

Oracle[®] Enterprise Manager

Administrator's Guide

Release 1.6.0

June, 1998

Part No. A63731-01

Oracle Enterprise Manager Administrator's Guide

Part No. A63731-01

Release 1.6.0

Copyright © 1993, 1997, Oracle Corporation. All rights reserved.

Contributors: Robert Abbott, Eric Belden, Diarmuid Cawley, Larry Chen, Valerie Kane, John Kennedy, Keri Kruse, Dennis Lee, Eric Lee, Priscilla Lee, Dimitris Nakos, Mark Osborn, Jay Rossiter, Phil Schwartz, Vipul Shah, Michael Stern, Dave Stowell, Kirill Yurchenko, Kim Zabora, and the entire SMP team.

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

