drop objects from the Navigator into the Group view. See Chapter 3, "Groups" for information.

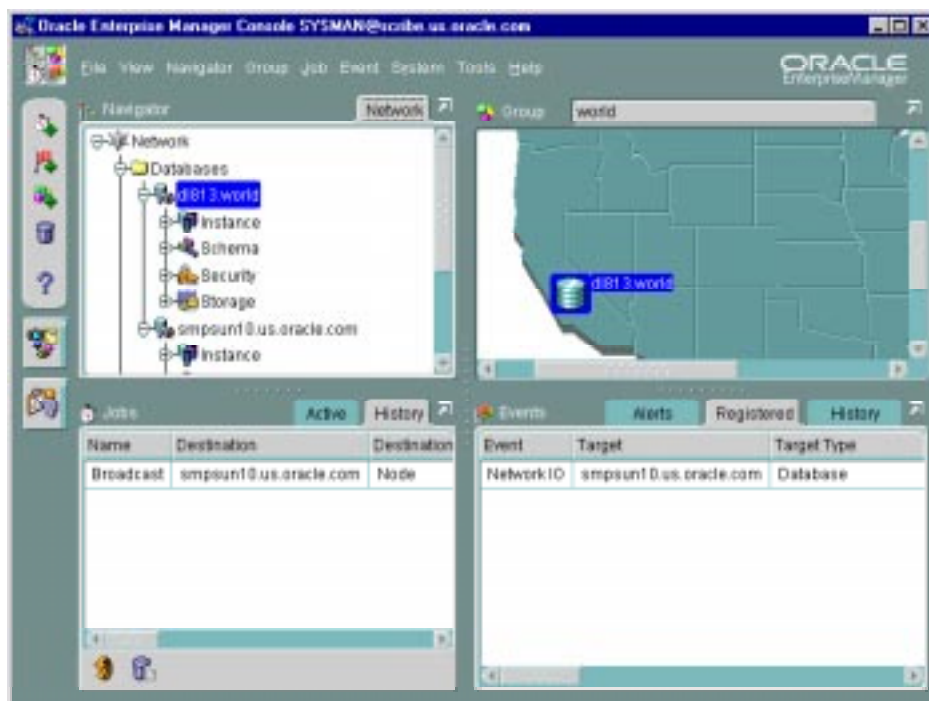
Jobs

Jobs are standard tasks, such as executing a SQL script or backing up a database, that can be automated through Oracle Enterprise Manager. The Job system allows you to manage and schedule jobs and remote operations for the services of your network. Jobs can be scheduled at various times, such as daily or weekly, and at single or multiple destinations. See Chapter 4, "Jobs" for more information.

Events

The Event system allows you to monitor and display the status of events—such as a loss of service—occurring on the managed services in your network system. See Chapter 5, "Events" for more information.

Figure 1-1 Console Window



Console Menus

The Console menu bar provides access to the following menus:

Figure 1–2 Console Menu Bar



The File, View, System, Tools, and Help menus are described in this section.

- See "Navigator Menu" on page 2-5 for information about the Navigator menu.
- See "Group Menu" on page 3-3 for information about the Group menu.
- See "Job Menu" on page 4-10 for information about the Job menu.
- See "Event Menu" on page 5-20 for information about the Event menu.

When using the Console menus, note that:

- 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 or the active window in the Console. 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 some windows of the Console to display a context-sensitive, or short-cut, menu. This menu usually contains a subset of the options that are available through a menu in the main menu bar. For example, if you click the right mouse button on a group in the Navigator pane, a menu appears with the menu options from the Group menu.

Note: See "Discovering Services" on page 2-7 for information on context-sensitive menu options in the Navigator.

File Menu

The File menu items allow you to exit the system.

View Menu

The View menu customizes the Console layout, allowing you to show or hide toolbars or windows. At least one window must be visible in the Console.

Show/Hide Navigator Pane

Hides or shows the Navigator pane.

Show/Hide Job Pane

Hides or shows the Job Scheduling pane.

Show/Hide Event Pane

Hides or shows the Event Management pane.

Show/Hide Group Pane

Hides or shows the Group pane.

Show/Hide Toolbar

Hides or shows the Oracle Enterprise Manager toolbar.

System Menu

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

Preferences...

Displays the Administrator Preferences property sheet, which allows the current administrator to change preferences, including permissions and login credentials. See "Preferred Credentials" on page 1-18 for more information.

Manage Administrators

Displays the Administrator Manager property sheets for adding, modifying, and removing administrators. This option is only available to Super Administrators. See "Managing Enterprise Manager Administrators" on page 1-7 for more information.

Configure SMTP Gateway

The System menu contains the Configure SMTP Gateway option. Use the Configure SMTP Gateway option in the Console System menu to specify the SMTP gateway that is used for the SMTP email system.

Add Services to tnsnames.ora

Updates the local network configuration file (tnsnames.ora) with discovered services information. You should run this after a Discovery operation so that any Value-Added Applications you may have can access the discovered services.

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.

Database Applications

This menu option lists the database administration tools and utilities that you can execute from the Console. See "DBA Management Pack" on page 1-8 for more information.

Service Management

This menu option lists integrated applications such as the Oracle Application Server Manager.

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. See "Online Help" on page 1-9 for more information about the online help system.

Managing Enterprise Manager Administrators

Enterprise Manager, Release 2.1 is a multi-administrator system: every person who is administering systems using Enterprise Manager has his or her own administrator account which he or she uses to log into the Console or stand-alone application.

The installation of Enterprise Manager creates a single super administrator. This super administrator account name is `sysman`, and the initial password is `oem_temp`.

The super administrator `sysman` creates administrators using the Manage Administrators item in the System menu. In addition to an administrator name and password, each account can be tagged as a "Super Administrator" account or a "regular" account.

Super administrators can bypass all security in the system and automatically have full privileges for all objects. 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, or checking other administrator's schedules. Using the `sysman` account for daily administration work is not recommended.

Regular administrators can have full access to all other Console operations but they will only see or be able to modify those jobs, events, or groups to which they have access.

Administrators share a single Enterprise Manager, Release 2.1 repository. 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 or stand-alone application, the preferred credentials used are those defined explicitly for that administrator.

Managing Administrators


Enterprise Manager administrators are created, edited, and deleted with the Manage Administrators option of the Console System menu. The Manage Administrators option displays the Manage Administrator Accounts dialog (Figure 1-3, "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:


- Select the Add or Edit button in the Manage Administrator Accounts dialog to display the Administrator Account dialog. The Administrator Account dialog allows you to add a new administrator or modify an existing one. See "Creating or Editing an Administrator Account" below for more information.
- Select the Delete button to delete 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 1-9 for more information.

Figure 1–3 *Manage Administrator Accounts*



Creating or Editing an Administrator Account

In the Manage Administrators dialog, enter a unique administrator name and password for the Enterprise Manager administrator. Note that the administrator names is not the same as a database username.

Figure 1-4 Create Administrator Account

Check the access available to the administrator:

Super Administrator

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

Access to Job System

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

Access to Event System

Allows the administrator access to the Event system component.

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


If you select the Edit button, you will get the Preferences window, as described in "Administrator Preferences" on page 1-10. You can then edit the preferences for the selected administrator.

Reassigning Object Ownership

If you delete an administrator who owns any object, the Object Ownership Management dialog in appears. See Figure 1-6. The dialog allows you to reassign the objects to another Enterprise Manager administrator. If you do not reassign the objects, the objects are deleted.

If the administrator has active objects, submitted jobs and events, those objects must be made inactive before they can be reassigned.

Figure 1–5 Object Ownership Management



Administrator Preferences

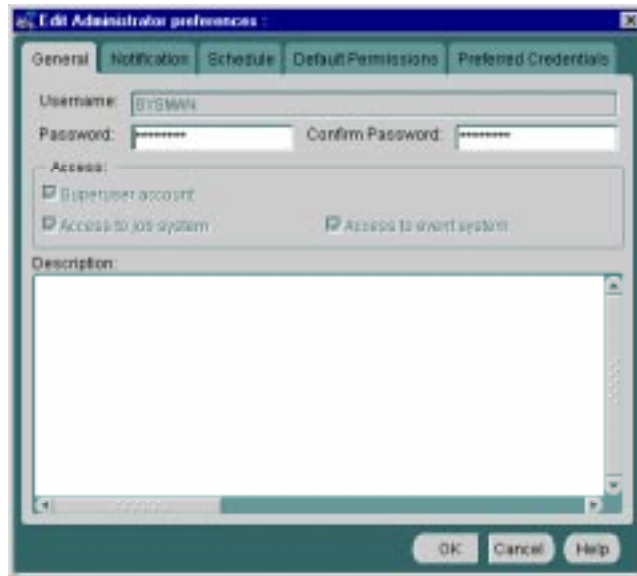
The Administrator Preferences property is displayed with the Preferences option of the System menu. This property sheet allows you to configure the preferences for an Enterprise Manager administrator. The pages in the property sheet are:

- General
- Notification
- Schedule
- Permissions
- Preferred Credentials

General

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

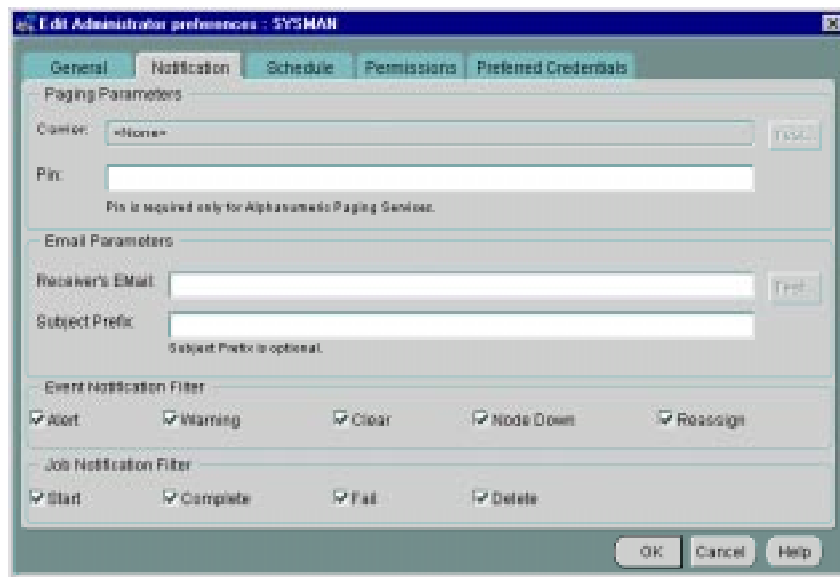
Figure 1–6 Administrator General Preferences



Notification

Complete this property sheet page to set up paging and email notification methods for the administrator. You can also send a test page or email with the Test options on this page.

Figure 1–7 Administrator Notification Preferences



Note: If you plan to use a paging service for notification, you need to configure the paging and modem settings for your system before you set up administrators. For information on paging setup, see the *Oracle Enterprise Manager Configuration Guide*. See "System Menu" on page 1-5 for information on email setup.

Paging Information

Carrier:

After a node containing paging services has been discovered, select the name of the paging carrier service from the pull-down list.

Pin:

Enter the paging personal identification number (PIN) for your page. A PIN is only required for alphanumeric pagers, and not for numeric pagers. This entry is required for the FLEX-TD, GSM and TAP protocols for alphanumeric pagers. For GSM, the PIN is the actual phone number of the administrator.

Test:

To test the validity of the paging information, select the paging service from the pull-down menu, enter the PIN number (if needed), and click the Test button. A message appears informing you of the status of the page.

If the test fails, check the paging services configuration file (paging.cfg). Refer to the *Enterprise Manager Configuration Guide* for explicit instructions on paging setup and configuration. 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 which has the Intelligent Agent Paging Extensions.

Note: Both the Management Server and the paging service must be running for paging to work.

Email Parameters

Receiver's Mail:

Enter the email address of the administrator to whom email notification will be sent.

Test:

Click the Test button to check the validity of the mail information. A message displays to inform you of the status of the test email.

If the test fails, check the email configuration. If tracing has been enabled for the Management Server, you can also view the mail trace log file in the `ORACLE_HOME\sysman\log` directory.

Subject Prefix:

Enter an optional prefix that is added to the email notifications. This allows you to further differentiate notifications from the Enterprise Manager Console that are sent to administrators.

Event Notification Filter

The Event Notification Filter allows you to filter email/pages sent to administrators according to the event's level of severity. Filtering is set at the user level through Administrator Preferences dialog. You can select any combination of the following levels of event severity. Selecting all levels of severity provides no filtering.

- Alert
- Warning
- Clear
- Node Down
- Reassign

Job Notification Filter

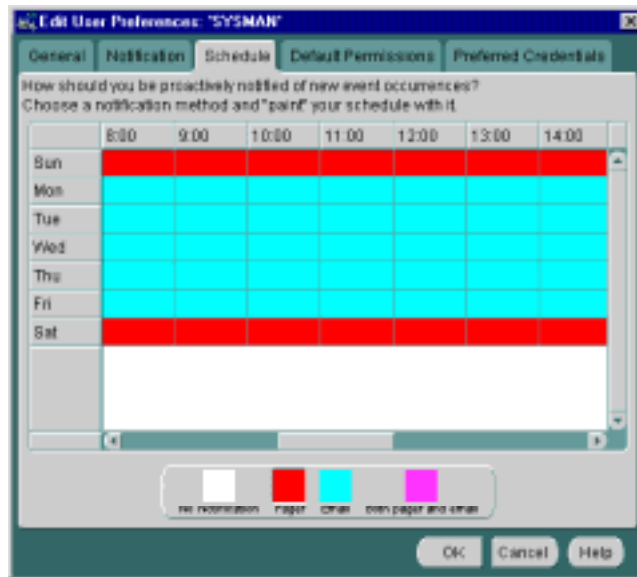
The Job Notification Filter allows you to filter email/pages sent to administrators according to job status. As with the Event Notification Filter, filtering is set at the user level through the Administrator Preferences dialog. Selecting all job statuses provides no filtering.

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

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 1–8 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:

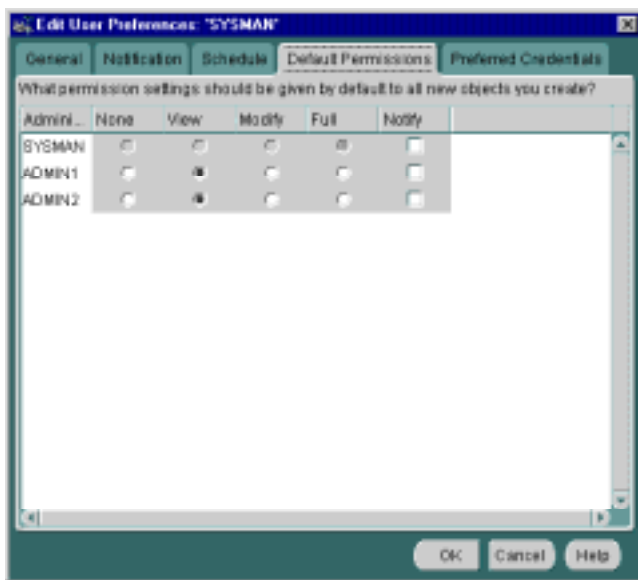
1. Select None, Pager, Email, or Email and Pager at the bottom of the dialog.
2. Click in the appropriate day/hour to set the notification method for that time period.
3. Repeat steps 1 and 2 for each notification method.

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.

Permissions

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 1–9 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 Permissions Page" on page 3-6
- "Job Permissions Page" on page 4-18
- "Event Permissions Page" on page 5-30

Note: Changing your default permissions does not retroactively change permissions on existing objects.

The levels of permission that you can assign to an Enterprise Manager administrator are:

None

This permission does not allow the administrator to view this object anywhere.

View

This permission 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 permission allows the administrator to edit the object's properties except those reserved for Full permission.

Full

This permission allows the administrator to delete the object, modify permissions for other administrators, and change the ownership of the object.

Notify

This permission allows the administrator to receive enhanced 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 services in the network, along with the service type and the administrator name for accessing the service. The property sheet page is accessed with the Administrator Preferences option of the System menu. Each row in the list of the property sheet includes:

- Service Name
- Service 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 1–10, "Administrator Preferred Credentials" for an illustration of the property sheet.

Note: Individual instances of an Oracle Parallel Server are listed. Oracle recommends that all instances of a Parallel server use the same preferred credentials as the Oracle Parallel Server.

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 4-4 and "Registering Events" on page 5-3 for more information.

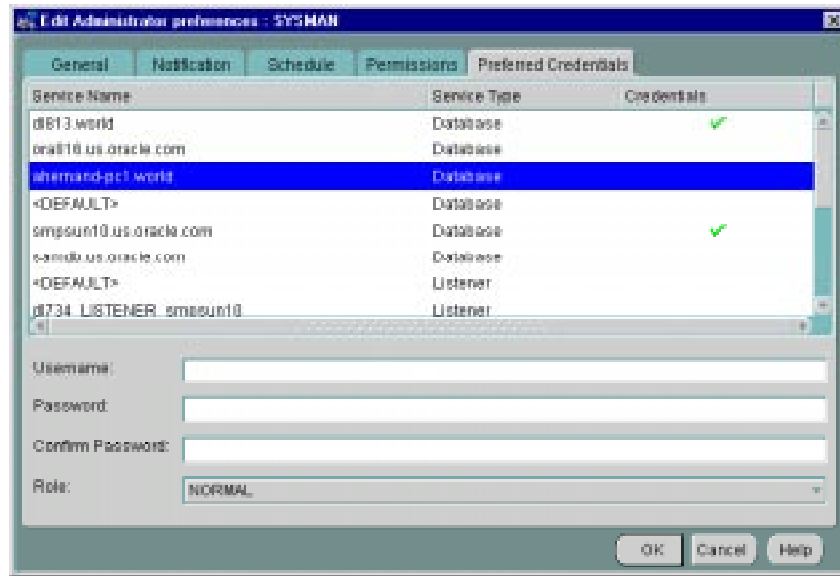
Oracle recommends that if some services of a particular type require the same credentials, use the <DEFAULT> credentials selection for that service type to have those credentials used by all services of that type.

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 1–10 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. Use the mouse to select other lines on the page. 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 list. By selecting the Save Event History from the Events 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 desire 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.

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 Services
- Manipulating Objects in the Navigator
- Paging/Email Blackout

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.

The Navigator window can be hidden or shown by selecting Show Navigator Pane in the Console View menu. You can also hide or show the pane by clicking on the Navigator icon in the Console toolbar.

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 include:

- Databases
- Groups
- Listeners
- Nodes

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 2-1, "Navigator Menu and Window" for an illustration of a Navigator window.

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

- Databases folder

- Groups folder
- Listeners folder
- Nodes folder
- Folders for additional services discovered

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 view datafiles, in-doubt transactions, profiles, redo log groups, roles, rollback segments, schema objects, tablespaces, and users. When a connection is made to a database, an icon displays on the database in the tree. See "DBA Studio" on page 7-2 for more information on integrated database administration tools that are accessible from specific Navigator objects.

information and optionally save that information as a preferred credential. See "Preferred Credentials" on page 1-18 for information.

Groups

Groups are created with the Create Group option of the Group menu. You can use groups to organize related objects and other discovered types in the tree. You can drag and drop nodes, databases, groups, and listeners from the Navigator tree to a group folder or the Group window to add those objects to a group. See Chapter 3, "Groups" for information on groups.

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 2-1, "Navigator Menu and Window" for an illustration of the Navigator menu.

Clicking on an object with the right mouse button, displays a *context 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 7, "DBA Management Pack" for an overview of the database application tools.

Create

Allows you to create an object of a specified type.

Create Like

Allows you to create a copy of the selected object in the Navigator.

Edit

Allows you to alter or edit the object selected in the Navigator.

Remove

Allows you to delete all occurrences of the selected object in the Navigator.

If you remove a node from the Navigator tree, the node and all its services are removed from the Navigator tree.

If you remove a database, listener, or other service object from the Navigator tree the object appears again in the Navigator tree when the node that contains the service is rediscovered.

Discover Node

Discovers services on a node in the network. You can also display the status of network services. See "Discovering Services" on page 2-7 for more information.

Refresh Node

Refreshes services on any node selected in the Navigator.

Paging/Email Blackout

Allows a super-user to specify (enhanced notification) paging and email blackouts for specified targets and/or services that have been previously discovered in the Navigator.

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.

Split View

Creates a new view page in the Navigator window that contains only the currently selected object, such as a single folder, group, or user. This feature allows you to create separate views that display specific objects. For example, to create a new view that only contains the DEMO database:

1. Select the DEMO database in the tree.
2. Select the Split View menu option to create a separate view page in the Navigator pane.
3. Click on the tab of the page to move the page and its contents to the front of the Navigator pane.

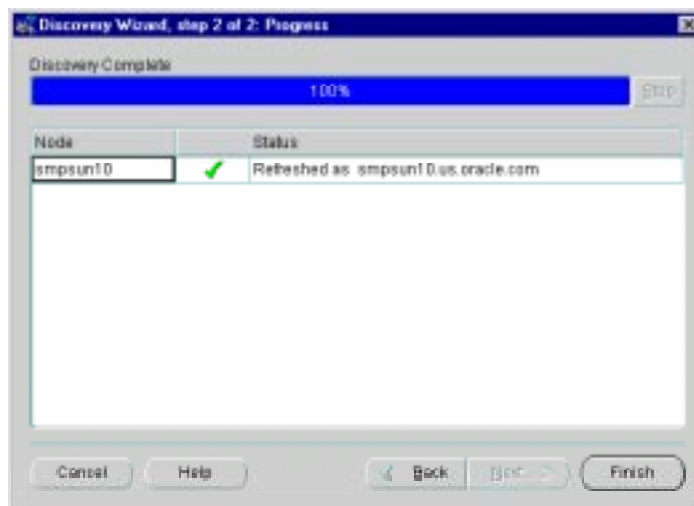
Remove Split View

Removes the current split view from the Navigator. Make sure the view that you want to delete is at the front of the Navigator pane.

Discovering Services

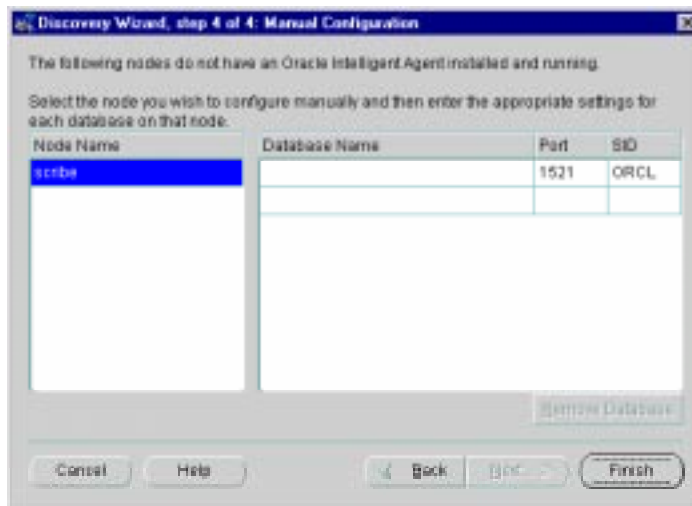
The Navigator provides service discovery functions for identifying network services and populating the Navigator tree through the Discovery/Refresh Wizard.

Figure 2–2 Discovery Wizard



The Discovery Wizard searches the network for the nodes you specify. If that node 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 node. You will be able use these services as targets in jobs and events.

If there is no Oracle Intelligent Agent running, you may perform a manual configuration. This allows you to define an Oracle database on that node 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.

Figure 2–3 Discovery Wizard: Manual Configuration

When refreshing an existing node, the Discovery Wizard verifies and updates the list of services on a node that had been previously discovered. If the node was previously Manually Configured, the Discovery Wizard will check if that node 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 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 node 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 Service

Refreshing a node verifies and updates the list of services on a node that has been previously discovered. You can refresh the discovery of a node by selecting a node

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 node selected in the Console Navigator. This menu option is greyed out until a specific node is selected from the Navigator Nodes folder.

Removing a Service

Select a node or service from the Console Navigator. Using the right-mouse button, choose Delete from the context-sensitive menu.

Ping 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 to check 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 also.
- 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).

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 2-5 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 an administrator or role from one database to another to add that role to another user. Note, however, that 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.

To make an identical copy of a database user, select the user you want to copy from one database and drag and drop the user to the database where you want the user to appear. To add a role or privilege to a user in a database, select the role or privilege in the tree and drag and drop the role or privilege onto the user.

Paging/Email Blackout

Paging/Email Blackout allows an administrator with super-user 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-user administrator can schedule a paging/email blackout for that target to prevent enhanced notification during the maintenance period.

Specifying Total Paging/Email Blackout

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

1. Choose Paging/Email Blackout from the Navigator menu.
2. Check Total Paging/Email Blackout.
3. Click OK.

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

Defining a Paging/Email Blackout

1. Select a target from the Navigator and then choose Paging/Email Blackout from the Navigator menu. Optionally, you can select an object in the Navigator and access 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

1. Choose Paging/Email Blackout from the Navigator 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 Navigator 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.

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 Pane
- Group Menu
- Managing Groups

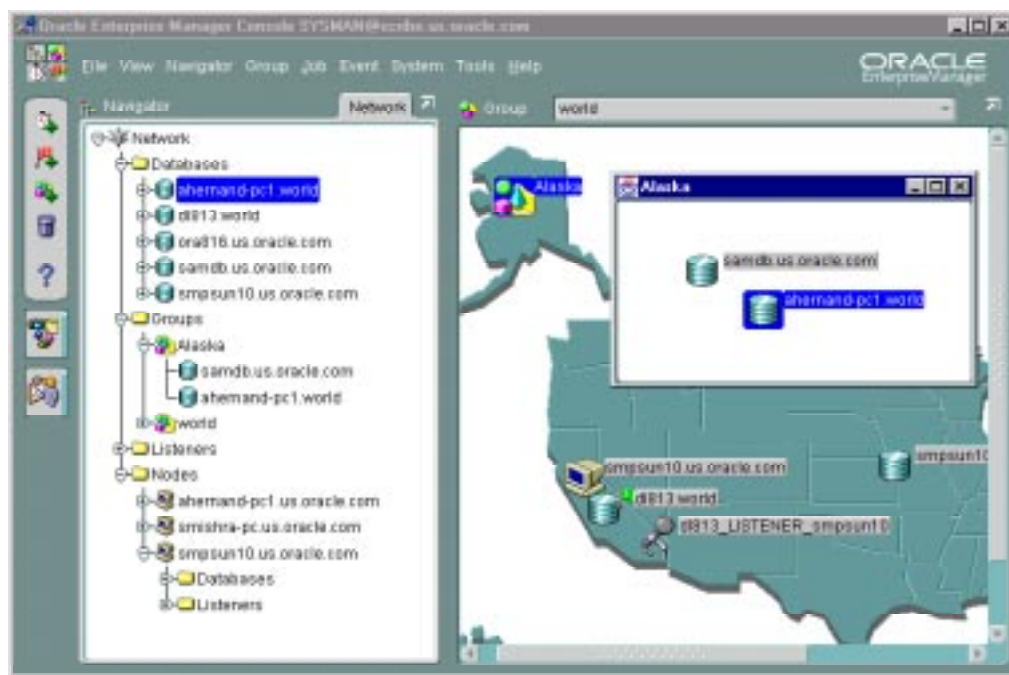
Group Pane

The Group pane provides an area for constructing a graphical custom view of the network environment. When all four panes are displayed, the Groups pane is located in the upper right corner of the console. See Figure 3–1, "Navigator and Group Panes" for an illustration of a Group pane.

The Group pane can be hidden or shown by selecting Show/Hide Group pane in the Console View menu. You can also hide or show the pane by clicking on the Show/Hide Group Display button above the Console toolbar.

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

Figure 3–1 Navigator and Group Panes



Group Menu

The Group menu options allow you to perform such actions as create, modify, and remove groups..

Create Group

Adds a new group to the current group or a selected group. See "Adding Objects to a Group" on page 3-9 for more information.

Create Group Like

Copies an existing group.

Edit Group

Modifies an existing group. See "Modifying a Group" on page 3-10 for more information.

Remove Group

Removes an existing group from the current group, or deletes the group from the entire system. See "Removing Groups" on page 3-10 for more information.

Remove Item From Group

Deletes the selected item from the group pane. You can also delete an object from the current group or group by selecting the object and pressing the Delete key. See "Removing Groups" on page 3-10 for more information.

Managing Groups

A group is a collection of objects, such as databases, listeners, or nodes, 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 pane.

A group is represented on 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 pane gives you a more 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 an object 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.

Creating a Group

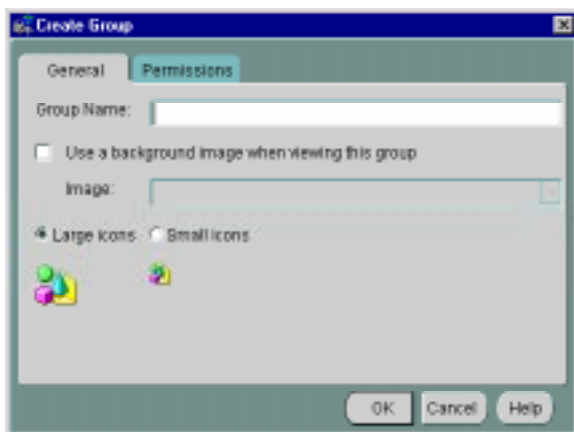
To create a group, select the Create Group menu item to display the Create a Group property sheet. The property sheet contains the General and Permissions pages (see Figure 3-2 and Figure 3-3). After you have completed the property sheet pages,

click the OK button to save the group. The new group appears under Groups in the Navigator. The Group pane will remain empty.

The new group is empty until you add objects to it. Dragging an object from the Navigator menu copies and pastes it into a group. Note that this action does not remove the object from the Navigator. In fact, you can add an object to any number of groups, and still be able to manipulate it individually, as well.

To add an object to a group, select it and drag it over to the Group pane, or drag it over to a group folder in the Navigator pane. To add an existing group to the new group, or to another existing group, drag and drop the group from the Navigator into the group folder or the Group pane. An icon representing the object or group appears in the Navigator and in the Group pane.

Figure 3–2 Create Group General Page



Group General Page

Group Name

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

Use a background image when viewing this group

Check this option if you want to display an image in the background of the group pane. Sample files are provided in the `oracle_home\classes\oracle\sysman\resources\images` directory on the machine where the Management Server is located. If you want to use your own images, Oracle recommends that you add them to this directory.

The sample images included are:

- asia.gif
- europe.gif
- japan.gif
- usa.gif
- world.gif

Do not specify a directory for graphics located in this directory. Specify only the file name, such as `usa.gif`. You can use `.GIF` or `.JPG` graphic files. If you check the box, enter the file name of the image in the Bitmap Relative URL field. When an image is added to a group, the file name is stored in the repository.

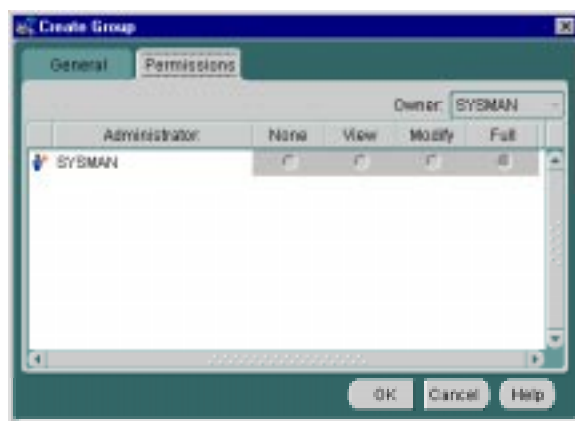
Icons

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

Group Permissions Page

Determine the permissions that you want to assign to the group with the Permissions Page. This allows other users to view or modify the group. Notifications are not assigned with this page. Any permissions assigned on this page over-write any user default permissions. See "Permissions" on page 1-16 for an explanation of the permissions that can be assigned.

Figure 3–3 *Group Permissions Page*



Manipulating Group Views

Groups are populated by dragging and dropping databases, groups, nodes, listeners or any other discovered services from the Navigator into the Group pane. 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.

Double-clicking on a group in the Navigator pane opens the hierarchical list of its members. You can display groups in the Group pane by selecting the group name from the pull-down list in the Console toolbar.

Monitoring Status

The Group system 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 signal 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).

Groups inherit the worst state of any of the members. If one service in a group is down, the group displays a grey flagt.

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 5, "Events" for more information on events.

Expanding Objects

Group objects can be expanded in the Group pane 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.

See Figure 3-1, "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 group expands to display the objects in the group. You can drag and drop objects from the Group, the Navigator, and other groups to add them to the group. You can delete objects from the group by selecting them and pressing the Delete key, or clicking the right mouse button and selecting Remove.

Nodes

When you double-click on a node icon in the Navigator pane, the node expands to display the databases and listeners on the node.

Databases and Other Discovered Services

When you double-click on a database or other discovered service, 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 1-18 for information on preferred credentials.

The database instance property sheet contains the Status, Startup, and Shutdown pages. You cannot start up or shut down a database instance unless you have connected to the database as SYSDBA or SYSOPER.

Launching Tools from a Group

You can launch a database tool using an object in the group. Select a database icon in the Group pane, then select a DBA tool from the Tools Applications menu or launch palette. You are connected to the database according to the user credentials that have been set up for the system. See "Preferred Credentials" on page 1-18 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—either the group folder on the Navigator or the group shown in the Group pane.

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 in a Group

To delete objects in a group pane:

1. Double-click on a group icon in a group to expand the group.
2. Select an object in the group and press the Delete key to delete an object from a group.

To delete objects from a group in the Navigator:

1. Expand the group in the tree.
2. Select an object in the group and press the Delete key to delete an object from a group.

Note: When you delete 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.

Modifying a Group

You can modify the properties of a group.

1. Select a group in the Navigator.
2. Select the Edit Group menu item to display the Edit Group dialog box. 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 OK button when you have finished.

Removing Groups

There are two ways to remove a group:

- 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 four ways to remove it:

- Delete the group completely by selecting Remove Group from the Group menu.
- Delete the group completely by choosing a group at the top level (a group that is not inside another group) and selecting the Remove menu item from the group's context menu.
- Delete the group completely by choosing a group at the top level (a group that is not inside another group) and selecting the <DELETE> key.
- Remove a sub-group from within another group by selecting Remove Item from the Group 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 Remove 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 selecting 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.

Example: Creating a Group

The following example illustrates how to create a new database group.

1. Select Create Group from the Group menu to display the Create a Group dialog.
2. Enter a name for the new group in the Group Name field.
3. Check the Use Background Bitmap box to add a graphic to the background. Enter a graphic name, such as europe.gif, in the Bitmap Relative URL field.
4. Choose small icons for the group.
5. Set the permissions so all Enterprise Manager administrators can view but not modify the group.
6. Click on the Create button to create the new group. A folder for the new group appears in the Navigator.
7. Select the new group from the pull-down group list over the Group pane to display the group.
8. Expand the Databases branch of the Navigator.
9. Drag a database into the Group pane.
10. Drag and drop a listener into the new group folder in the Navigator. This listener should appear in the group pane.

The new group can be added to other groups by dragging and dropping the group from the Navigator.

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 Pane
- 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 destination (target type) of the task:

- Databases
- Nodes
- Listeners

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 4-23 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 5, "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 *Oracle Enterprise Manager Configuration Guide* 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 service 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 node. 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 1-18 for information on administrator preferences.

Important: You must set up valid user credentials for the nodes that you want to run jobs on. 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 users 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. Default Service
4. Null Credentials (Lowest Priority) **Note:** Jobs and events will fail.

For more information on configuration issues, see the *Oracle Enterprise Manager Configuration Guide*.

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 destinations 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 destinations 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 to the node on which the job will be run. You only need to submit the job from the Console and specify the destinations on which it should run. The destinations 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.

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 Pane

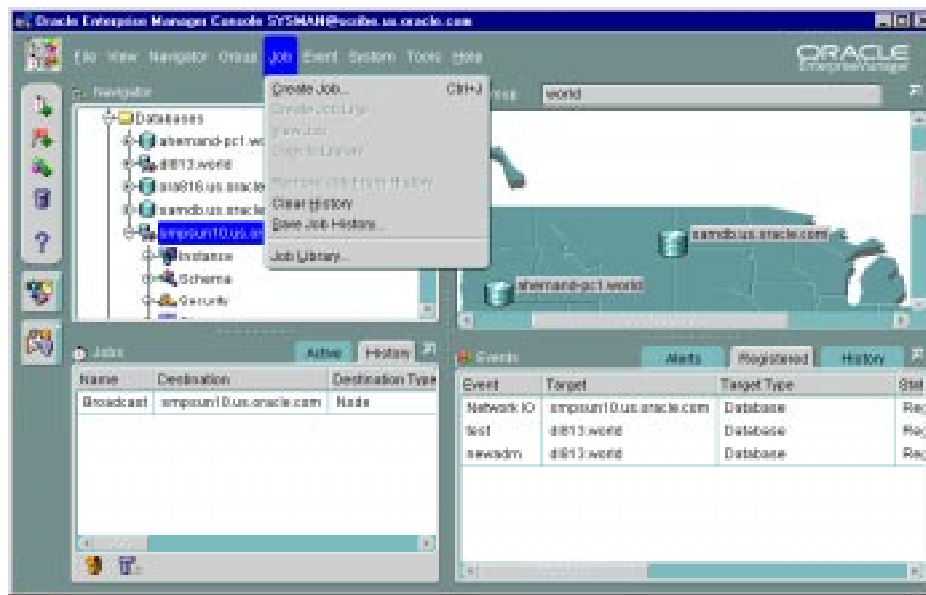
You can view the different pages of job information by selecting the page tabs in the Job pane. The pages in the Job pane are:

- 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.

The Job pane can be hidden or shown by selecting Job in the Console View menu. You can also hide or show the pane by clicking the Job system icon in the Console toolbar.

Figure 4–1 Job Pane in the Console



Active 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 View Job menu option to display the details of the selected job. You cannot modify these jobs.

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

Name

Name of the job.

Destination

Destination of the job.

Destination Type

database, listener, node

Owner

Administrator submitting the job.

Status

Status of job is one of the following:

- *Submitted*: The job has been submitted to the Management Server at the job destination.
- *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.

Edit/View Job

Displays the property sheet for the selected job. Only the permissions, notifications, and owner 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. This page also lists deleted jobs. You cannot modify these jobs.

- Job Name is the name of the job.
- Destination of the job.
- Owner of the job.
- 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 Page" on page 4-7.
- Finish Time is the time when the job finished, failed, or was deleted.

Refresh History Icon

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 click the Refresh History icon located at the bottom left of the Job pane.

Clear History Icon

Clears all entries in the Job History list.

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. Only the permissions, notifications, and owner 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.

You can save the jobs from the History page to a text 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. Jobs are saved to a text 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 4-11 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 4-1, "Job Pane 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 4-11 for more information.

Create 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 4-11 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. See "Creating, Modifying, or Viewing a Job" on page 4-11 for more information.

Copy 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 4-11 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 4-5 for specifics on job cancellation.

Job Library

Displays the Job Library dialog. See "Job Library" on page 4-11 for more information.

Save Job History

Saves the contents of the Job History page to a text file. Enter a valid file name or select an existing file in the file selection dialog box.

Clear History

Clears the jobs listed in the Job History page.

Remove Job from History

Clears the selected job listed in the Job History page.

Job Library

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 4-3, "Job Tasks Page" for an illustration of a Job property sheet. The property sheets contains:

- General Page
- Task Page

- Parameters Page
- Schedule Page
- Permissions Page
- Progress Page

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 OK to submit the job to the Intelligent Agents at the selected destinations. The job appears in the Active page.
- 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 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.

Modifying a Saved Job

Note: Only saved jobs that are in the Job Library can be fully modified. For submitted jobs in the Active Jobs page, you can only modify permissions, owners, or notifications.

1. Select Job Library from the Job menu. The Job Library dialog appears.

2. Select a job from the library.
3. Click Edit. The Job property sheet appears.
4. 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.

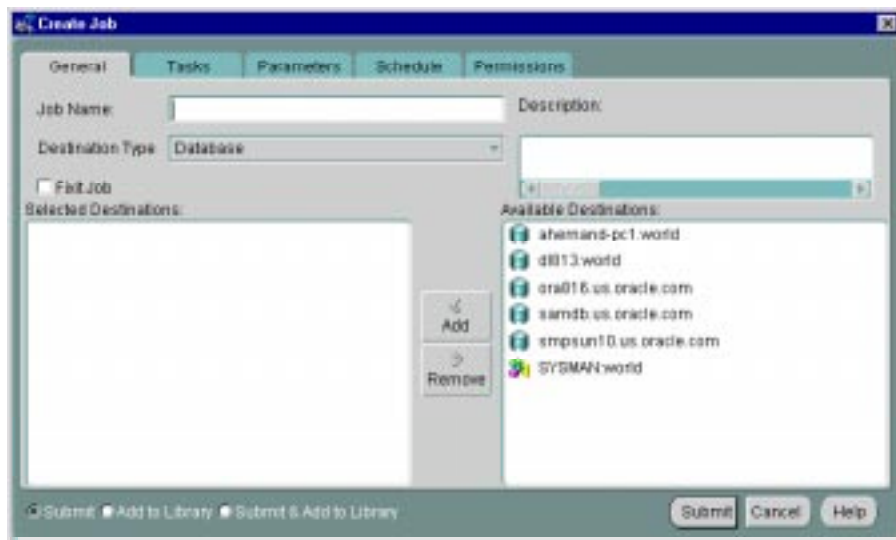
Viewing Job Details

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

General Page

The General page allows you to determine job name, description, type, fixit job status, and destinations.

Figure 4–2 Job General Page



Owner Name

Select the owner from the list of administrators.

Job Name

Enter the name of the new job.

Description

Enter the description of the job.

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 destination 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 Parameters page after it has been successfully submitted to an Intelligent Agent. See "Creating, Modifying, and Registering an Event" on page 5-24 for more information.

Destination Type

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

Job destinations

Select the destinations of the job in the Available Destinations list and click the << (Add) button to move the destination to the Selected Destinations list.

The destinations are determined by the Destination Type. The destinations include databases, listeners, nodes, and groups of these objects.

Note: Only network objects that have been discovered correctly and have an Intelligent Agent running are included in the list of available destinations. See "Discovering Services" on page 2-7 for more information.

Job Task Page

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

Available Tasks

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 4-23 for information on Oracle predefined job tasks and their parameters.

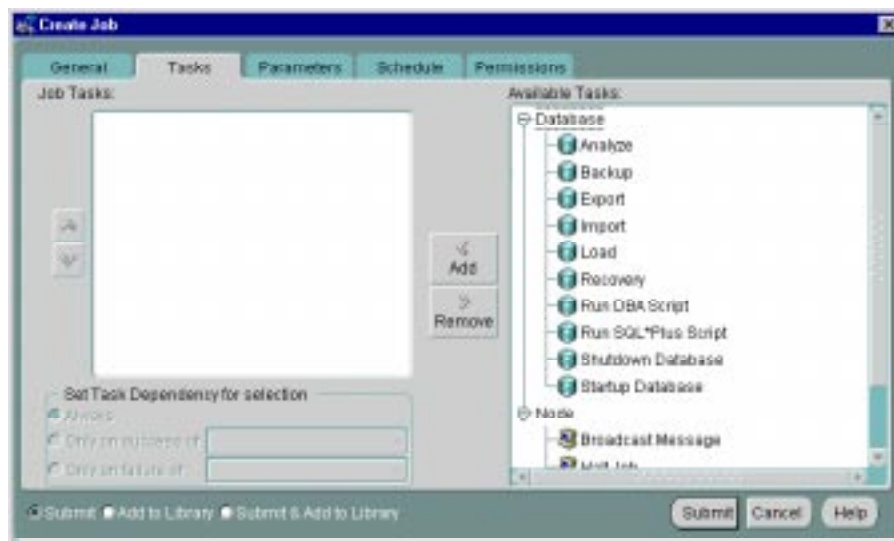
Selected Tasks

You can remove tasks from this list. Select the task and click 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. You can change the order in which the tasks are executed. Select a task in the Selected Tasks list and click on the up or down arrow button to position the task.

Figure 4–3 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: 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.

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.

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 4-23 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 4-4 and "Preferred Credentials" on page 1-18 for information on administrator preferences.

Job Schedule Page

The Schedule page allows you to schedule the execution of the job task.

Run Job

Select the frequency that you want the task executed. The choices are Immediately, Once, On Interval, On Day of Week, and On Date of Month.

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.

Stop 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.

Note: When scheduling a job, the Intelligent Agent time zone (the local time zone of the destination where the job will be executed) is used.

Job Permissions Page

Determine the permissions that you want to assign to the job with the Permissions Page. This allows other administrators to view or modify the job. Notifications can be assigned with this page.

The levels of permission that you can assign to an Enterprise Manager administrator are:

None

Does not allow the administrator to view this job anywhere.

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 edit the job's properties except those reserved for administrators with Full permission.

Full

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

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. This option must be selected in order for the job to appear in the Notification Schedule.

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 "Permissions" on page 1-16 for more information.

Show Notification Schedule

Show Notification Schedule displays the notification schedule for the job. You can view the administrators that are notified by moving the cursor over any time block in the grid.

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.

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 notifications 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, Running, and Completed or Failed notifications for that execution. If you display an execution from the Active Jobs page, the Progress page displays only those notifications that have been received.

When you display the Progress page, the page displays the notifications only for the destination and execution time of the job occurrence selected. To view the notifications associated with other destinations or execution times, select other destinations or execution times from the Destination or Execution pull-down lists. You can also select <All> in either list to view all notifications. If the job has been Deleted on a destination, the Deleted notification always displays at the top of the Progress Page.

The following options are available on the Progress page:

Destination

Select the destination 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 As

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

Print

Select the Print button to print the list of job occurrences in the Progress page.

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 Page" on page 4-7 and "History Page" on page 4-8 for information on the status of a job.

Destination

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

Notification Time

This is the time the Oracle Management Server was notified.

Output Dialog Box

The Output dialog box 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 that states that 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.

Print

Select the Print button to print the contents of the dialog box.

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.joboutput.maxsize=128K
```

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

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. Do not change the Owner Name.
3. 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.
4. Select Database from the Destination Type pull-down list.
5. Do not check the Fixit Job box for this job. You can create jobs that are used specifically as fixit jobs. Fixit jobs can be run automatically in response to an event occurrence. See "Event Parameters Page" on page 5-29 for information on fixit jobs with events.
6. Select a destination from the Available Destinations list, then click on the << (Add) button to add the destination to the Selected Destinations list. These are

destinations where an Intelligent Agent is running. The Job Credentials must be set up correctly for these destinations. See "Job Credentials" on page 4-4 for more information.

7. Repeat the previous step for another destination. You are choosing the destinations where the job will be run.
8. Click on the Task Page tab of the Create Job property sheet.
9. 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.
10. 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.
11. Check the box to send an enhanced notification if the job task fails.
12. Click on the Parameters Page tab of the Create Job property sheet to set the parameters:
 - Enter "SELECT * FROM dba_users;" in the 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 4-4 for permissions needed to run job tasks.
13. Click on the Schedule Page tab of the Create Job property sheet to schedule the execution of the job.
14. Select the On Interval under Run Job to execute the job on a specific interval.
15. Set the Start Running Date of the job to 7/1/2000. 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.
16. Repeat the process in the previous step for the day and year values of the Start Running Date.
17. Change the Start Running Time to 12:00 AM. Use the same procedure that you used to set the date.
18. Check the Stop Running box to set a final execution date for this job.
19. Set the Stop Running Date of the job to 12/30/2000. Set the Stop Running Time of the job to 12:00 AM.

20. 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.
21. 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.
22. Click on the Permissions Page tab of the Create Job property sheet to assign permissions on this job for other administrators.
23. Allow all other Enterprise Manager administrators to view and receive notifications for this job.
24. 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.
25. 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 destination 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 7/1/2000 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)

Oracle Job Tasks

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

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

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
- Halt Job

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:

- 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 Details 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 1-18 for more information.

3. Script Text. You can copy and paste your script into 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 1-18 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 1-18 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 1-18 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 restrict this session of the database.
 7. **Force Startup.** Check this box if you want to force the startup of a database that is running.

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 Host Tasks

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

- Broadcast Message
- Run OS Command
- Run Tcl Script

Broadcast Message

This job allows you to submit a message to the selected destinations using the platform-specific mechanism. To send the message to a destination, 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: `-l /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 1-18 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. Copy and paste the Tcl script into the Script Text box.

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 $curl]
while {$oramsmsg(rc) == 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@or803.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
    orarepotevent /user/host/file/alert $oramsmsg(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 5-26 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 1-18 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 1-18 for information on user preferences.

Parameters:

None

Important: See "Numeric Pager Job/Event Ids" on page 5-36 for specific ids.

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 Categories and Types
- Event Pane
- Event Menu
- Creating, Modifying, and Registering an Event
- Oracle Event Tests

Event System Overview

The Event system allows you to efficiently monitor a large system. 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 destination and frequency of the event.
 - b. Selecting the event tests that you want to run.
 - c. Determining the parameters for the event tests.
 - d. Assigning permissions to allow other administrators to share the event.
2. Saving and modifying an event.
3. Registering, or submitting, an event to a destination.
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 service type, for instance:

- Database
- Listener
- 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 5-6 for more information.

The events are created with information entered in the Event property sheet. You determine parameters such as the destination that is monitored, the specific tests to perform, the frequency that the event test is executed, and whether other administrators can share the events. See "Permissions" on page 1-16 for more information. Some event tests have parameters with threshold values that you can customize for your system. See "Event Parameters Page" on page 5-29 for more information.

Registering Events

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

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 1-18 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 Manager 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 "Frequency" on page 5-27 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 5-6 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 5-1, "Event Menu and Pane" 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 email 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 setups and permissions. You need to set up the notification services and determine the administrators that need to be notified for the events. See "Event Permissions Page" on page 5-30 and "Permissions" on page 1-16 to determine the administrators that receive notifications. See "Notification" on page 1-12 and "Schedule" on page 1-15 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 are set up properly:

- The settings for the system modem.
- The mail and paging services that are used to contact administrators.
- The mail address and paging numbers for each administrator. See the *Oracle Enterprise Manager Configuration Guide* for more information on setting up enhanced notification.

Interpreting Events

All events return values and some events produce output messages. The events return color values depending on the severity of the event. These severity levels are determined by parameters you set for the event tests during event creation. The colors are displayed on the event severity flag that is located:

- Next to the event name listed in the Alerts page.
- On the object in the Group pane. See Chapter 3, "Groups" for information on the Group pane.

The colors of the event severity flags are:

- Event cleared (green)
- Warning (yellow)
- Critical (red)
- Unknown (gray) -- Intelligent Agent down or node inaccessible

If an UpDown event test is specified for a particular event, and the node to which the event is assigned goes down, the event will return an "unknown" severity (gray flag). The original event severity will be restored when the node comes back up. At least one UpDown event test (node, listener, database, etc.) must be specified for an event in order for the "unknown" severity to be returned. Events not containing an UpDown event test will not provide any node down notification.

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. These situations are noted in the online help for Oracle events.

In other cases, the solution requires careful attention by 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.

For information on how to correct event conditions, refer to the appropriate documentation. To correct 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 events

Only the UpDown event tests are included with Oracle Enterprise Manager. Additional advanced events are available with the optional Oracle Diagnostics Pack.

See the online help for Oracle predefined event tests and "Oracle Event Tests" on page 5-34 for information on events and their parameters. 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. 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.

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 event tests 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 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

To register interest in an unsolicited event, check the Unsolicited Event Test box in the Event General page and complete the information in the succeeding property sheet pages. Select the Unsolicited event in the Test page and complete the Parameters pages. 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 "Creating, Modifying, and Registering an Event" on page 5-24 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 following field entries:

- **Event Name:** Name of the event
- **Object:** Name of the object that the event is monitoring
- **Severity:** Level of severity of the event
- **Message:** Message to be displayed in the Console

You can enter a wildcard "*" (no filtering) or a character string in these fields to

define the filtering mechanism.

Event Name

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

`/vendor/product/category/name`

You can enter any character strings 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.

Enter a "*" if you do not want to filter on an event name.

Object

An object can be any service the unsolicited event is monitoring, such as a database, listener, or node name. **Important:** The object must be a service that the Intelligent Agent has already discovered.

Enter a wildcard character ("*") if you want to receive unsolicited events from any type of object. Or, you can use the percent sign "%" for partial filtering. For example, `ORCL%` for events fired on objects whose name starts with "ORCL."

Severity

Enter the numeric severity level (-1, 1, 2) corresponding to the severity of the event you want to filter through.

Severity levels map to different flags in the Console Event Pane:

Table 5–1 Severity Level/Event Flag Mapping

Flag Color	Severity	Severity Level
Green	Cleared	Severity -1
Yellow	Warning	Severity 1
Red	Critical or Alert	Severity 2

Message

Message is a quoted text string that is displayed in the Console, such as "File not found." If a message filter is specified, then the event is passed only if the message matches what is specified in the message filter.

You can specify wildcard characters as part of the filter i.e. Message Filter : ERROR%. In this case the unsolicited event should fire on anything with a message that starts with "ERROR:"

Raising an Unsolicited Event Test

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` 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 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 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 an example of using OraTcl to trigger a third-party event, see *Tcl Script Examples* in the Run Tcl job task. For information on OraTcl and event scripts, see the *Oracle Intelligent Agent User's Guide*.

Raising Unsolicited Events Via Enterprise Manager Jobs

Typically, unsolicited events are evaluated and raised by third-party software. Enterprise Manager allows you to implement 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. Using techniques like the ones scribed below, users of Enterprise can implement 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 event raised by `oemevent`, will receive the event notification. The event has to be known to the administrator submitting the job that raises it.

However, 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 OEM event tests.

Unsolicited events (their logic actually) are evaluated in their own process and within the proper OS security envelope. For this reason, no longer pose security or robustness threats to the system. What we have introduced here is a procedure where one needs to submit a job in order to monitor for an event.

Consider as an example how to implement an event test which triggers when a particular file is found. Lets 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 $oramsmsg(oraobject) 2 "$file_name found"
}
```

In order to receive this event, a user needs to register an event with the Unsolicited test selected and configured to filter in events with name `/myevents/node/files/filefound`. This event should be registered against a node and will trigger against it. The message attached to the event occurrence will contain the values of all parameters passed into `orareporevent`.

Although this event is fairly straightforward, there are two things wrong::

1. The event does not ever clear – it should, if the file disappears

2. The event will trigger every time the above script is evaluated, even if it has triggered the time before – you will receive the same alert multiple (possibly infinite) times. 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 preferable 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 oramsq  
    append event_lock [tempdir] "/" $oramsq(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 } {
    orareportevent $event_name $oramsg(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/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 in 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 section contains an example of unsolicited event tests. In this case, the test evaluation involves connecting to an Oracle instance and performing 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 $oramsg(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. Notice that this script allows for an optional overwrite of the preferred credentials via an optional forth input argument.

Unsolicited Event Caveats

- To monitor for an unsolicited event against multiple targets on the same node (say the event above against multiple databases on the same node), you need to register for the unsolicited event multiple times so that you may have the flexibility of proper filtering.
- Unsolicited Events cannot be used with Fixit Jobs.

Event Pane

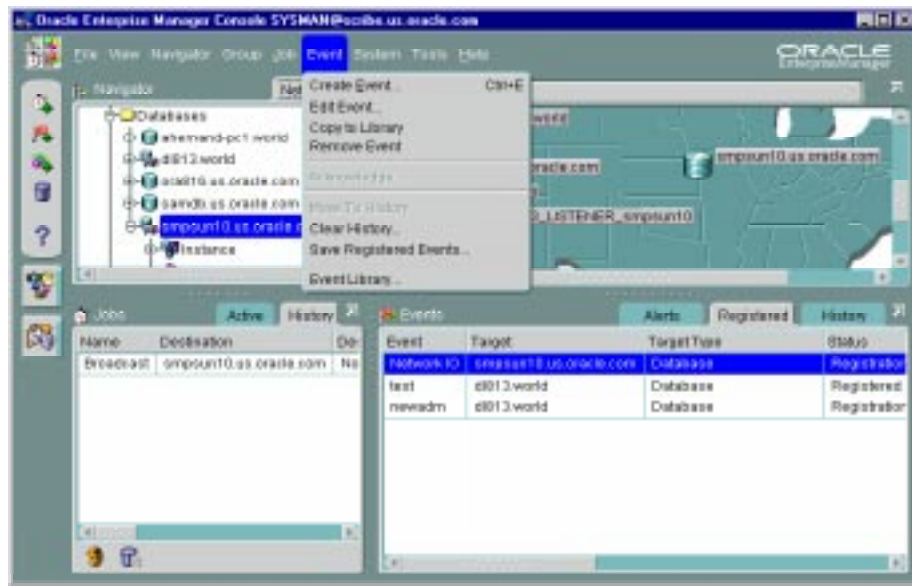
The Event pane contains the following pages:

- 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 5-1, "Event Menu and Pane" for an illustration of the Event pane.

The Event pane can be hidden or shown by selecting Show/Hide Event Display in the Console View menu. You can also hide or show the pane by clicking on the flag icon in the Console toolbar.

Figure 5–1 Event Menu and Pane



Alerts Page

The Alerts page displays event tests that have been triggered.

Severity

Severity of the event occurrence: critical (red), warning (yellow), clear (green), or unknown (gray).

Event

Name of the event.

Target

Target where the event was triggered.

Time/Date

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 to display the Event Viewer property sheet. 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 5-22 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 History icon located at the bottom left of the event History page.

To clear all entries in the History page, click the Clear History icon located next to the Refresh History icon.

Registered Page

The Registered page displays the events that have been registered, or submitted, to monitor test conditions on any network objects. The Registered page contains the following information:

Event

Name of the event.

Target

Target where the event was monitored.

Target Type

Type of event destination: database, node, listener

Status

Current registration status of the event: Registered, Registration Pending, De-Registration Pending, and Registration Failed.

Owner

Administrator who owns the event.

Note: Under certain circumstances, an event will remain in a Registration Pending state. There are two probable causes:

1. 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.
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 depending on the items selected in the Event pane. See Figure 5-1, "Event Menu and Pane" 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 "Creating, Modifying, and Registering an Event" on page 5-24 for more information.

Edit Event Occurrence

Displays the definition of an existing event. See "Creating, Modifying, and Registering an Event" on page 5-24 for more information.

Acknowledge

Acknowledges the selected event in the Alerts page.

Copy to Library

Copy the selected event in the Event pane to the Event Library.

Remove Event

Removes the selected event from the Event pane. 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.

Clear History

Clears the contents of the Event History page to a file.

Save History

Saves the contents of the Event History page to a text file.

Event Library

Displays the Event Library dialog. See "Event Library Dialog" on page 5-21 for more information.

Register Event

Registers, or submits, the selected event to a destination. This action monitors the destination, or network object, with the event tests in the event. This menu option is used to retry a registration that has failed.

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.

Description

Description of the event.

Owner

Administrator who created the event.

Editing an Event

Select an event and click Edit to display the property sheet for the event. The property sheet allows you to view and modify the events. Editing most event parameters can only be done via the Event Library dialog box.

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. You can add these tests to an event. The tests include:

- Oracle DB UpDown: checks whether a database is up or down.
- Oracle Host UpDown: checks whether a node is up or down.
- Oracle Listener UpDown: checks whether a listener 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.

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 and "Oracle Event Tests" on page 5-34 for information on Oracle events and their parameters.

Event Viewer

The Event Viewer property sheet displays details on a selected event in the History or Alerts page. To view details of an event, double-click on an event to display the property sheet. You can enter optional comments in the Log page. Once cleared, events are automatically moved to the History page. The pages of the Event Viewer include:

- General
- Log

- Notification Details

General

The Event Viewer General page displays statistics and owner information on a selected event. You can display the property sheet for the event with the Show Event Definition option. The following statistics are displayed on the General page:

Target

Destination of the event.

Target Type

Destination type, such as database, listener, or node.

Last Updated

Time of last update.

Owner

Administrator that created the event.

Assigned To

List of administrators that are receiving notifications for this event.

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.

Log

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.

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.

Time/Date

Time and date when the comment was entered.

Notification Details

The Event Viewer Notification Details displays details of email and paging notifications sent for a selected event. The information displayed in Details page includes:

Administrator

Administrator that was notified.

Time/Date

Time and date of the notification.

Method

Email address or pager number.

Status

Status of the notification, indicating whether the notification was sent, is pending, or has failed.

Creating, Modifying, and Registering an Event

Events include the destination type and the event information that you want to monitor. Events can consist of multiple event tests. To create or modify 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 or modify the fields in the General, Events, Parameter, and Permissions pages of the property sheet to create a new event. If you modify an event that has been registered, those changes are not used by the registered event.

Note: Modification of event parameters can only be carried out by editing the event from the Event Library dialog.

3. When you have completed the Event property sheet:
 - a. Choose Register, or submit, to send the event to the selected destinations and select the OK button at the bottom of the property sheet. The new event is not saved and does not appear in the Event Library dialog.
 - b. Choose Add to the Library to save the event but not submit it to the destinations at this time. Select the OK button at the bottom of the property sheet. The new event appears in the Event Library dialog.
 - c. Choose Register and Add to the Library to submit the event to the selected destinations and save the event. Select the OK button at the bottom of the property sheet. The new event appears in the Event Library dialog.

If you registered an event, the Intelligent Agent for a destination processes the event and the event appears in the Registered page of the Event pane. Each destination service is listed separately with the event.

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. This notification changes the color of the severity flag for the event in the Alerts page. If the destination database icon is displayed in the Group pane, the flag on the icon changes color. The colors and their meaning are:

- Unknown (gray)
- Event cleared (green)
- Warning (yellow)
- Alert (red)

Cases where an event notification is Unknown (gray flag) indicate the target or node where the event is registered is unavailable.

Warning: Do not register an UpDown event (included in the Oracle DB Fault event) against the database or node where the repository table is stored. The communication between the Intelligent Agent and Management Server can interfere with the repository database connection. If the repository database or node is down, the event cannot notify the Management Server.

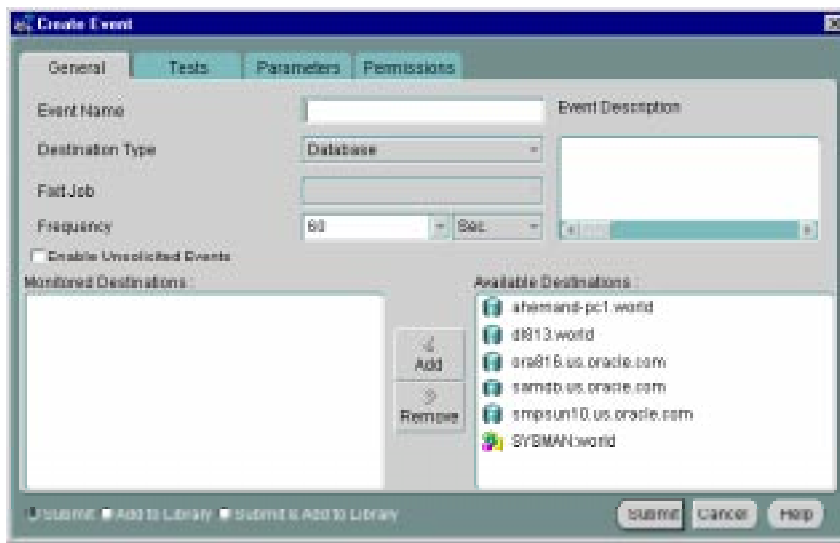
The property sheet for creating a new event is the same as the property sheet for modifying an event, except that the event name field is read-only. See Figure 5-2, "Event General Page" for an illustration of the Event property sheet.

You can also use events that have been installed with Oracle Enterprise Manager. See "Event Categories and Types" on page 5-6 for more information.

Event General Page

On the General page, you determine the set name, destination type, description, and whether this event should monitor third-party events.

Figure 5-2 Event General Page



Event Name:

Enter an event name.

Destination Type:

Select the destination 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.

Event Description:

Enter a description or comment for the event

Fixit Job

A fixit job is designed to correct a problem. 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 have been created with the Job system and have been designated as fixit jobs. The jobs must be submitted and running on the same destination that the event is set on. You must select destinations to make fixit jobs available in the pull-down list. See "General Page" on page 4-13 for more information.

Fixit job options are:

- To choose a fixit job, select an existing fixit job from the pull-down list.
- Select <None> if you do not want to use a fixit job.

To turn off a fixit job after an event has been registered, you must remove the event registration, create a new event, and register it.

Note: Each event must use a unique fixit job on each destination where the event is registered. Also, when a single Intelligent Agent is monitoring multiple databases at a destination, create a separate event and fixit job for each database.

Frequency

Determine the frequency that you want the test to run on the selected destinations. The frequency determines how often the event condition is checked. For example, if the frequency is set to 30 seconds, the event condition is checked every 30 seconds. To ensure timely notification, set the frequency to the smallest interval possible. To set the frequency:

- Enter a value in the frequency field or select a value from the pull-down list.
- Select the interval units from the pull-down list.

See "Event Notifications" on page 5-4 for details on the notification frequency.

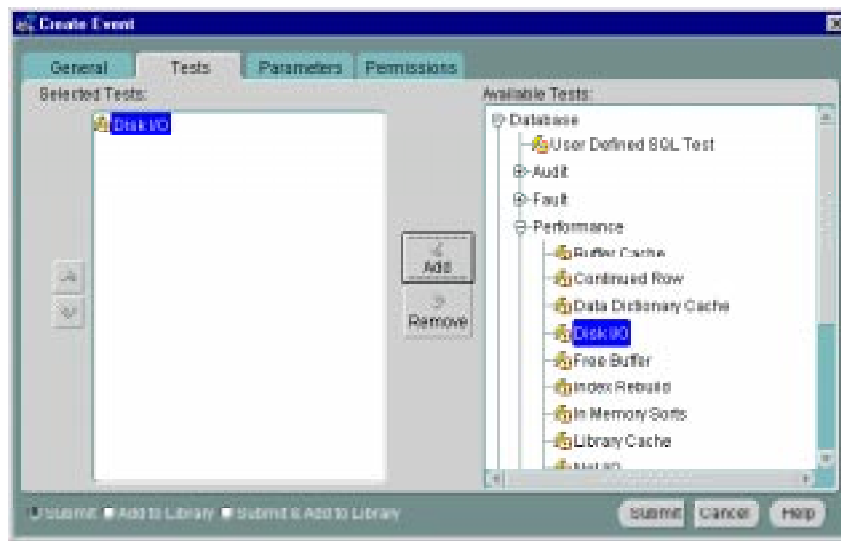
Enable Unsolicited Events

Check this box to monitor unsolicited events. See "Unsolicited Event Tests" on page 5-8 for more information.

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 5-3 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.

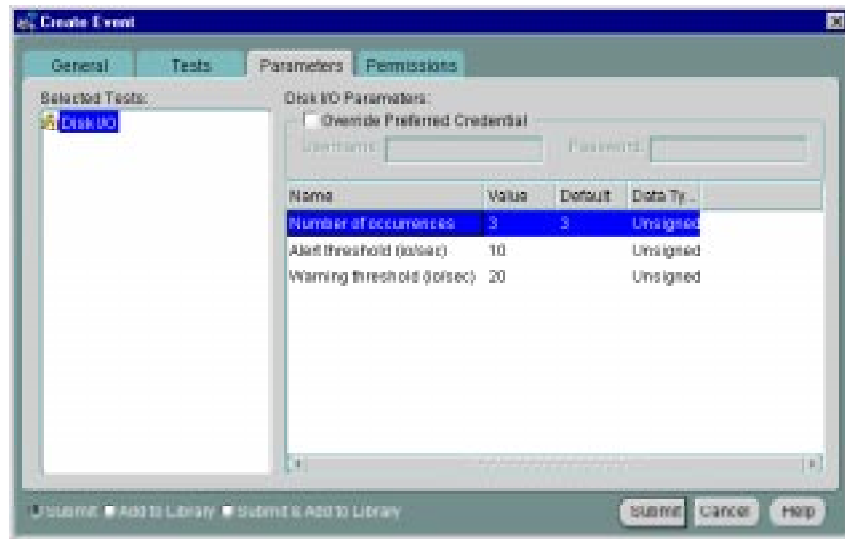
Selected Tests:

Select the event tests in the list you want to remove from this event, then click on the >> (Remove) button.

Event Parameters Page

The fixit job, frequency, and parameter settings are entered in the Event Information section of 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 5-34 for information on tests and their parameters.

Figure 5–4 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.
- Double-click on an item in a scrolling list to display the parameter dialog box. Enter a value or filter in 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 Permissions Page

Determine the permissions that you want to assign to the event with the Permissions Page. This allows other administrators to view or modify the event. Notifications can be assigned with this 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 5-22), and edit the event's properties except those reserved for Full permission.

Full

Allows the administrator to delete the event, modify permissions for other administrators, and change the ownership of the event.

Notify

Allows the administrator to enhance 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 "Permissions" on page 1-16 for more information.

Enable SNMP Notifications

When checked, permits external notification (SNMP traps) for third-party event tests at the destination where the Intelligent Agent is located.

Show Notification Schedule

Show Notification Schedule displays the notification schedule for the event. You can view the administrators that are notified by moving the cursor over any time block in the grid. 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.

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, or Registration Pending. In addition, the destination 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 destinations, you can view the status of this event for those destinations by selecting the desired target from the Destination pull-down list. The status of the event displays for that destination. To view the status of this event for all destinations simultaneously, select <All>.

The following options are available on the Progress page:

Destination (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, or Registration Failed.

Destination

Network destination for the event.

Time/Date

Time and date the event was submitted for registration.

Show Output

Displays the Event Status Message dialog. This button is active only when event registration has failed, thus generating an error message.

Save List

Saves the contents of the list to a text file.

Example: Creating and Registering an Event

In the following example, a new event for monitoring extents in a database is created and registered at several databases. When creating an event, the General, Tests, Parameters, and Permissions Pages are completed to define the event. You need to register, or submit, an event to run event tests on specific destinations in the network environment.

1. To create a new event, select Create Event from the Event Menu to display the General Page of the Event property sheet.
2. Enter an event name in the Event Name field, such as MyEventSet.
3. Enter a brief description of and comments about the event in the Description field. The event in this example will monitor Maximum Extents. This event is only available with the optional Performance Pack.
4. Select Database from the Destination Type pull-down list.
5. Do not change the frequency that the event monitors the destination. The default frequency interval is acceptable.
6. Do not select a fixit job (if any are available). These are jobs that have been created with the Job system and have been designated as fixit jobs.
7. Select a database from the Available Destinations list, then click on the << (Add) button. The database is moved to the Selected Destinations list and will be monitored for the events in the event.
8. Repeat the previous step for several additional databases.
9. Click on the Tests page tab to display the Tests page.
10. Select the Maximum Extents event test. This test checks the number of available extents for table, index, cluster, or rollback segments. Because this test accesses system tables, you must set up Administrator Preferred Credentials with an administrator that has system privileges. See "Preferred Credentials" on page 1-18 for more information.

Note: The Maximum Extents event test is only available with the optional Diagnostics Pack. The standard Up/Down event tests provided with Enterprise Manager do not require parameter settings.

11. Click on the << (Add) button. The event test is moved to the Selected Tests list and will be added to the event.
12. Click on the Parameters page tab to display that page. If you had selected several tests, you would need to select a specific test in the Selected Tests list to display the Parameters for the event test.
13. Click on Warning threshold in the scrolling list, to change the value of the parameter. Enter 3 in the text field.
14. Click on Alert threshold in the scrolling list, to change the value of the parameter. Enter 2 in the text field.
15. Do not change the other parameters. Note that the * in the Tablespace name, Segment name, and Segment type parameter fields signifies all existing values. To monitor a specific tablespace, such as USERS, you would enter = 'USERS' in the Tablespace name field.
16. Click on the Permissions page tab to display that page. Assign permissions to all administrators to allow them to view the event and receive notifications.
17. After you have completed the event property sheet, select the Submit and Add to the Library button at the bottom of the property sheet to submit the event to the selected destinations.
18. Click the OK button at the bottom of the property sheet to exit. The new event is saved and appears in the Event Library dialog.
19. If you want to modify the event, you can edit this event in Event Library dialog. In order for new values to affect event monitoring, you need to re-register the events.

When an event is submitted, each destination is validated and the Intelligent Agent for a destination processes the event. If the registration is successful, the event appears in the Registered page of the Event pane. Each destination database is listed separately with the event.

If an event condition occurs or threshold values are exceeded for an event test, a notification is sent to the Alerts page of the Console Event pane. If administrators have been selected to be notified by email or paging, an email or page is sent. The event notification changes the color of the severity flag for the event in the Alerts page. If the destination database icon is displayed in the Group pane, the flag on the icon changes color. The colors and their meaning are:

- Unknown (gray)
- Cleared (green)

- Warning (yellow)
- Alert (red)

After an event condition is fixed, the event is cleared automatically. You can also clear an event by acknowledging and moving the event to the Event History page.

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 1-7. For information on configuring and email and paging services, see the *Oracle Enterprise Manager Configuration Guide*.

Oracle Event Tests

This section lists the Event system event tests with their parameters and return values. See "Event Parameters Page" on page 5-29 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 5-36 for more information.

Event tests are specified for database, listener, 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.

The event test scripts are written in the Tool Command Language (Tcl) enhanced with Oracle Tcl commands (OraTcl). For information on using Tcl and OraTcl to write event scripts, see *Oracle Enterprise Manager Application Developer's Guide*.

Some of the event tests require special tables in the database. For example, the `catblock.sql` script needs to be run to use the User Blocks event test. See the *Oracle Enterprise Manager Configuration Guide* for scripts that need to be executed against the database that will be monitored.

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 "Preferred Credentials" on page 1-18 for more information.

Database Fault Management Event Tests

This category of event test tracks severe problems that require immediate attention.

UpDown (Database)

This event test checks whether the database being monitored is running. If this event test is triggered, other database events are not ignored.

Parameters

none

Action

The Startup Database job task can be set up as a fixit job for automatically correcting the problem.

Note: If the listener serving a database is down, this event test may be triggered because the Intelligent Agent uses the listener to communicate with the database.

Node Fault Management Events

This category of events monitors for catastrophic conditions on the system. Immediate action needs to be taken by the administrator.

Note: If an Up/Down event test (Database Up/down, Listener Up/Down) has been registered on a node, then the UpDown Node event test is implicitly registered and is triggered when the Intelligent Agent is down, or when the Management Server loses contact with the node.

UpDown (Node)

This event checks whether the Intelligent Agent on a node can be accessed from the Console via periodic "heartbeating" between the Management Server and the Intelligent Agent. Since this event does not actually register with the Intelligent Agent, no fixit jobs can be associated with it. If the defaults are accepted, this event will trigger within five minutes of the Intelligent Agent becoming inaccessible and will clear within three minutes of the Intelligent Agent becoming accessible again.

Parameters

none

Listener Fault Management Events

This category of events monitors for catastrophic conditions on the system. Immediate action needs to be taken by the administrator.

SQL*Net UpDown (Listener)

This event checks whether the listener on the node being monitored is available.

Parameters

none

Action

The Startup Listener job task can be set up as a fixit job for automatically correcting the problem.

Data Gatherer Fault Management Events

The Oracle Data Gatherer Alert and UpDown events are both Fault Management events for the node destination type. This category of events monitors for catastrophic conditions on the system. Immediate action needs to be taken by the administrator.

UpDown (Data Gatherer)

This event checks whether the Intelligent Agent data gathering service on a node can be accessed from the Console. If the Intelligent Agent data gathering service is down, this event is triggered.

Parameters

none

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.

The event status and page 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 Alert

800 = Event Node Down

900 = Event Unknown

1000 = Event Assignee Changed

If JJJ= job id then,

JJJ100 => job started for job# JJJ

JJJ200 => job completed for job# JJJ

JJJ300 => job failed for job# JJJ

For example, if 001100 is displayed on your numeric pager, then you know that job id 001 just started. If 001300 appears, then you know job id 001 just failed.

If EE= event ID then,

EE500 => event cleared for event # EE

For example, if 7500 is displayed, you know that event number 7 had a status change and the status is now cleared.

General Event System Troubleshooting Procedures

The following procedures will assist you and/or Oracle Support if problems are encountered with the Event subsystem.

Obtaining Enterprise Manager Repository Status

If you need to contact Oracle Support, you must provide specific information on the status of your Oracle Enterprise Manager repository. To assist you, a SQL script has

been provided that extracts pertinent Repository data and generates a debug log file.

To generate the log file, perform the following procedure:

1. Save `oms.log.*` for every Management Server connected to the repository. The log files are, by default, located in
`$ORACLE_HOME/sysman/log.`
2. Log on as the repository user and connect to the repository using SQL*Plus. For example:
`connect <repository name>/<repository password>`
3. Run the following SQL script:
`@$ORACLE_HOME/sysman/admin/evtdbg.sql`
4. This SQL script will generate a file called `evtdbg.log` in the directory where Server Manager or SQL*Plus is invoked.
5. Submit the Management Server log files along with `evtdbg.log` to Oracle Support.

Note: The size of the `evtdbg.log` file can be quite large.

Removing Old Job and Events

If you submit jobs or events to a node running an Intelligent Agent, and the Intelligent Agent unexpectedly goes out of service and is not restarted, you will be unable to delete active jobs and events from the Oracle Enterprise Manager repository: Jobs and events will remain in a "pending deletion" state. Consequently, the node itself cannot be removed.

To remove these "stale" entries from the Oracle Enterprise Manager repository:

1. Shutdown the Management Server.
`$ oemctrl stop oms <username>/<password>`
2. Logon to the database that hosts the repository as a valid user
`$ sqlplus <user_name>/<user_pwd>@<repository_database>`
You may alternatively use Server Manager.

3. Execute the following SQL statement:

```
execute smp_maintenance.delete_operations_on_node('<target  
nodename>');
```

The node name should be identical to what appears in the Console Navigator (it may include the domain name as well).

4. Log out of the repository.
5. Restart your Management Server.

```
$ oemctrl start oms
```

This will remove all currently active jobs and events against the specified node. No historical data will be affected. The node itself will not be deleted. To remove the node, log on to the Management Server through the Console and remove the node from the Navigator.

Management Server and Intelligent Agent

Enterprise Manager, Release 2.1 deploys a Management Server and an Oracle Intelligent Agent to allow clients to manage objects, such as databases, listeners, nodes, and additional discovered services.

- Oracle Intelligent Agent
- Management Server

Oracle Intelligent Agent

The Oracle Intelligent Agent is a daemon-like process running on a remote node in the network. An agent resides on the same node as the service it supports.

However, the agent can support more than one service on a particular node. For example, if two databases are installed on one machine, a single agent can support both databases. The agents are responsible for:

- Providing local services or calling operating system dependent services to interact locally with the managed targets.
- Running jobs, collecting their results and output, and queueing them for the Management Server.
- Checking for events, and queueing the resulting event reports for the Management Server.
- Returning job and event reports to the appropriate Management Server.
- Canceling jobs or events as directed by the Console or other applications.
- Handling Simple Network Management Protocol (SNMP) requests, if SNMP is supported on the agent's platform.

For information on configuring the agent, see the *Oracle Enterprise Manager Configuration Guide* and the Oracle server platform-specific installation documentation for your system. For information on general Intelligent Agent operation, troubleshooting and Tcl scripting, see the *Oracle Intelligent Agent User's Guide*.

Characteristics

Intelligent Agents are autonomous because they function without requiring that the Console or Management Server be running. An agent that services a database can run when the database is down, allowing the agent to start up or shut down the database. The Intelligent Agents can independently perform administrative job tasks at any time, without active participation by the administrator. Similarly, the agents can autonomously detect and react to events, allowing them to monitor the system and execute a fixit job to correct problems without the intervention of the administrator.

The agents operate independently and are able to execute jobs and monitor events when the administrator has logged out of the Console and/or the Management Server has been shut down. The agents queue any job or event messages destined for that administrator, and deliver them when the administrator logs in to a Console again. Information about jobs and events are stored in files on the agent's node.

Note: The agent queues a maximum of 500 messages. After the limit is reached, the oldest messages are dropped.

Event and Job Support

The agents are responsible for executing jobs and monitoring for events. Jobs and events are implemented as Tcl scripts. When the agent executes a job or tests for an event, it runs the appropriate Tcl script. Because jobs can be long-running or complicated tasks, such as a database backup job, the agent does not execute the job in its process space. Jobs are run in a separate process. When the job is completed, the agent sends the results to the Management Server. In contrast, event scripts are typically run directly by the agent. Event scripts are used for detecting exceptions and are expected to have short execution times.

Network Encryption

Network encryption for communication between the Management Server and Agent can be accomplished with the Advanced Networking Option (ANO) option of Net8. ANO uses a sophisticated algorithm to provide encryption on the Transparent Network Substrate (TNS) connections. For information on ANO, see the *Net8 Administrator's Guide*. For SQL*Net installations, network encryption can be accomplished with the Oracle Secure Network Services (SNS) option.

SNMP Support

The agents support SNMP, so applications can communicate directly with the agent using SNMP protocol. The agents provide access to Oracle's database Management Information Base (MIB) variables. Although the agent supports SNMP, the Management Server does not use that protocol to communicate with the agent. You can submit jobs or events that access Oracle MIB variables even when the database resides on a platform that does not support SNMP. For more information on SNMP, see the *Oracle SNMP Support Reference Guide*.

Management Server

The Management Server communicates with first tier clients, then processes all system management tasks, administrates the distribution of these tasks across the entire network system to the Agents, and provides scalability across the system. The Management Server uses a repository as persistent storage to maintain system data, application data, and state for many entities distributed throughout the architecture. The Management Server is responsible for such activities as:

- Communicating with agents on nodes to schedule and manage jobs.
- Communicating with agents on nodes to submit and monitor events.
- Queueing operations that fail and retrying them periodically.
- Discovering the network services in the system.

The Management Server is implemented as a multi-threaded process. For example, separate threads in the Management Server perform activities such as submitting jobs and events to agents, discovering new services in the network, or receiving messages from agents. Because the Management Server's threads operate independently, the Management Server can perform various operations simultaneously and perform efficiently in large busy distributed environments.

Communication between the Management Server and the agents is vital to the Job and Event systems. The Management Server must be able to send messages to the agents in order to submit jobs and events. The agents must be able to send messages to the Management Server to report results and status messages for the jobs and events.

The Management Server uses CORBA to communicate with components of the Oracle Enterprise Manager system. When the Management Server communicates with V1.x Oracle Intelligent Agents, an Agent Gateway is used. The Agent Gateway and V1.x agents communicate using Oracle Remote Operations, which is a remote procedure call mechanism based on the Transparent Network Substrate (TNS). The

Management Server and agents can also use Oracle Secure Network Services (SNS) to maintain the security of their network transmissions.

Message Queues

The Management Server and agents use message queues for the messages they send. Using queues ensures that no messages are lost even when the Management Server or agent is down. The Management Server maintains several queues for messages. The operations queue contains job and event requests sent by the Console. For example, when you submit a new job to the Job system, the Console queues the new job request on the Management Server's operations queue.

Failed Queue

When the Management Server retrieves a job or event request from its operation queue, it tries to contact the agent that should receive the request. If it cannot contact the agent, the Management Server places the request in its failed queue. Periodically the Management Server tries to contact the agent which is responsible for the operation request in the failed queue. If the Management Server successfully contacts the agent, the operation request is removed from the failed queue and sent to the responsible agent.

Notification Queue

The Management Server maintains a notification queue for job and event notification. The notification queue contains messages about the status of jobs and events. When the Management Server receives a message from an agent regarding a job or event, it places the message in the notification queue. When the Management Server changes the status of a job or event it also places a message in the queue.

For example, when the Management Server has successfully submitted a new job to an agent, it places a message in the notification queue updating the job's status to submitted. Messages in the notification queue are used to update the job and event information stored in the repository.

Connection Cache

In order for an Management Server and an agent to pass messages, they must establish a connection. Rather than requiring that the Management Server or agent create a new connection each time it wants to send a message, the Management Server maintains a cache of connections. If a connection is needed and it already exists in the cache, it can be reused. This reduces the overhead involved in creating

new connections. Connections in the cache are aged out using a least recently used algorithm.

Part II

The Database Administration Tools

- Chapter 7, "DBA Management Pack"
- Chapter 8, "Managing Backup and Recovery"

DBA Management Pack

The primary administrative components of Oracle Enterprise Manager are contained in the DBA Management Pack, which is a suite of database administration tools and wizards bundled with Oracle Enterprise Manager. This functionality can be accessed with the new DBA Studio or launched separately from the Oracle Enterprise Manager Console or operating system. DBA Management Pack features can be used to perform most of your database administration tasks, and support all versions of Oracle databases.

DBA Studio or individual DBA Management Pack tools can be launched by either connecting to the Oracle Management Server using Oracle Enterprise Manager's three-tier framework, or by connecting directly to a database.

DBA Management Pack includes the following functionality:

- Instance Management
- Schema Management
- Security Management
- Storage Management
- SQL*Plus Worksheet

With this functionality you can:

- Complete database administration operations quickly and conveniently using a mouse rather than manually entering SQL commands
- Concurrently perform multiple tasks in multiple Oracle Enterprise Manager windows
- Administer multiple databases simultaneously

- Perform database administration tasks from a central location for both local and remote databases running on any Oracle platform in any location worldwide

DBA Management Pack also includes wizards for assisting with such tasks as importing, exporting, loading, backing up, and recovering data.

This chapter contains the following sections:

Topic	See Page
DBA Studio	7-2
Common Features of DBA Management Pack Functionality	7-2
DBA Management Pack Functionality and Wizards	7-4

DBA Studio

DBA Management Pack features can now be accessed from the new DBA Studio, as well as from the Oracle Enterprise Manager Console or operating system. DBA Studio is an interface from which all DBA Management Pack features can be accessed (with the exception of SQL*Plus Worksheet) and can either be run with a connection to an Oracle Management Server, which utilizes Oracle Enterprise Manager's three-tier framework, or connected directly to a database.

With DBA Studio, you can manage database security, schemas, instances, and storage from a unified tree view without launching separate features. This integrated access to DBA functionality simplifies switching between DBA tasks and provides an accurate view of database configuration status.

DBA Studio can be launched from the operating system or from the Enterprise Manager Console.

Common Features of DBA Management Pack Functionality

This section discusses the common features shared by the different tools in DBA Management Pack.

Tree Views and Property Sheets

Most DBA tools display a tree view of connected databases, which can be expanded to show subordinate objects. When you select a database or 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

When you select a folder in a tree list, a multi-column list of all the folder's objects appears on the right side of the window, providing a quick summary of information about each object in the selected folder.

Database Version Awareness

All DBA Management Pack tools 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.

Show SQL

Though one of the benefits of DBA Management Pack 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, 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 DBA Management Pack, 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.

Advanced Mode

DBA Management Pack features can be run in Advanced Mode, which displays additional tabbed pages in property sheets from which a tool's advanced database administration features can be accessed.

DBA Management Pack Functionality and Wizards

This section describes DBA Management Pack's features and wizards.

DBA Studio Tool	Task	See Page
Instance Management	Manages instances and sessions	7-4
Schema Management	Manages schema objects	7-8
Security Management	Manages security parameters	7-10
Storage Management	Manages database storage	7-13
SQL*Plus Worksheet	Executes SQL and PL/SQL commands	7-15
Wizards	Assist with importing, exporting, loading, backing up, and recovering data, as well as analyzing and creating tables and views	7-15

Instance Management

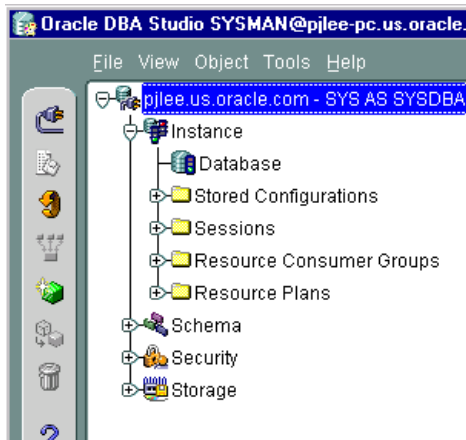
Instance Management functionality helps you manage database instances and sessions in your Oracle environment. With Instance Management features you can:

- Start up and shut down a database
- View and edit the values of initialization parameters
- View current session SQL and the explain plan
- Manage stored configurations (when connected to the Oracle Management Server)
- Manage resource allocations (with Oracle8i)
- Manage users' sessions
- Monitor long running operations (with Oracle8i)

When you select the Instance icon under the database in the tree view, the following list of objects and folders appears:

- Database
- Stored Configurations (only when connected to Oracle Management Server)
- Sessions
- Resource Consumer Groups (with Oracle8i)
- Resource Plans (with Oracle8i)

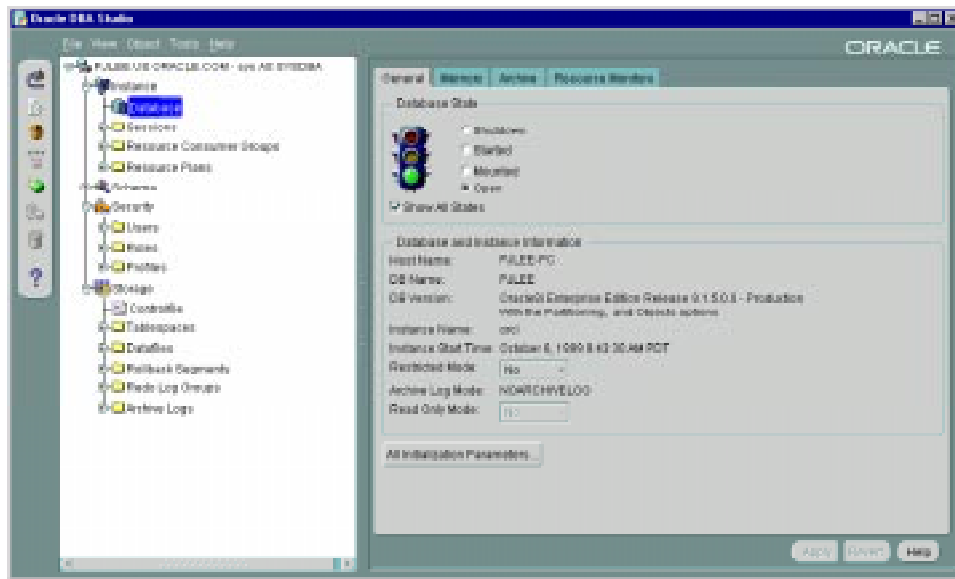
Figure 7-1 Instance Management Tree



Database Operations

When you select the Database icon under Instance, a property sheet of tabbed pages appears on the right for viewing information about the database instance and editing database properties. In this property sheet, the following information can be viewed and/or edited:

- Status of the database, including the database version and any installed options
- Memory usage of the current database
- Current state of redo log archival
- Resource management (with Oracle8i)

Figure 7-2 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.

Sessions

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.

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

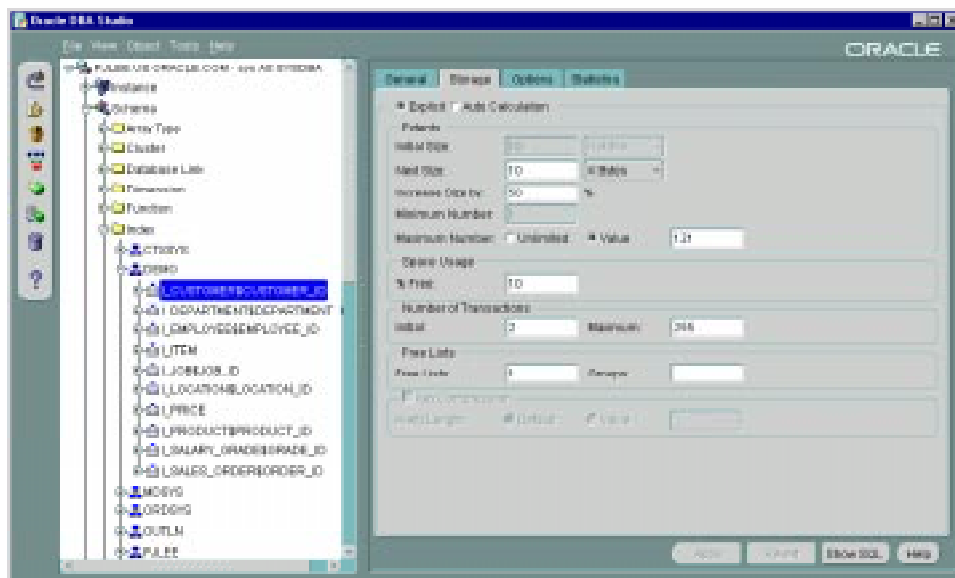
The Resource Plans folder lists objects that represent resource plans, which are ways of allocating resources among consumer groups. Resource plans contain directives that specify the resources to be given to each group and can be specified in hierarchical fashion using subplans.

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, and then activate the plan. You can also view the degree of parallelism for each consumer group.

Schema Management

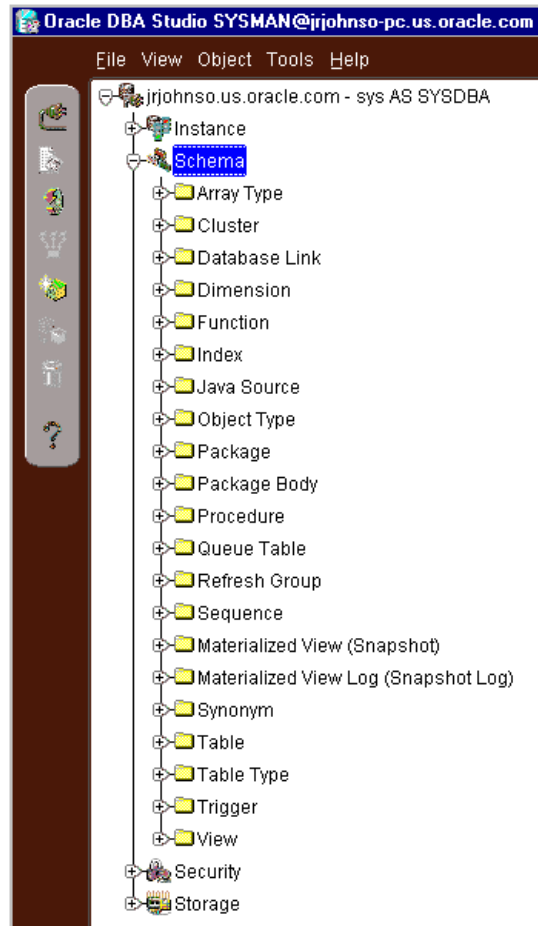
With DBA Management Pack's Schema Management functionality, you can create, alter, or drop database schema objects such as clusters, indexes, snapshots, tables, and views, as well as view dependencies of schema objects. Oracle Schema Management features also support Oracle8i objects, which include Index Organized Tables, partitioned tables and indexes, advanced queues, and Java stored procedures.

Figure 7-3 Schema Management Window



When you select the Schema icon under the database in the tree view, a list of schema objects appears.

Figure 7-4 Schema Objects



Tree List by Schema

Databases contain at least one named schema for each database user. Each schema object belongs to one of these named schemas, regardless of object type. Several objects belonging to the same schema can be edited by opening the View menu and

selecting By Schema. The tree view then reorders all schema objects within a set of schema folders listed alphabetically by name.

Editing an Object

To edit an object, such as a cluster, you use the cluster's property sheet which appears when you select the cluster folder in the tree view. You can then modify the cluster's parameters such as its name, type, and key column attributes.

Schema Management features also include the Table Data Editor, which allows you to view, update, and delete the contents of a table by selecting a right-mouse command on a table in the tree view.

Creating Objects

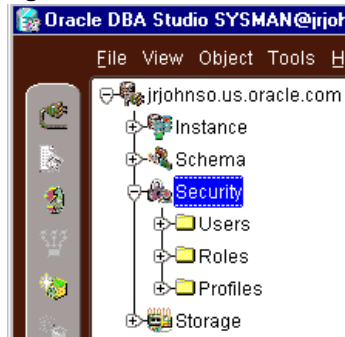
Schema Management allows you to create an object or a clone of an object by selecting Create or Create Like from the Object menu. When creating a clone of an object, all attributes are identical except for the name. Parameters for new objects and cloned objects are specified in property sheets which appear when you select Create or Create Like from the Object menu.

Schema Management features also include a Table Wizard which walks you through each step of creating a table, providing defaults for most options.

Security Management

In a large network environment, security parameters for objects, administrators, and users are in constant change. With DBA Management Pack's Security Management features, an administrator can make these necessary changes quickly and efficiently.

When you select the Security icon under the database in the tree view, folders for users, roles, and profiles appear.

Figure 7-5 Managing Users, Roles, and Profiles**Figure 7-6 Security Management Window**

 A screenshot of the Security Management window in Oracle DBA Studio. The window title is "Oracle DBA Studio". The left pane shows the tree view with "Security" selected. The main pane displays a table of users with the following columns: Username, Account Status, Expiry Date, Default Tablespace, and Temporary Tablespace. The table contains the following data:

Username	Account Status	Expiry Date	Default Tablespace	Temporary Tablespace
ALFORD@OPENAUTHENTICATED	OPEN	NONE	SYSTEM	SYSTEM
OTINGYS	OPEN	NONE	SYSTEM	SYSTEM
demo	OPEN	NONE	SYSTEM	SYSTEM
DEMO	OPEN	NONE	SYSTEM	SYSTEM
MOEYS	OPEN	NONE	SYSTEM	SYSTEM
MYOYS	OPEN	NONE	SYSTEM	SYSTEM
ORDWUSER	OPEN	NONE	SYSTEM	SYSTEM
ORDWUS	OPEN	NONE	SYSTEM	SYSTEM
ORUM	OPEN	NONE	SYSTEM	SYSTEM
PUDE	OPEN	NONE	DEM_REPOSITORY	TEMP
SCOTT	OPEN	NONE	USERS	TEMP
SYS	OPEN	NONE	SYSTEM	SYSTEM
SYSTEM	OPEN	NONE	USERS	TEMP

User Operations

Security Management features help you manage the database users in your network by helping you create users and clones of users, add and remove user permissions and roles, and alter user properties, including account status and default profiles.

Security Management capabilities also allow you to easily see users' dependents and dependencies.

Role Operations

With Security Management role operations features, 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.

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.

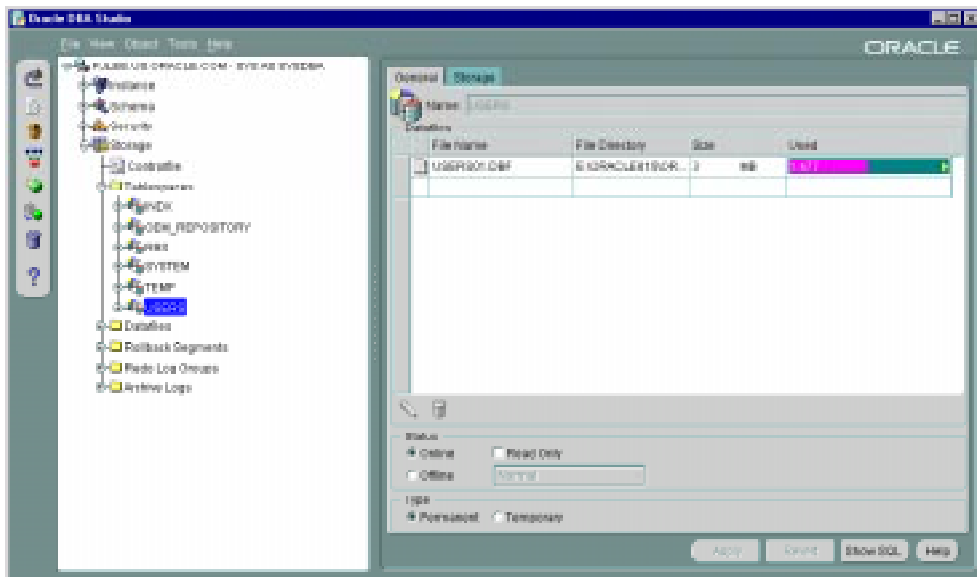
Profiles, roles, and the users to which they are assigned can easily be seen in security lists provided by Security Management. Administrators can then use property sheets to determine security parameters, simplifying the process of making changes.

Security Management features also support Oracle8 password management, which increases system security. Supported features include: account locking, password lifetime and expiration, password history, password complexity, verifications, and export/import of passwords.

Storage Management


DBA Management Pack's Storage Management features help you administer tablespaces, datafiles, redo logs, and rollback segments for optimum database storage.

Figure 7-7 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 7–8 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, 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.

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.

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.

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.

SQL*Plus Worksheet can be accessed from the Enterprise Manager Console.

Note: Additional information on all tools and wizards in DBA Management Pack can be found in the Oracle Enterprise online Help.

Wizards

For help with database administration tasks such as importing, exporting, loading, backing up, and recovering data, the DBA Management Pack offers Backup Management tools and Data Management tools. These wizards increase the ease-of-use of the DBA applications.

The wizards can be launched from the Schema Management or Storage Management windows through the Oracle Enterprise Manager Console, or from DBA Studio if it is connected to the Oracle Management Server.

Backup Management Tools

Oracle Enterprise Manager backup management tools include the Backup Wizard and the Recovery Wizard, which assist with backing up and recovering Oracle8 databases. These wizards use the Oracle8 Recovery Manager (RMAN) subsystem to perform tasks.

The Backup Wizard is used for backing up one or more databases, tablespaces, datafiles, or archivelogs. The wizard prompts you to specify the time and frequency to run the backup, as well as the backup configuration and target database(s), then sends your specifications to the Job Scheduling system for execution.

The Recovery Wizard is used for restoring and recovering a database, or one or more tablespaces or datafiles. The wizard prompts you to select the objects you need to recover and the backup configuration to use, then sends your specifications to the Job Scheduling system for execution. You can also rename datafiles if required.

Data Management Tools

Data Management tools automate the transfer of data to and from Oracle databases. The Data Management tools consist of the Export Wizard, Import Wizard, and Load Wizard.

- The Export Wizard transfers data from an Oracle database to an Oracle binary operating system file.
- The Import Wizard transfers data from an exported binary file to an Oracle database.
- The Load Wizard transfers data from operating system files, such as text files, to an Oracle database.

All three wizards submit jobs to perform tasks. Jobs can be scheduled for a specific time or at recurring intervals. Once a job is submitted with one of the wizards, you can monitor its progress and view its properties in the console's Job pane, just like any other job.

Note: For more information on wizards, see the DBA Management Pack Quick Tour or the Oracle Enterprise Manager online Help.

Managing Backup and Recovery

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

The following topics are covered:

- Accessing the Backup Management Wizards and Property Sheets
- Backup Wizard
- Recovery Wizard
- Create Backup Configuration
- Catalog Maintenance Wizard
- Backup Configuration Library
- Configuring a Remote Database for Backup or SYSDBA Administration
- Recovery Catalog Setup and Registration for Oracle8i

The Oracle Enterprise Manager Backup, Recovery, and Catalog Maintenance wizards, Backup Configuration Library, and the Create Backup Configuration property sheets are only available for Oracle8 and above databases.

The wizards provide a graphical user interface to Recovery Manager, an Oracle utility you use to back up, restore, and recover database files.

For more information on backup and recovery strategies, see your *Oracle8i Backup and Recovery Guide*.

Accessing the Backup Management Wizards and Property Sheets

The Backup Management wizards and property sheets consist of:

- Backup
- Recovery
- Catalog Maintenance
- Create Backup Configuration
- Backup Configuration Library

You can access the Backup Management wizards and property sheets using one of the following methods:

Console

- From the Console Navigator tree, select the database, tablespace, or datafile you want to administer; then, choose the tool from the context-sensitive Backup Management menu.
- From the Console's Tools->Database Applications menu.

DBA Studio

- From DBA Studio's Tools menu, choose Backup Management.
- From DBA Studio, select the database, tablespace, or datafile you want to administer in the Navigator tree; then, choose the tool from the context-sensitive Backup Management menu.

Storage Manager

- From Storage Manager's Tools menu, select Backup Management
- From Storage Manager, select the database, tablespace, or datafile you want to administer in the Navigator tree; then, choose the tool from the context-sensitive Backup Management menu.

Note: The Backup Management wizards and property sheets are only available when you are connected to a Management Server.

Backup Wizard

The Backup Wizard provides an intuitive graphical user interface for the backup facility of Recovery Manager, an extremely powerful and flexible command line backup and recover utility designed for the Oracle8 and above databases. This wizard guides you through the entire backup process.

Before attempting to use the Backup Wizard, make sure the following requirements are met:

- The Oracle target database you want to backup is version 8 or above
- The application is connected to a Management Server
- The preferred credentials is SYSDBA or you have created a backup configuration and overridden the preferred credentials. If your target database is not configured for remote sysdba operation, see "Configuring a Remote Database for Backup or SYSDBA Administration" on page 8-31.
- The job and event systems are fully functional
- A service name entry exists for the target database in its tnsnames.ora file which matches the entry in the tnsnames.ora file in the Oracle Enterprise Manager home.

You can access the wizard through the Console and Storage Manager if you are connected to a Management Server.

For more information on backups and related strategies, see your *Oracle8i Backup and Recovery Guide*.

Backup Wizard: Introduction

The Backup wizard helps you perform a backup of the various objects like database, datafiles, tablespaces, and archivelogs. It guides you through the process of specifying the backup strategy and submits a backup job through the Enterprise Manager to complete the operation.

A checkbox "Display this page next time" is available where you can specify not to show this page once you are familiar with the Backup Wizard. By default, the checkbox is selected. If you deselect the checkbox, the wizard starts on the Strategy Choice page the next time.

Backup Wizard: Strategy Choice

The Strategy Choice page allows you to specify whether you want to use a predefined backup strategy or customize your own. There are two options from which to choose:

- **Predefined backup strategy:** Select the appropriate backup strategy. For backup frequency guidelines, refer to the *Oracle8i Backup and Recovery Guide*.
- **Customize backup strategy:** Select the information you want to backup and the schedule for the execution of the backup. Refer to the guidelines documented in the *Oracle8i Backup and Recovery Guide*.

After choosing your strategy choice, press Next to continue.

Proceed to Predefined Backup Strategy on page 8-5 if you have chosen Predefined Backup Strategy as your strategy choice.

Proceed to Customize Backup Strategy on page 8-8 if you have chosen Customize Backup Strategy as your strategy choice.

Predefined Backup Strategy

The following pages appear if you choose "Predefined backup strategy" as your strategy choice:

- Backup Wizard: Backup Frequency
- Backup Wizard: Backup Time
- Backup Wizard: Configuration
- Backup Wizard: Multiple Destination

Backup Wizard: Backup Frequency

The Backup Frequency page allows you to specify how often you want to perform a backup and on which day of the week.

Choose the type of description that best fits your database.

- My database is not frequently updated (DSS--Decision Support System)
 - Full backup once-a-week
 - Backup day:
Choose the backup day from the scroll-down list.
- My database is moderately updated (OLTP--On Line Transaction Processing) and is not very large
 - Full backup everyday
- My database is frequently updated and is large or medium
 - Full backup once-a-week and incremental everyday
 - Full backup day:
Choose the full backup day from the scroll-down list.

After choosing your backup frequency, press Next to continue. Proceed to Backup Wizard: Backup Time on page 8-6.

Backup Wizard: Backup Time

The Backup Time page allows you to specify the start time for performing the backup.

If your database is running in ARCHIVELOG mode, the backup is taken when the database is open. When you run a database in ARCHIVELOG mode, you allow Oracle to archive the online redo log. You can either perform the archiving manually or enable automatic archiving.

If the database is running in NOARCHIVELOG mode, the backup job closes the database, takes the backup, and then reopens the database. When you run your database in NOARCHIVELOG mode, you disable the archiving of the online redo log. The choice of whether to enable the archiving of filled groups of online redo log files depends on the availability and reliability requirements of the application running on the database.

Refer to the *Oracle8i Backup and Recovery Guide* for information on choosing between NOARCHIVELOG and ARCHIVELOG mode.

Refer to the Instance Manager online help for information on setting up backup modes.

After specifying your backup time, press Next to continue. Proceed to Backup Wizard: Configuration on page 8-6.

Backup Wizard: Configuration

The Configuration page allows you to choose the configuration for backup and view the details that you have set up previously using the Create Backup Configuration option.

Choose the configuration for this backup

Allows you to choose a configuration you have previously set up using the Create Backup Configuration Property Pages.

Description of the configuration

The description of the configuration.

Summary of the configuration:

- Name of the configuration
- Type of Destination (disk or tape)
- Channel Name (User-specified name of the channel)
- File Name (Unique backup set name)

- Drive and path where backup sets are stored.
- Recovery Catalog Credentials if you are using a recovery catalog

Refer to "Create Backup Configuration" on page 8-20 for information on creating a backup configuration.

After choosing your configuration, press Next to continue. Proceed to Backup Wizard: Multiple Destination on page 8-7.

Backup Wizard: Multiple Destination

The Multiple Destination page allows you to specify the destination database(s) to submit the job. You can select from the list of databases shown on the Available Databases window.

Move the selection(s) to and from the Available databases and Selected databases windows by using arrow buttons between them. By default, the database to which you are connected is shown in the Selected Databases window.

Note: To submit the same backup to multiple destinations, the databases must have the same structure and the same disk/tape configuration.

Backup Wizard: Summary

The Summary page appears when you click the Finish button. Information about the backup management job you just created using the Wizard is summarized before submission to the Oracle Enterprise Manager job system.

Clicking OK submits the job the Oracle Enterprise Manager job system.

Clicking Cancel returns you to the wizard.

Customize Backup Strategy

Customize backup strategy allows you to elect the information you want to backup and the schedule for the execution of the backup.

Refer to the guidelines documented in the *Oracle8i Backup and Recovery Guide*.

Backup Wizard: Backup Selection

Your database contains a wide variety of types of data. When developing your backup strategy, you must decide what information you want to backup. 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.

Note: If your database is in NOARCHIVELOG mode, you can only perform a full database backup. If the database is in NOARCHIVELOG mode and open at the time of the backup job, it is shut down during the backup, started in MOUNT mode, backed up, and then brought up once the backup is complete.

Choose from the following basic backup types:

Entire database

Entire database backup contains the control file along with all database files that belong to a database. Whole database backups are the most common type of backup. This option backs up the whole database at once. If you operate in ARCHIVELOG mode, you also have the option of backing up different parts of the database over a period of time, thereby constructing a whole database backup piece by piece.

Tablespace

A tablespace backup is a backup of a subset of the database. Tablespace backups are only valid if the database is operating in ARCHIVELOG mode.

Datafile

A datafile backup is a backup of a single datafile. Datafile backups, which are not as common as tablespace backups, are only valid if you run the database in ARCHIVELOG mode.

Archived redo log

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 destinations. Archive redo logs do not change, for example, but they are crucial for recovering your database, so you should maintain multiple copies if possible. This option is available only in ARCHIVELOG mode.

After choosing your backup type, press Next to continue.

Proceed to Backup Wizard: Archive Logs on page 8-10 if you have chosen Entire database or Archive Logs as your backup type.

Proceed to Backup Wizard: Tablespaces on page 8-9 if you have chosen Tablespace as your backup type.

Proceed to Backup Wizard: Datafiles on page 8-9 if you have chosen Datafiles as your backup type.

You can submit a backup job directly from the Job System, but the recommended method is through the wizard.

Note: If the database was OPEN and in ARCHIVELOG MODE when the backup job was created and the database is in MOUNT state when the job runs; then, the backup job will fail because of "Alter System Switch Logs."

Backup Wizard: Tablespaces

A tablespace backup is a backup of a subset of the database. Tablespace backups are generally valid if the database is operating in ARCHIVELOG mode.

The Tablespace page allows you to select the tablespaces you want to back up. Choose from the list of tablespaces shown on the Available Tablespaces window.

Move the selection(s) to and from the Available Tablespaces and Selected Tablespaces windows by using arrow buttons between them.

You may choose to include the control files with the backup.

After selecting the tablespaces you want to back up, press Next to continue.

Proceed to Backup Wizard: Archive Logs on page 8-10.

Backup Wizard: Datafiles

A datafile backup is the backup of a single datafile. Datafile backups, which are not as common as tablespace backups, are only valid if you run the database in ARCHIVELOG mode.

The Datafiles page allows you to select the datafiles you want to back up. Choose from the list of datafiles shown on the Available Datafiles window. Full path and name of the datafile are shown.

Move the selection(s) to and from the Available Datafiles and Selected Datafiles windows by using arrow buttons between them.

You may choose to include the control files with the backup.

After selecting the datafiles you want to back up, press Next to continue.

Proceed to Backup Wizard: Archive Logs on page 8-10.

Backup Wizard: Archive Logs

The Archive Logs page allows you to specify the following options:

- No (do not back up any archive logs)
- Yes, all archive logs
- Yes, selected archive logs (allows a range)

From

Date and time for the first archive logs to backup. These fields are enabled when you select "Yes, selected archive logs."

Until

Archive date and time for the last archive logs to backup. These fields are enabled when you select "Yes, selected archive logs."

delete archive logs after they have been backed up

Deleting archive logs saves space.

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 destinations. Archive redo logs do not change, for example, but they are crucial for recovering your database, so you should maintain multiple copies if possible.

After specifying the archive log options, press Next to continue.

Proceed to Backup Wizard: Backup Options on page 8-11 if you have chosen Entire database, Tablespace, or Datafile as your backup type.

Proceed to Backup Wizard: Configuration on page 8-12 if you have chosen Archive Logs as your backup type.

Backup Wizard: Backup Options

The Backup Options page allows you to choose from a full or an incremental backup.

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.

Incremental backup

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 back up only blocks that have changed since previous incremental backups. Blocks which have not changed will not be backed up.

- Non-cumulative backup

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 is longer.

- Cumulative backup

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.

Proceed to Backup Wizard: Configuration on page 8-12 if you have chosen your backup option.

Backup Wizard: Configuration

The Configuration page allows you to choose the configuration for backup and view the details that you have set up previously using the Create Backup Configuration option.

Choose the configuration for this backup

Allows you to choose a configuration you have previously set up using the Create Backup Configuration Property Pages.

Description of the configuration

The description of the configuration.

Summary of the configuration:

- Name of the configuration
- Type of Destination (disk or tape)
- Channel Name (User-specified name of the channel)
- File Name (Unique backup set name)
- Drive and path where backup sets are stored.
- Recovery Catalog Credentials if you are using a recovery catalog

After choosing your configuration, press Next to continue. Proceed to Backup Wizard: Schedule on page 8-12.

Backup Wizard: Schedule

The Schedule page enables you to schedule the execution of a backup.

Run Job

Select the frequency of backup. The choices are Immediately, Once, On Interval, On Day of Week, and On Date of Month.

- Immediately submits the task as soon as you finish the wizard. The task executes only one time.
- Once schedules the task only one time at the date and time you choose.
- On Interval enables 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 the scroll buttons. You can also type in a new value.
- On Day of Week enables you to schedule the task on one or multiple days (Sunday, Monday, etc.) of the week. Click the days of the week to select the days you want the task scheduled.

- **On Date of Month** enables you to schedule the task on one or multiple days (1 - 31) of the month. Click 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. Choose the first day of the month.

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 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 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 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 the scroll buttons to change the value. You can also type in new values.

Time Zone

Shows the time zone. The agent schedules the task execution at each destination based on the actual system time of each agent. Tasks are not necessarily run simultaneously.

Run on Multiple Databases

The Run on Multiple databases option allows you to access the Multiple Destination page so that you can specify the destination database(s) to submit the job.

After specifying the scheduling options, press Next to continue.

Proceed to Backup Wizard: Job Information on page 8-14.

Backup Wizard: Job Information

The Job Information page allows you to specify the name and description for the backup job and what you would like to do with the job as it will appear in the Oracle Enterprise Manager Console Job System.

The choices for what to do with the job are listed below:

- Submit the job now. The job is submitted to the Management Server.
- Do not submit the job now, but add it to the Job Library for later use. The Enterprise Manager job is added to the Job Library. From the Console, you can edit or submit the job at a later date.
- Submit the job now and add it to the Job Library

You can view the status of the job by selecting the Active and History page tabs in the Console's Job window.

Note: If a service name entry does not exist for the target database in its `tnsnames.ora` file which matches the entry in the `tnsnames.ora` file in the Oracle Enterprise Manager home, the following error message appears:

```
ora-12154: TNS:could not resolve service name
```

```
Last Message : The job failed
```

For the backup job to run, you must update the `tnsnames.ora` file in the target database node, bounce the agent, and rediscovering the target node in the Console.

Backup Wizard: Summary

The Summary page appears when you click the Finish button. Information about the backup management job you just created using the Wizard is summarized before submission to the Oracle Enterprise Manager job system.

Clicking OK submits the job the Oracle Enterprise Manager job system.

Recovery Wizard

The Recovery wizard helps you restore and recover various objects like databases, datafiles, and tablespaces.

Before attempting to use the Recovery Wizard, make sure the following requirements are met:

- The Oracle target database is version 8 or above
- The application is connected to a Management Server
- The preferred credentials is SYSDBA or you have created a backup configuration and overridden the preferred credentials. If your target database is not configured for remote sysdba operation, see "Configuring a Remote Database for Backup or SYSDBA Administration" on page 8-31.
- the job and event systems are fully functional.

You can access the wizard through the Console and Storage Manager if you are connected to a Management Server.

Note: You can submit a recovery job directly from the Job System, but the recommended method is through the wizard.

Recovery Wizard: Introduction Page

The Recovery Wizard 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.

A checkbox "Display this page next time" is available where you can specify not to show this page once you are familiar with the Recovery Wizard. By default, the checkbox is selected. If you deselect the checkbox, the wizard starts on the Recovery Selection page the next time.

Press Next to continue. Proceed to Recovery Wizard: Recovery Selection on page 8-16.

Recovery Wizard: Recovery Selection

The Recovery Selection page of the Recovery Wizard allows you to select from the following types of recovery operations:

- Database
- Tablespaces
- Datafiles

Note: If the database is mounted, but not open, all three options are available.

If the database is open and in ARCHIVELOG MODE, only tablespaces and datafiles are available.

The type of recovery you select determines which procedure the Restore Wizard guides you through.

Note: Depending on the status of the target database (ARCHIVELOG, NOARCHIVELOG, mounted and open, mounted only, or started), some options will be disabled or not displayed.

For more information on recovery operations and related strategies, see the *Oracle8i Administrator's Guide* or the *Oracle8i Backup and Recovery Guide*.

Database

A recovery of an entire database is a recovery of all database files that belong to a database. This selection only appears if the database is in the "MOUNT" state.

Note: Use Instance Manager or the Console to open the database once the job has completed. You can view the status of the job by selecting the Active and History page tabs in the Console's Job window.

Tablespace

A tablespace recovery is a recovery of a subset of the database. Tablespace recoveries are generally valid if the database is operating in ARCHIVELOG mode.

Datafile

A datafile recovery is a recovery of a single datafile. Datafile backups, which are not as common as tablespace recoveries, are generally valid if you run the database in ARCHIVELOG mode.

Note: If you are in NOARCHIVELOG mode, your database must be mounted in order to perform a restore. To perform a full database recovery, your database must be in the MOUNTED mode. Use Instance Manager or the Console to open the database after the job has completed. You can view the status of the job by selecting the Active and History page tabs in the Console's Job window.

Perform Recovery Without Restoring the Datafiles

A checkbox "Perform recovery without restoring the datafiles" is available. By default, the checkbox is deselected. If you select the checkbox, a recover is performed without a restore. The Rename page will not appear if you have selected the "Perform recovery without restoring the datafiles" option.

Use the "Perform recovery without restoring the datafiles" option if the datafile exists and is not corrupted and the database only requires to recover (apply the redo logs) to the datafile.

After choosing your recovery selection, press Next to continue.

Proceed to Recovery Wizard: Restore Until on page 8-17 if you have chosen Entire database as your recovery selection.

Proceed to Recovery Wizard: Tablespaces Page on page 8-17 if you have chosen Tablespace as your recovery selection.

Proceed to Recovery Wizard: Datafiles Page on page 8-18 if you have chosen Datafiles as your recovery selection.

Recovery Wizard: Restore Until

If you press the Finish button now, the database will be restored to the most recent time possible.

If you want to restore to a point previous to now, then you must check the Until checkbox and enter a date and time when the date and time fields are activated.

Note: Use Instance Manager or the Console to open the database after the job has completed. You can view the status of the job by selecting the Active and History page tabs in the Console's Job window.

This feature is available if your database is in the MOUNTED mode.

Press Next to continue. Proceed to Recovery Wizard: Rename Page on page 8-18.

Recovery Wizard: Tablespaces Page

A tablespace backup is a backup of a subset of the database. Tablespace backups are generally valid if the database is operating in ARCHIVELOG mode.

The Tablespace page allows you to select the tablespaces you want to recover. Choose from the list of tablespaces shown on the Available Tablespaces window.

Note: Tablespaces are displayed with a different icon if they require recovery.

Move the selection(s) to and from the Available Tablespaces and Selected Tablespaces windows by using arrow buttons between them.

Press Next to continue. Proceed to Recovery Wizard: Rename Page on page 8-18.

Recovery Wizard: Datafiles Page

A datafile recovery is the recovery of a single datafile. Datafile recoveries, which are not as common as tablespace recoveries, are generally valid if you run the database in ARCHIVELOG mode.

The Datafiles page allows you to select the datafiles you want to recover. Datafiles which require recovery are displayed with a different icon. Choose from the list of datafiles shown on the Available Datafiles window. Full path and name of the datafile are shown.

Move the selection(s) to and from the Available Datafiles and Selected Datafiles windows by using arrow buttons between them.

You may choose to include the control files with the recovery.

Press Next to continue. Proceed to Recovery Wizard: Rename Page on page 8-18.

Recovery Wizard: Rename Page

The Rename page enables you to restore the selected datafile(s) to a new location. When datafiles are restored to a new location, they are considered datafile copies. For this reason, a switch is automatically performed.

Original Name

The full path and filename of the selected datafile(s). This field is not editable.

New Name

Enter a new filename and/or a new path.

You can copy all or sections of the Original Name entry and paste it into the New Name field by highlighting the desired text and selecting the desired action from the context-sensitive right mouse menu.

Press Next to continue. Proceed to Recovery Wizard: Configuration Page on page 8-19.

Recovery Wizard: Configuration Page

The Configuration page allows you to choose the configuration for recovery and view the details that you have set up previously using the Create Backup Configuration option.

Choose the configuration for this recovery

Allows you to choose a configuration you have previously set up using the Create Backup Configuration Property Pages.

Description of the configuration

The description of the configuration.

Summary of the configuration:

- Name of the configuration
- Type of Destination (disk or tape)
- Channel Name (User-specified name of the channel)
- Drive and path where backup sets are stored or recovered.
- Recovery Catalog Credentials if you are using a recovery catalog

Recovery Wizard: Summary

The Summary page appears when you click the Finish button. Information about the backup management job you just created using the Wizard is summarized before submission to the Oracle Enterprise Manager job system.

Clicking OK submits the job the Oracle Enterprise Manager job system.

Clicking Cancel returns you to the wizard.

Create Backup Configuration

A configuration is a set of defaults you set up for backup and recovery. You can use the configuration for one database or many databases depending if the systems are the same.

Before attempting to use the Oracle Enterprise Manager Backup Management wizards, make sure the following requirements are met:

- The Oracle target database is version 8 or above
- The application is connected to a Management Server
- The preferred credentials is SYSDBA or you have created a backup configuration and overridden the preferred credentials. If your target database is not configured for remote sysdba operation, see "Configuring a Remote Database for Backup or SYSDBA Administration" on page 8-31.
- the job and event systems are fully functional

You can access the wizards through the Console and Storage Manager if you are connected to a Management Server.

Create Backup Configuration: General Page

The General page allows you to give a name and description to a set of defaults used for backup, recovery, and catalog maintenance operations.

You can access the wizards through the Console and Storage Manager if you are connected to a Management Server.

Create Backup Configuration: Channels Page

The Channels page allows you to specify a channel or channels. A channel is an output device for backup and restore operations. Multiple channels can be created to allow parallel backup/recovery by a single job.

Note: At least one channel must exist before performing a backup, restore, or recover operation.

Attention: Oracle8 can only allocate one Recovery Manager channel at a time, thus limiting the parallelism to one stream. The Oracle8 Enterprise Edition allows unlimited parallelism. See the *Oracle8i Installation Guide for Windows NT* and the *Oracle8i Administrator's Guide for Windows NT* for more information about the features available with Oracle8i and Oracle8i Enterprise Edition.

Choose the Type of Destination for Backups

Select from one of the following types:

- Disk: Backup sets written to disk.
- Tape: Backup sets written to tape.

Note: You can have multiple channels, but the channels defined must have the same type of destination.

Channel Name

User-specified name of the channel.

Note: 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.

File Name

Unique backup set name. The following parameters can be used:

- %p 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 Unique name.

Note: 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 Backup set number. The counter value starts at 1 and is unique for the lifetime of the control file.
- %t Backup set timestamp. **Note:** The combination of %s and %t can be used to form a unique name for the backup set.

Directory (for Disk)

Drive and path where backup sets are stored. The Directory field is active when channel type is set to Disk.

Parameters (for Tape)

Parameters regarding the device to allocate. The Parameters field is active when channel type is set to Tape. For information on `parms` parameter, see the *Oracle8i Backup and Recovery Guide*.

Limits

Press the Limits button to access the Limits page in order to set the limits for any backup.

Insert

Allows you to create another channel definition.

Note: The channels defined must have the same type of destination.

Remove

Deletes the selected channel.

Limits

The Limits page allows you to set the limits for any backup or copy operation. 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.

Maximum Capacity (Mb)

When checked, it allows you set the maximum number of megabytes that a backup operation can write to a single backup piece.

Maximum Read Rate (Kb/s)

When checked, it 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.

Maximum Open Files

When checked, it allows you to control the maximum number of input files that a backup 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.

Create Backup Configuration: Recovery Catalog Page

The Recovery Catalog page allows you to choose where you want the backup information to be stored.

The recovery catalog contains the following information:

- datafile and archivelog backup sets and backup pieces
- datafile copies
- archived redo logs and copies of them
- tablespaces and datafiles in the target database
- named user-created sequences of commands called stored scripts.

The choices are listed below.

In the target database's control file

The target database's control file. For information regarding the pros and cons of using a recovery catalog rather than a control file, see the *Oracle8i Backup and Recovery Guide*.

In the Recovery Catalog

A repository of information that is used and maintained by the Recovery Manager. For information regarding the pros and cons of using a recovery catalog rather than a control file, see the *Oracle8i Backup and Recovery Guide*.

Username

A user which contains the schema for the recovery catalog.

Password

The password of the user which contains the schema for the recovery catalog.

Service Name

The service name of the database where the recovery catalog resides.

If you have chosen to use a recovery catalog and clicked OK, the Create Backup Configuration checks that:

- the user exists
- the recovery catalog is installed
- the database is registered. If the database is not registered, it registers it.

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.

See the *Oracle8i Backup and Recovery Guide* for information on operating with a recovery catalog.

See the *Oracle Enterprise Manager Configuration Guide* for complete information on setup and configuration.

Create Backup Configuration: Backup Parameters Page

The Backup Parameters page allows you to set storage parameters for the current backup set.

Override Recovery Manager defaults

Check the "Override Recovery Manager defaults" box to use this page to override the Recovery Manager default settings that the Recovery Manager has calculated.

Maximum files at a time in a backup set

When checked, it 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.

Maximum size of a backup set

When checked, it 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.

Maximum size of a backup set for archivelogs

When checked, it 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.

Create Backup Configuration: Preferred Credentials Page

The Preferred Credentials page allows you to specify preferred credentials for running backup jobs. These configurations will take precedence over the preferred credentials set in the Console.

Use Preferred Credentials if you do not want to set the login credentials to SYSDBA in the Console. These credentials will only be used for backup and recovery jobs.

If you want to configure your database for remote sysdba operation, see "Configuring a Remote Database for Backup or SYSDBA Administration" on page 8-31.

Note: The Backup Configuration you have set become your default backup configuration if you use the Backup and Recovery wizards and submit a job.

Catalog Maintenance Wizard

The Catalog Maintenance wizard helps you perform basic recovery catalog operations like register the database, reset the database, and re-synchronize the catalog.

Before attempting to use the Catalog Maintenance wizard, make sure the following requirements are met:

- At least one backup configuration set exists which uses a recovery catalog.
The recovery catalog is a separate database normally residing on a remote system that contains information about the structure and previous backups of the target database.
One recovery catalog can manage many databases. It is recommended that you use resync catalog for every database that uses a recovery catalog on a regular basis.
- The Oracle target database is version 8 or above
- The application is connected to a Management Server
- The preferred credentials is SYSDBA or you have created a backup configuration and overridden the preferred credentials. If your target database is not configured for remote sysdba operation, see "Configuring a Remote Database for Backup or SYSDBA Administration" on page 8-31.
- the job and event systems are fully functional

Note: You can submit a catalog maintenance job directly from the Job System, but the recommended method is through the wizard.

Catalog Maintenance Wizard: Introduction Page

The Catalog Maintenance wizard guides you through the process of defining the recovery catalog operation and submits a catalog maintenance job through the Enterprise Manager to complete the operation.

A checkbox "Display this page next time" is available where you can specify not to show this page once you are familiar with the Catalog Maintenance Wizard. By default, the checkbox is selected. If you deselect the checkbox, the wizard starts on the Operation Choice page the next time.

You can view the status of the job by selecting the Active and History page tabs in the Console's Job window.

Press Next to continue. The Catalog Maintenance Wizard: Operation Choice Page appears.

Catalog Maintenance Wizard: Operation Choice Page

The Operation Choice page allows you to manage the recovery catalog.

In order for the Backup wizard to use the recovery catalog, the database must be registered first.

Register database

Register the target database with the Recovery Catalog. 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.

Resync catalog

Re-synchronizes 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. The recovery catalog must be periodically resynced from the target database in order to ensure the catalog is up-to-date. The recovery catalog is not updated automatically when a log switch occurs or when a log is archived. Also, any structural changes to the target database would require re-synchronization of the Recovery Catalog.

Depending on how often your target database structure changes (including archive logs), you must decide how often to resynchronize the catalog. It is recommended that you should resynchronize the catalog at least once a day. If your site generates many archived logs each day, then you should resynchronize the catalog more frequently.

Reset database

Reset the Recovery Catalog for the target 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 "Coordinating Time-Based and Change-Based Distributed Database Recovery" section of the *Oracle8i Backup and Recovery Guide* for information on the RESETLOGS option.

Press Next to continue. The Catalog Maintenance Wizard: Configuration Page appears.

Catalog Maintenance Wizard: Configuration Page

The Configuration page allows you to choose the configuration which has the recovery catalog information that you have set up previously using the Create Backup Configuration option.

Choose the configuration

Allows you to choose a configuration you have previously set up using the Create Backup Configuration Property Pages.

Description of the configuration

The description of the configuration.

Summary of the configuration:

- Name of the configuration
- Type of Destination (disk or tape)
- Channel Name (User-specified name of the channel)
- File Name (Unique backup set name)
- Drive and path where backup sets are stored.

For users registering a database or resynchronizing a catalog, press Next continue.

Proceed to Catalog Maintenance Wizard: Multiple Destinations Page on page 8-27 if you have chosen to register a database.

Proceed to Catalog Maintenance Wizard: Schedule on page 8-28 if you have chosen to resynchronize the catalog.

For users resetting the database, press Finish to submit a catalog maintenance job through the Enterprise Manager to complete the operation.

Catalog Maintenance Wizard: Multiple Destinations Page

The Multiple Destination page allows you to specify the destination database(s) to submit the job. You can select from the list of databases shown on the Available Databases window.

Move the selection(s) to and from the Available databases and Selected databases windows by using arrow buttons between them. By default, the database to which you are connected is shown in the Selected Databases window.

Press Finish to submit a catalog maintenance job through the Enterprise Manager to complete the operation.

Catalog Maintenance Wizard: Schedule

The Schedule page enables you to schedule the execution of a catalog maintenance job.

Run Job

Select the frequency of catalog maintenance job. The choices are Immediately, Once, On Interval, On Day of Week, and On Date of Month.

- Immediately submits the task as soon as you finish the wizard. The task executes only one time.
- Once schedules the task only one time at the date and time you choose.
- On Interval enables 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 the scroll buttons. You can also type in a new value.
- On Day of Week enables you to schedule the task on one or multiple days (Sunday, Monday, etc.) of the week. Click the days of the week to select the days you want the task scheduled.
- On Date of Month enables you to schedule the task on one or multiple days (1 - 31) of the month. Click 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. Choose the first day of the next month instead.

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 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 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 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 the scroll buttons to change the value. You can also type in new values.

Time Zone

Shows the time zone. The agent schedules the task execution at each destination based on the actual system time of each agent. Tasks are not necessarily run simultaneously.

Run on Multiple Databases

The Run on Multiple databases option allows you to access the Multiple Destination page so that you can specify the destination database(s) to submit the job.

After specifying the scheduling options, press Next to continue.

Proceed to Catalog Maintenance Wizard: Multiple Destinations Page on page 8-27.

Catalog Maintenance Wizard: Summary

The Summary page appears when you click the Finish button. Information about the backup management job you just created using the Wizard is summarized before submission to the Oracle Enterprise Manager job system.

Clicking OK submits the job the Oracle Enterprise Manager job system.

Backup Configuration Library

The Backup Configuration Library page displays configurations created and saved using the Backup Configuration option.

Name: The name of the configuration

Description: The description of the configuration

Use Recovery Catalog: Whether the configuration uses a recovery catalog

This is useful if you want to perform the following tasks:

- use the same configuration at different times
- make minor modifications to an existing configuration without having to redefine it from scratch
- create a configuration similar to an existing configuration without having to define it from scratch
- delete an existing configuration

Double-click an entry in the Backup Configuration Library list to start the Backup Configuration option with the stored parameters.

Configuring a Remote Database for Backup or SYSDBA Administration

Before using one of the backup management wizards or starting up or shutting down a database through Oracle Enterprise Manager, you need to set up your database for remote sysdba management.

The following procedure outlines the operations you may need to perform on both the server and client if they have not been performed already through the Database Configuration Assistant.

Additional Information: Refer to the operating system-specific documentation for your Oracle database for the exact name of the password file creation utility for your system.

Server Side (NT)

1. From the command prompt, run the ORAPWD executable to create a new password file.

```
$ORACLE_HOME/bin/orapwd file=<fname> password=<password> entries=<users>
```

where

fname is the name of password file (mandatory),

password is the password for SYS (mandatory),

entries is the maximum number of distinct DBAs and OPERs (optional),

There are no spaces around the equal-to (=) character.

The <fname> will be created in the Oracle_Home/Database directory.

2. Edit the `init.ora` file located in the `Oracle_Home/Admin/<SID>/pfile` directory. Add or change the following parameter
`remote_login_passwordfile=exclusive` in the `init.ora` file.
3. Shut down your database and then restart it from Control Panel > Services.
4. Grant SYSDBA and/or SYSOPER roles to users who will be performing remote DBA operations.

Server Side (UNIX)

1. Login as ORACLE and set your environment variables.

```
setenv ORACLE_HOME your_oracle_home
setenv ORACLE_SID your_oracle_sid
```

Note: This example assumes that you have set up your database as ORACLE.

2. Change to the dbs directory in your Oracle home.

```
cd $ORACLE_HOME/dbs
```

3. Run the ORAPWD executable to create a new password

```
file.$ORACLE_HOME/bin/orapwd file=<fname> password=<password>
entries=<users>
```

where

fname is the name of password file (mandatory),

password is the password for SYS (mandatory),

entries is the maximum number of distinct DBAs and OPERs (optional),

There are no spaces around the equal-to (=) character.

The <fname> will be created in the \$Oracle_Home/dbs directory.

4. Edit the *init.ora* file. Add or change the following parameter
remote_login_passwordfile=exclusive in the *init.ora* file.

5. Shut down your database and then restart it.

- a. Start SQL*Plus Worksheet.
- b. Log in as a user with the SYSDBA role.
- c. Type

```
shutdown immediate
startup
```

6. Grant SYSDBA and/or SYSOPER roles to users who will be performing remote DBA operations

Recovery Catalog Setup and Registration for Oracle8i

The *recovery catalog* is a repository of information that is used and maintained by Recovery Manager. Recovery Manager uses the information in the recovery catalog to determine how to execute requested backup and restore actions.

Because most of this information is also available from the target database's control file, you have the option of using the control file instead of creating a recovery catalog. This operational mode is appropriate for small databases. No setup is required if you are using the control file. However, Oracle recommends using a recovery catalog. If the control file is lost, the database cannot be recovered.

To use Oracle8i Recovery Manager with a recovery catalog, you must first install the recovery catalog and then register your database. It is recommended that you install the recovery catalog on a different location than the target database.

For more information on recovery catalog maintenance, Recovery Manager, and backup and recovery strategies, see the *Oracle8i Backup and Recovery Guide*.

Important: Your Oracle Enterprise Manager Job system must be functioning properly before attempting to install the recovery catalog. You will need the Job system to register the target database in the recovery catalog.

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 and later) on the Target Machine

Create a Tablespace

1. Start Storage Manager.
2. Log in as a user with the DBA role.
3. Select Create from the Tablespace menu. The Create Tablespace property sheet appears.
4. In the Name field, type the name of the new tablespace. For example, `rcvcat`.
5. Click the Add button in the Datafiles section. The Create Datafile property sheet appears, allowing you to specify a new datafile which will belong to the new tablespace.
6. In the Name field, type in complete path and name of the datafile. For example:
`c:\orant\oradata\rcvcat01.dbf`.
7. In the Size section, enter the size of the new datafile. For example, 10 M. The datafile size will depend on the number of databaes and how often they are backed up.
8. Click the OK button in the Create Datafile property sheet.
9. Click the Create button in the Create Tablespace property sheet.

Create a User

1. Start Security Manager.
2. Log in as a user with the DBA role.
3. Choose Create from the User menu. The Create User property sheet appears.
4. In the General Page, fill in the following information:

In the Name field, enter the name of the new user. For example, `rman`.

In the Password and Confirm Password fields, enter a password. For example, `rman`.

Choose the default tablespace. For example, `RCVCAT`.

Choose the temporary tablespace. For example, `TEMPORARY_DATA` or `TEMP`.
5. In the Roles Page, grant the `RECOVERY_CATALOG_OWNER` role to the user.
6. In the Quotas Page, specify an unlimited quota for the default tablespace. In this example, the default tablespace is `RCVCAT`.
7. Click the Create button after specifying the requisite parameters.

Run the catrman Script (for pre-8i only)

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 so that errors will be written in a log file.
4. Execute the catrman script located in the Oracle_Home/rdbms/admin or Oracle_Home\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>/rdbms/admin/catrman and press the Execute button.

Create the Recovery Catalog (for 8i and later) on the Target Machine

1. Connect to the recovery catalog from the operating system command line:

```
%> rman rcvcat 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 rcvcat;
```

- For Windows NT, type:

```
RMAN> create catalog tablespace "RCVCAT";
```

Registering the Recovery Catalog

To register the recovery catalog for the first time, follow the steps 1 through 7.

1. Start the Oracle Enterprise Manager Console.
2. Click the right mouse button on the database in the Navigator tree. The context-sensitive Backup Management menu appears.
3. Modify a backup configuration you are already using or choose Create Backup Configuration from the Backup Management menu so that you can create a backup configuration that uses a recovery catalog. The Create Backup Configuration property sheet appears.
4. Fill out the General and Channels property pages.
5. In the Recovery Catalog Page, fill in the following information:
Choose "in a recovery catalog" as where you want your backup information stored.
In the Username field, enter the name of the user. For example, `rman`.
In the Password field, enter a password. For example, `rman`.
In the Service field, enter the database where the Recovery Catalog resides.
6. Fill out the Preferred Credentials page.
7. Press the Create button. The database you selected will automatically be registered with the recovery catalog.

Note: You only need to register the database once.

To use the same configuration on subsequent databases, follow steps 8 through 16.

8. Click the right mouse button on the database in the Navigator tree. The context-sensitive Backup Management menu appears.
9. Choose Catalog Maintenance from the Backup Management menu so that you can register the target database in the recovery catalog. The Catalog Maintenance appears. Click Next.
10. In the Operation Choice Page, choose Register Database. Click Next.
11. In the Configuration Page, ensure that you are using the Backup Configuration you have created/modified. Click Next.

12. In the Multiple Destinations Page, choose the destination database(s) to submit the job.
13. Click the Finish button. At this point, the database registration is sent as a job to the Oracle Enterprise Manager Job system.
14. When the summary screen appears, click OK to complete the operation.
15. In the Console, click the registration job in the Active Jobs list to see the current state of the job.
16. Once the job is completed, check the job history in the Jobs property sheet to make sure that the job completed successfully.

When the job completes successfully, your backup and recovery environment is configured. For more information on recovery catalog maintenance, Recovery Manager, and backup and recovery strategies, see the *Oracle8i Backup and Recovery Guide*.

Dropping the Recovery Catalog

If you do not want to maintain a recovery catalog, then you can drop the recovery catalog schema from the tablespace in 8i. The drop catalog command deletes all information from the recovery catalog. Hence, if you have no backups of the recovery catalog schema, then all backups of all target databases managed by this catalog become unusable.

The drop catalog command is not appropriate for "unregistering" a single database from a catalog that has multiple target databases registered. If you try to delete the information for one target database by dropping the catalog, you thereby delete the information for all target databases.

To drop the recovery catalog schema:

1. Use RMAN to connect to the target and recovery catalog databases.

```
%> rman rcvcat rman/rman@rcvcat
```

2. Issue the drop catalog command:

```
RMAN> drop catalog
```

For drop catalog command syntax, see the *Oracle8i Backup and Recovery Guide*.

Upgrading the Recovery Catalog

If you use a version of the recovery catalog that is older than that required by the RMAN executable, then you must upgrade it. For example, you must upgrade the catalog if you use an Oracle8i RMAN with an Oracle8i recovery catalog.

Note: You will receive an error if the recovery catalog is already at a version greater than that required by the RMAN executable. RMAN permits the command to be run if the recovery catalog is current, however, so that you can re-create packages if necessary. Check the message log for error messages generated during the upgrade.

To upgrade the recovery catalog:

1. Use RMAN to connect to the target and recovery catalog databases.

```
%> man rcvcat rman/rman@rcvcat
```

```
RMAN-06008: connected to recovery catalog database  
RMAN-06186: PL/SQL package rcat.DEMS_RCVCAT version 08.00.04 in RCVCAT  
database is too old
```

2. Issue the upgrade catalog command:

```
RMAN> upgrade catalog
```

```
RMAN-06435: recovery catalog owner is rman  
RMAN-06442: enter UPGRADE CATALOG command again to confirm catalog upgrade
```

3. Enter the update catalog command again to confirm:

```
RMAN> upgrade catalog
```

```
RMAN-06408: recovery catalog upgraded to version 08.01.05
```

For upgrade catalog command syntax, see the *Oracle8i Backup and Recovery Guide*.

Enterprise Manager Event Handler

Important: The Event Handler is a pre-production feature in this release of Enterprise Manager and should not be used in a production environment.


Applications accessing the Oracle Management Server event services may require trapping of specific event conditions and the ability to handle them appropriately. This can be achieved using the Event Handler, an application that listens for event notifications from the Oracle Management Server and allows these events to be further processed by event "adapters". An adapter is an additional piece of software that takes events from the Event Handler for further processing.

Enterprise Manager 2.1 ships with 2 event adapters:

- *Event Logger* - allows the Event Handler to log events to designated log files and
- *Command Executor* - allows simple operating system commands to be executed in response to an event occurrence.

These adapters permit other types of applications to integrate with the Enterprise Manager event system and trigger other operations based on these trapped events. You can have either or both adapters working with the Event Handler.

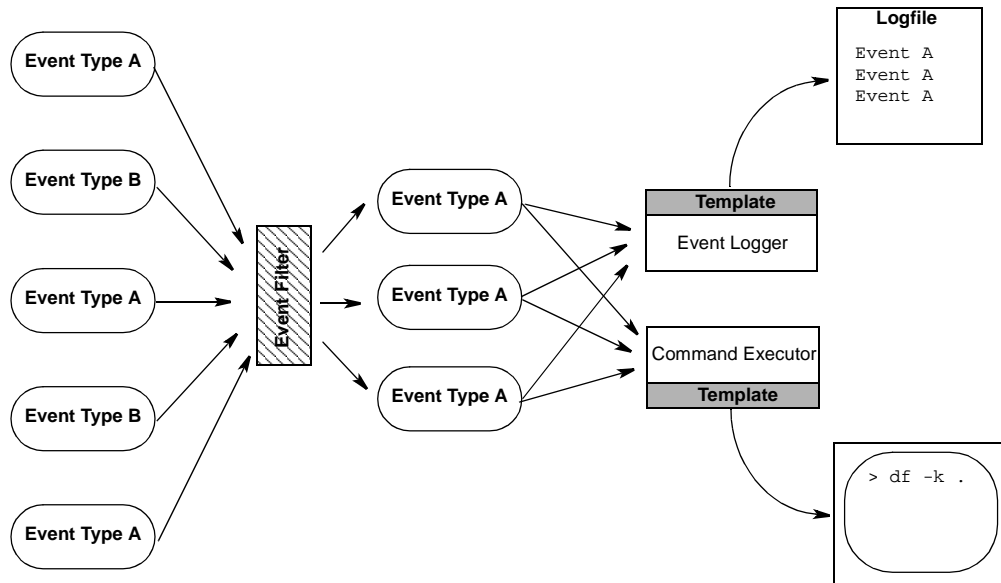
Prior to passing event notifications for further processing, the Event Handler also provides a simple filtering mechanism that allows system administrators to specify the conditions by which the events are passed to either or both event adapters.

Figure A-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 adapters. Once the adapters receive the event, they further process the event based on the *templates* set up for the adapter. Templates provide a way to customize the behavior of the adapters. For the Event Logger, this means specifying which events should be logged. 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 A-2 Event Handler Process



How to Set Up the Event Handler

To set up and implement an Event Handler, you must perform the following steps:

1. Define the filters.

Filters determine which events are passed on to the Event Logger and Command Executor adapters.

2. Define the templates.

Templates tell the Event Logger and Command Executor adapters how to process the event. For the Event Logger, a template specifies which events should be logged and how the information should be formatted. For the Command Executor, a template specifies the events the adaptor should respond to and the operating system command it should execute in response to that event.

3. Determine how to start up the Event Handler

There are two ways to start the Event Handler: As a Management Server service, or as a client application.

4. Start up the Event Handler

Each of these steps are discussed in detail later in this appendix.

Defining Event Filters

When the Event Handler starts, it uses "filters" to determine which events are passed on to the Event Logger and/or Command Executor adapters. The system administrator can designate what events, if any, are sent to the Event Logger and Command Executor adapters. This is called "filtering." Events may be filtered on a global basis (i.e. applied to both adapters) or on a per-adapter basis. Filters are set up by means of entries in the *omsconfig.properties* file located in the `$ORACLE_HOME/sysman/config` directory. Each entry specifies the conditions by which the events are passed to the adapters. The conditions are based on 'eventname' and/or 'node' on which the event occurred. An event passes thru the filter if it meets the conditions specified by the filter.

Each filter is distinguished by having a common name associated with it and the entries for that filter are grouped under this common name.

Filter Syntax

Actual filter entries are defined in the *omsconfig.properties* file. The syntax for filters are as follows:

Syntax for global filters:

```
/com/oracle/sysman/em/eventHandler/global_filters/<filter-name>/<'eventname' or 'node'> = <value>
```

Syntax for filters for the Event Logger only:

```
/com/oracle/sysman/em/eventHandler/eventlogger_filters/<filter-name>/<'eventname' or 'node'> = <value>
```

Syntax for filters for the Command Executor only:

```
/com/oracle/sysman/em/eventHandler/commandexecutor_filters/<filter-name>/<'eventname' or 'node'> = <value>
```


The character "*" can be used as a wildcard character to specify all possible values. 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 be 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 all event adapters.

2. The event must pass at least one filter
3. An event is forwarded to an adapter (Event Logger or Command Executor) if it passes thru either a global filter or a filter for that particular adapter.

By default, all events are suppressed, and at least one filter (global or adapter-specific) must be present in order for events to be forwarded to any of the adapters. See Sample Filters and Templates: on page A-14 for examples of filters.

Assuming the event has passed thru the filters, they are forwarded on to the event adapters, Event Logger or Command Executor. We will describe what happens to them next.

Important: Another level of filtering occurs when you start the Event Handler itself. In order to start the Event Handler, you need to provide a valid Management Server administrator and password. This Management Server administrator will have associated events that it will be able to receive. If you start the Event Handler using a particular Management Server administrator, then the Event Handler will only be able to receive event notifications to which that Management Server administrator has been given at least VIEW permissions. If you start the Event Handler using the sysman or superuser account, it will receive ALL events. Hence all events that the Management Server administrator is not entitled to receive will *not* be received by the Event Handler.

The Event Logger Templates

The Event Logger logs events received from the Management Server to a file.

The default log file is `$ORACLE_HOME/sysman/log/event.log`. 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, read as entries 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 `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.

You can further customize Event Logger operation by using templates. Templates are strings that generically represent the format of a message as a way to customize the format of event notifications written to a log file. A template contains one or more place holders, which are symbolic representations for pertinent pieces of information about the event. Place holders are enclosed within "%" characters to distinguish them from ordinary words in the template string. Available place-holders are:

Table 8–1 Event Handler Place Holders

Place Holder	Definition
<code>%eventname%</code>	The name of the event.
<code>%severity%</code>	The severity of the event as a string. (Alert, Clear, Warning, Node Down)
<code>%timestamp%</code>	The timestamp of the event occurrence. (MM-dd-yyyy hh:mm:ss Example: 05-22-00 05:22:00 AM)
<code>%targetname%</code>	Name of the target.
<code>%targettype%</code>	Type of the target. (oracle_sysman_node, oracle_sysman_database, oracle_sysman_listener)

Table 8–1 Event Handler Place Holders

Place Holder	Definition
%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.
%output%	The output associated with the event occurrence.

Templates are declared by making entries in the *omsconfig.properties* file. Events can be assigned to templates based on event names or originating node (the same criteria used by the Event Handler filtering system). Templates use this format to specify which events the template should apply:

```
/com/oracle/sysman/em/eventlogger/templates/<template-name>/<'event_name' or 'node'>=<value>
```

To specify the message format that the template should use, the following format must be used:

```
/com/oracle/sysman/em/eventlogger/templates/<template-name>/message=<message format>
```

An example of a template entry in the *omsconfig.properties* 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%
```

The template definition above assigns all events named *cputest* to the template named *foo*. Whenever a *cputest* event fires, the logged output will appear as follows:

```
cputest fired on smptest16: Cpu usage on smptest16 is high! occ_no: 21 Severity: Alert Time: 10-21-1999 02:39:29 PM
```

Note: Templates override the default message format.

As with Event Handler filters, the wildcard character "*" can be used for event names and node names. If more than one template criterion matches an event

occurrence, only one of them is randomly chosen. For this reason, care should be used when assigning events to templates.

The Command Executor Templates

The Command Executor adapter 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 an event occurrence. A template can contain several place-holders, which are symbolic representations for interesting pieces of information about the event. Place-holders are enclosed within '%' characters to distinguish them from ordinary words in the template string. For a list of currently used place-holders, see Table 8-1, "Event Handler Place Holders" on page A-6

As with the Event Logger, templates are declared by making entries in the *omsconfig.properties* file. Events can be assigned to templates based on event names or originating node (the same criteria offered by the event-handler filtering system).

Templates use this format to specify the events to which the template should apply.

```
/com/oracle/sysman/em/eventlogger/templates/<template-name>/<'event_name' or  
'node'>=<value>
```

Templates use the following format to specify the command to which the template should apply:

```
/com/oracle/sysman/em/commandexecutor/templates/<template-name>/command=<value>
```

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
```

As with event-handler filters, the character '*' can be used to wildcard event names and node names. If more than one template criterion matches an event occurrence, only one of them is randomly chosen. It is important, therefore, to be careful when assigning events to templates.

Starting up and Running the Event Handler

There are 2 ways to start the Event Handler

- As a Management Server service
- As a client application

In both cases, make sure you have the following entries in the *omsconfig.properties* file. This is located in the `%ORACLE_HOME%\sysman\config` directory. `%ORACLE_HOME%` refers to the Oracle Home directory where Enterprise Manager was installed.

Note: For convenience, these entries are in the `%ORACLE_HOME%\sysman\admin\omsconfig.properties.EventHandler` file. Simply copy the entries from this file into the *omsconfig.properties* file.

The following configuration file entries relate to the Event Logger:

```
/com/oracle/sysman/em/eventHandler/handlers/eventlogger/entrypoint=oracle.sysman.vda.vdae.VdaeEventLogger
/com/oracle/sysman/em/eventHandler/handlers/eventlogger/processOnStartup=true
```

The following configuration file entries relate to the Command Executor:

```
/com/oracle/sysman/em/eventHandler/handlers/commandexecutor/entrypoint=oracle.sysman.vda.vdae.VdaeCommandExecutor
/com/oracle/sysman/em/eventHandler/handlers/commandexecutor/processOnStartup=true
```

Important: There should be **NO carriage returns or linefeed** characters within each entry. This is true for all entries in any of the Enterprise Manager configuration files.

Starting the Event Handler

You need to supply a valid Management Server administrator and password in order to start the Event Handler.

1. Start the Event Handler as a Management Server Service

Before you can start the Event Handler as a Management Server service, you must add the following entries to the *OMSExternalRegistry.registry* file (located in %ORACLE_HOME%/sysman/config), substituting the appropriate values for <oracle-home>, <OMS User>, <OMS password>, and <OMS nodename>:

```
/com/oracle/sysman/em/oms/services/external/EventHandlerSession/external_
service_name = EventHandlerSession
/com/oracle/sysman/em/oms/services/external/EventHandlerSession/name =
EventHandlerSession
/com/oracle/sysman/em/oms/services/external/EventHandlerSession/oms_service_
interface = oracle.sysman.vda.vdae.VdaeEventHandlerSession
/com/oracle/sysman/em/oms/services/external/EventHandlerSession/startup_
mode=1
/com/oracle/sysman/em/oms/services/external/EventHandlerSession/name_value_
pair=ORACLE_HOME=<oracle-home>
/com/oracle/sysman/em/oms/services/external/EventHandlerSession/args=-s,<OMS
User>,<OMS Password>,<OMS nodename>
/com/oracle/sysman/em/oms/services/external/EventHandlerSession/service_
classpath=
/com/oracle/sysman/em/oms/services/external/EventHandlerSession/outfile_
destination=<oracle-home>/sysman/log
/com/oracle/sysman/em/oms/services/external/EventHandlerSession/errfile_
destination=<oracle-home>/sysman/log
/com/oracle/sysman/em/oms/services/external/EventHandlerSession/logfile_
name=eventHandler
```

Note: For convenience, these entries can be found in the `%ORACLE_HOME%/sysman/admin/OMSEExternalRegistry.registry.EventHandler` file. Simply copy the entries from this file and paste them into the `OMSEExternalRegistry.registry` file located in your `%ORACLE_HOME%/sysman/config` directory.

2. Start the Management Server.

For UNIX and Windows NT, enter the following at the command line:

```
oemctrl start oms
```

For Windows NT, you can start and stop the Management Server as a Service in the Control Panel (OracleORANTManagementServer).

3. Start the Event Handler as a client application

Use the `oemctrl` utility to start the Event Handler.

```
oemctrl start eventhandler <oms-user> <oms-password> <nodename>
```

where:

<oms-user> is a valid Management Server administrator

<oms-password> is the corresponding password

<nodename> is the node where the Management Server is running.

Stopping the Event Handler

There are 2 ways to stop the Event Handler, each method corresponding to how it was started.

- If the Event Handler was started as an Management Server service, then it will be stopped when the Management Server stops. Use the OEMCTRL utility to stop the Management Server.

```
oemctrl stop oms <oms-user>/<oms-password>
```

- If the Event Handler was started as a client application using the `oemctrl` utility, then stop the utility using `Ctrl-C`.

Troubleshooting Tips

- When starting the Event Handler as a client application, make sure that the Management Server is up and running when you do so.
- If the logger/executor doesn't seem to respond to events, check the following:
 1. Make sure you have the appropriate entries to start the logger and executor in the omsconfig.properties file. These are:

```
/com/oracle/sysman/em/eventHandler/handlers/eventlogger/entrypoint=oracle.sysman.vda.vdae.VdaeEventLogger
```

```
/com/oracle/sysman/em/eventHandler/handlers/eventlogger/processOnStartup=true
```

```
/com/oracle/sysman/em/eventHandler/handlers/commandexecutor/entrypoint=oracle.sysman.vda.vdae.VdaeCommandExecutor
```

```
/com/oracle/sysman/em/eventHandler/handlers/commandexecutor/processOnStartup=true
```

2. Make sure the Management Server administrator used to start the Event Handler receives the event notification via the EM console. The Management Server administrator needs to have at least VIEW permissions on the event.
3. Make sure you have set a filter that allows events to pass through to them. Try initially setting a filter that allows **all** events to pass through.

Example:

```
/com/oracle/sysman/em/eventHandler/global_filters/allEvents/eventname=*
```

4. Check the syntax of any templates associated with the logger/executor. Try using a simple template or any of the samples given below.
5. When using the logger on unix, make sure you have the appropriate permissions to create the log file. It will use the permissions of the operating system user who started the Event Handler. Try creating a file in the directory where the logger's log file should have been created.
6. When using the command executor on unix, make sure you have the appropriate permissions to execute the command. It will use the permissions of the operating system user who started the Event Handler. Try running the command directly on the operating system to make sure it works.

7. If you started the Event Handler as an Management Server service, make sure the service is still up. Check the screen where you started the Management Server and watch for messages that indicate the Event Handler is no longer running. If on a UNIX system, check to see if the service is running.
 8. Make sure there are no linefeed characters within any single entry in the *omsconfig.properties* and *OMSExternalRegistry.registry* file.
 9. 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)
- The Event Handler is an English-only release.

Tracing

Should come as a last resort to find out errors that have occurred with the Event Handler. Trace information can be directed either to the screen or to a log file in the `%ORACLE_HOME%\sysman\log` directory. To enable logging or tracing, turn on the options specified in the *omsconfig.properties* file and restart the Event Handler.

```
/com/oracle/sysman/em/eventHandler/tracing_enabled = true or false  
/com/oracle/sysman/em/eventHandler/logging_enabled = true or false
```

The trace information generated can either be generated to a file or to the screen. If `tracing_enabled` is set to true, then any trace information will be sent to the screen. If `logging_enabled` is set to true, then trace information will be generated in the log file. However, logging takes precedence, i.e. if `logging_enabled` is true, then trace information will always be sent to the log file.

These are the various combinations and their output:

```
tracing_enabled = false  
logging_enabled = false  
No trace on screen or in the log file.
```

```
tracing_enabled = false  
logging_enabled = true  
Trace information will be generated to the log file.
```

```
tracing_enabled = true
logging_enabled = false
Trace information will be generated to the screen.
```

Note: If the Event Handler was started as an external service, then the trace will be written to the file specified by the logfile_name entry and in the directory specified by outfile_destination entry in the *OMSExternalRegistry.registry* file

```
tracing_enabled = true
logging_enabled = true
Trace information will be generated to the log file.
```

Sample Filters and Templates:

The following are useful filters and templates. Some of these entries can be found in the %ORACLE_HOME%/sysman/admin/omsconfig.properties.EventHandler 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 on all events except "cputest" to both event logger and command executor

```
/com/oracle/sysman/em/eventHandler/global_filters/not-cputest/eventname =
!cputest
```

Pass on all events named cputest that did not originate on the node smptest16 to the Event Logger:

```
/com/oracle/sysman/em/eventHandler/eventlogger_
filters/pass-some-cputest/eventname=cputest
/com/oracle/sysman/em/eventHandler/eventlogger_
filters/pass-some-cputest/node=!smptest16
```

Pass all events except cputest to the Command Executor adapter:

```
/com/oracle/sysman/em/eventHandler/commandexecutor_
filters/pass-no-cputest/eventname=!cputest
```

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%
```

A sample log entry using this template would look like:

```
cputest fired on smptest16: Target smptest16: Output Cpu usage is high.
Severity: Alert Time: 10-21-1999 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 Solaris:

```
/com/oracle/sysman/em/commandexecutor/templates/allEvents/command=xterm
-display hqsun1:0 -e telnet
```

Here, whenever any event fires, a telnet session is opened.

Known Issues

- The Event Handler may not start if there are any carriage return characters within any of the required entries.

- If you specify incomplete template entries, for example:

```
/com/oracle/sysman/em/eventlogger/templates/allEvents/eventname=*
```

but do not specify the message format to use, for example:

```
/com/oracle/sysman/em/eventlogger/templates/allEvents/message= ...
```

You may see a `NullPointerException`.

- The default message format used by the Event Logger displays a numeric value indicating the severity. For example:

```
cputest;6;10-27-99 12:43:51 PM;;25
```

To obtain the actual severity level, define a message template and use the `%severity%` placeholder.

- For slower machines, starting the Event Handler as an External service may generate an error message stating the Event Handler is no longer available. This is due to discrepancies in synchronization speed. Nothing is wrong with the Event Handler. Wait a few minutes to allow the Event Handler to initialize.

Differences between Solaris and Windows NT

For Event Handler entries that require a directory path, use the syntax that is appropriate to the operating system. For example

Solaris:

```
/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
```

Index

Symbols

, 1-18

A

acknowledging events

Event menu, 5-20

Event window, 5-19

Active Jobs page, 4-7

add a group to a map

Map menu, 3-3

additional services folders

Navigator tree, 2-3

administration task

Navigator objects, 2-10

administrator

adding, editing, and deleting, 1-7

managing, 1-7

preferences, 1-10

reassigning objects when deleting, 1-9

super, 1-7

Administrator Notification

email and paging, 5-34

Advanced Mode

View menu, 1-5

agent

down, 5-35

intelligent, 6-2

process, 6-2

UpDown Node event, 5-35

apply (register) an event set, 5-32

Archive Log Operations in Storage Management

features, 7-15

archived redo log, 8-8

audience for this guide, xv

Available Destinations list

Job system, 4-15

B

Background images

adding to groups, 3-5

backup

cumulative, 8-11

full, 8-11

incremental, 8-11

non-cumulative, 8-11

tablespace, 8-17

backup and recovery, managing, 8-1

Backup Configuration Library, 8-30

backup configuration, creating, 8-20

backup sets, 8-21

Backup Wizard, 8-3

Archive Logs, 8-10

Backup Frequency, 8-5

Backup Options, 8-11

Configuration, 8-12

Datafiles, 8-9

Introduction, 8-3

Job Information, 8-14

Schedule, 8-12

Strategy Choice, 8-4

Summary, 8-14

Tablespaces, 8-9

blackouts, paging/email, 2-10

Broadcast Message task, 4-27

C

- Catalog Maintenance Wizard, 8-25
 - Configuration Page, 8-27
 - Introduction Page, 8-25
 - Multiple Destinations Page, 8-27
 - Operation Choice Page, 8-26
 - Schedule, 8-28
 - Summary, 8-29
- catalog, re-synchronizing, 8-26
- catrman Script, 8-35
- channel, specifying, 8-20
- clone (of a database object), 7-10
- Command Executor Templates, A-8
- communication daemon
 - populating Navigator, 2-2
- Components, 1-2
 - Event, 1-3
 - Group, 1-2
 - Job, 1-3
 - Navigator, 1-2
- composite job, 4-3
- Console
 - customize layout, 1-5
 - File menu, 1-5
 - Help menu, 1-6
 - menus, 1-4
 - user preferences, 1-18
 - using menus, 1-4
- console
 - list of tasks that can be performed from, xiv
- Console panes
 - Groups, 3-1
 - Navigator, 2-1
- context menu, 2-5
- context-sensitive menus
 - Navigator, 1-4
 - using in the Navigator, 1-4
- control file, 8-23
- Controlfile Operations of Storage Management
 - features, 7-14
- copying tree objects
 - Navigator, 2-10
- create
 - event set, 5-20

- job, 4-10
 - or modify a job, 4-11
 - or modify an event set, 5-24
 - specified object type in the Navigator, 2-5
- Create Backup Configuration
 - Backup Parameters Page, 8-24
 - Channels Page, 8-20
 - General Page, 8-20
 - Preferred Credentials Page, 8-24
 - Recovery Catalog Page, 8-22
- creating
 - tablespace, 8-34
 - user, 8-34
- cumulative backup, 8-11
- customize and manage network views
 - Map system, 3-3

D

- Data Gatherer events, 5-36
- database
 - connection made in tree, 2-3
 - recovery, 8-16
 - registering, 8-26
 - resetting, 8-26
- database administration applications
 - DBA, 1-6
- Database Fault Management Events, 5-35
- database instance property sheet
 - accessing from a map, 3-8
- Database Operations, 7-5
- Databases folder
 - Navigator tree, 2-2
- datafile
 - backup, 8-8, 8-9
 - recovery, 8-16, 8-18
- datafile operations of Storage Management
 - features, 7-14
- DBA Management Pack, 7-1
- delete
 - group, 3-3
- discovering services
 - Navigator, 2-7
 - SQL*Net network, 2-2
- discovery

- problems, 2-9
- discovery status, 2-9
- drag and drop
 - Navigator, 2-10
 - Navigator objects, 2-2

E

- edit
 - group, 3-3
- EMS
 - Event Management system, 5-1
- event
 - set, 5-2
 - unsolicited, 5-8, 5-28
- Event Filters, A-4
- Event Handler, A-1
- Event History page
 - Event window, 5-19
- Event Logger Templates, A-6
- Event Management, 5-1
 - categories of events, 5-6
 - creating an event set, 5-24
 - Enterprise Manager, 1-3
 - fault management, 5-6
 - introduction, 5-1
 - menu, 5-20
 - modifying a registered event set, 5-32
 - modifying an event set, 5-24
 - Oracle events, 5-34
 - parameter settings, 5-29
 - performance management, 5-7
 - process, 5-2
 - registering an event set, 5-32
 - resource management, 5-7
 - space management, 5-7
 - types of events, 5-6
 - viewing a registered event set, 5-32
 - viewing an event set, 5-21
 - window, 5-17
- Event Notification Filter, 1-14
- Event Set Library page
 - Event window, 5-21
 - predefined event sets, 5-21
- Event Set Management property sheet

- creating an event, 5-20
- modifying a new event, 5-20
- example
 - creating an event set, 5-32
 - group, 3-11
 - Job Scheduling, 4-21
 - registering an event, 5-32
- expanding
 - map objects, 3-8
- Expert Wizard, 7-16

F

- Fault Management events, 5-6
- fixit job, 4-3, 5-1
 - creating, 4-14
 - events, 5-27
 - removing, 5-27
 - submitting, 5-27
 - with events, 5-1
- full backup, 8-11

G

- General page
 - job, 4-13
- Group
 - pane, 3-2
- group, 3-4
 - add or delete, 3-9
 - adding background images, 3-5
 - create, 3-4
 - edit, 3-3
 - example, 3-11
 - job destination, 4-15
 - modify, 3-10
 - modifying properties, 3-10
- Group system, 1-2
- Groups folder
 - Navigator tree, 2-3

H

- history
 - Command History Button, SQL*Plus

Worksheet, 7-15

I

Import Wizard, 7-16
incremental backup, 8-11
Input pane, SQL*Plus Worksheet, 7-15
installing
 Intelligent Agent, 6-2
Intelligent Agent, 6-2, 6-3

J

job
 alter, 4-10
 composite, 4-3
 create, 4-10
 manage, 4-10
 modify, 4-10
 view, 4-10
Job History page, 4-8, 4-11
Job Library, 4-11
Job Notification Filter, 1-14
job output
 Progress Page, 4-20
Job Scheduling, 4-1
 credential preference, 4-4, 4-17
 destination requirements, 4-15
 Enterprise Manager, 1-3
 example, 4-21
 Failed to Authenticate User error, 4-4
 General page, 4-13
 job destination, 4-13
 job destinations, 4-15
 job name, 4-13
 job tasks, 4-15
 job type, 4-13
 managing jobs, 4-5
 menu, 4-10
 saving history, 4-11
 scheduling jobs, 4-5
 scheduling process, 4-2
 task parameters, 4-17
 window, 4-6
Job Tasks

wizards, 4-24
jobs, cancelling, 4-5

L

launching a database tool
 from a map, 3-8
 from the Console, 2-5
 from the Navigator, 1-2
 Navigator, 2-4
 with the right-mouse button, 2-4
levels of permission, 1-17
limits, backup, 8-22
Listener Fault Management Events, 5-36
Listeners folder
 Navigator tree, 2-3
Load Wizard, 7-16
LRS
 See Log roll-forward server (LRS), A-1

M

Management Information Base (MIB)
 with SNMP, 6-4
manipulation of the objects
 in the Navigator, 2-9
Map, 3-1
 bitmap background, 3-6
 create, 3-7
 creating with dragging and dropping, 3-7
 customized views, 3-1
 expanding objects, 3-8
 status of objects, 3-7
 user-defined views, 3-4
menus
 Console, 1-4
 File menu, 1-5
 Help menu, 1-6
 Job, 4-10
 View, 1-5
modify
 event set, 5-20
 group, 3-3
multiple tasks
 submitting a job, 4-12

N

- Navigator, 2-1
 - Enterprise Manager, 1-2
 - manipulating objects, 2-9
 - menu, 2-5
 - new view, 2-6
 - objects, 1-2
 - splitting the tree, 2-4, 2-6
 - window, 2-2
- network
 - container, 2-2
 - encryption, 6-3
- new view
 - Navigator tree, 2-6
- Node Fault Management Events, 5-35
- Nodes folder
 - Navigator tree, 2-3
- non-cumulative backup, 8-11
- Notification Filters, 1-14

O

- object
 - creating, 7-10
 - editing, 7-10
- Oracle DB UpDown
 - event profile, 5-22
- Oracle events, 5-34
- Oracle Host UpDown
 - event profile, 5-22
- Oracle Listener UpDown
 - event profile, 5-22
- Oracle Management Server, 6-4
- Oracle Multi-Protocol Interchange
 - and Oracle Management Server, 6-5
- Oracle Remote Operations
 - and Oracle Management Server, 6-5
- Oracle Secure Network Services (SNS)
 - and Oracle Management Server, 6-5
- Oracle Security Management, 7-10
- Oracle SQL*Plus Worksheet, 7-15
- Oracle8i objects, 7-8
- OraTcl
 - job scripts, 4-2

- Output Dialog box
 - Job Scheduling, 4-20
- Output pane, SQL*Plus Worksheet, 7-15
- Outstanding Events page
 - Event system, 5-18

P

- Paging/Email Blackout, 2-10
- parameters
 - events, 5-29
 - jobs, 4-17
- Parameters Page
 - job, 4-17
- password management, 7-12
- Performance Management events, 5-6
- Ping Agent, 2-6
- ping agent menu item, 2-9
- PL/SQL
 - using with Oracle SQL*Plus Worksheet, 7-15
- platform-specific installation documentation, 6-2
- Pre-defined Profiles
 - Event system, 5-22
- Preferred credentials
 - , 1-18
- preferred credentials
 - connecting in a map, 3-8
 - Console, 1-18
- problems
 - Navigator discovery, 2-9
- profile operations, Security Management, 7-12
- Progress page
 - job, 4-20
- purpose of this guide, xiv

R

- recovery
 - database, 8-16
 - datafile, 8-16, 8-18
 - tablespace, 8-16
- recovery catalog, 8-23, 8-26
- registering, 8-36
- setting up and registering, 8-33
- Recovery Wizard, 8-15

- Configuration Page, 8-19
- Datafiles Page, 8-18
- Introduction Page, 8-15
- Recovery Selection, 8-16
- Rename Page, 8-18
- Restore Until, 8-17
- Summary, 8-19
- Tablespaces Page, 8-17
- RECOVERY_CATALOG_OWNER role, 8-34
- redo log group operations of Storage Management
 - features, 7-15
- Refreshing Discovery
 - Refresh Topology Menu Choice, 2-8
- register database, 8-26
- registered destination
 - map objects, 3-7
- registering events, 5-25
- Registrations page
 - Event window, 5-19
- related publications, xx
- Remote SYSDBA Access
 - SYSDBA access, 8-31
- removing
 - a split page in the Navigator, 2-6
 - job with Remove menu item, 4-11
 - problems with a job, 4-11
 - problems with an event, 5-21
 - registered events, 5-21
- reset database, 8-26
- RESETLOGS option, 8-26
- Resource Management events, 5-6
- right mouse button
 - used in the Navigator, 1-4
- right-mouse button
 - Related Tools menu, 2-4
- role operations, Security Management, 7-12
- rollback segment operations of Storage Management
 - features, 7-14
- rollback segments, 7-13
- Run DBA Script task, 4-25
- Run OS Command task, 4-28
- Run SQL*Plus task, 4-24
- Run Tcl Script task, 4-28

S

- sample images, 3-6
- saving
 - job history, 4-10, 4-11
- schedule
 - and manage jobs, 4-5
 - execution of a job, 4-17
 - job, 4-5
- Schema Management, 7-8
- Security Management, 7-10
- Send Us Your Comments, xi
- Service Discovery
 - Navigator menu option, 2-6
- Service Name, 1-18
- Service Type, 1-18
- services and agents, 6-1
- setting a bitmap background
 - Map menu, 3-6
- short-cut menu, 1-4
- Shutdown Database task, 4-25
- Shutdown Listener task, 4-30
- signal flag
 - on map objects, 3-7
- SMTP mail
 - setup, 1-5
- SNMP
 - and agents, 6-4
- Space Management event, 5-6
- specifying service names, 1-19
- split (new) view
 - Navigator, 2-6
- splitting the Navigator tree, 2-4, 2-6
 - new view, 2-6
- SQL
 - Oracle SQL*Plus Worksheet, 7-15
- SQL error
 - job task, 4-25
- SQL*Net network
 - services, 2-2
- SQL*Plus Worksheet, 7-15
- Startup Database task, 4-26
- Startup Listener task, 4-30
- status
 - of a map object, 3-7

- of a registered event set, 5-3
- Storage Management, 7-13
- storage parameters, setting for backup, 8-24
- summary of the pending jobs
 - Active Jobs page, 4-7
- Super Administrator, 1-7
- SYSDBA
 - role, 3-8
- SYSOPER
 - role, 3-8

T

- tablespace
 - backup, 8-8, 8-9, 8-17
 - recovery, 8-16
- tablespace operations of Storage Management
 - features, 7-14
- tablespace, creating, 8-34
- tablespaces, 7-13
- task parameters
 - Job system, 4-17
- Tcl
 - job scripts, 4-2
 - script examples, 4-29
 - scripts executed by agents, 6-3
- third-party events
 - creating, 5-8
 - with job task, 4-29
- tnsnames.ora, 2-9
- Tools menu, 1-6
- Transparent Network Substrate (TNS)
 - and Oracle Management Server, 6-5

U

- unsolicited event, 5-8, 5-28
- unsolicited events, 5-6
- UpDown (Database) event, 5-35
- UpDown (Listener) event, 5-36
- UpDown (Node) event, 5-35
- user operations, Security Management, 7-11
- user preferences
 - jobs, 4-4, 4-17
 - when running jobs, 4-4, 4-17

- user, creating, 8-34
- user-defined groups
 - in the Navigator, 2-5
 - job destinations, 4-15
 - Navigator tree, 2-5
- user-defined map views
 - creating, 3-4

W

- Wizard Job Tasks, 4-24
- wizards
 - Expert Wizard, 7-16
 - Import Wizard, 7-16
 - Load Wizard, 7-16
- worksheet, Oracle SQL*Plus Worksheet, 7-15

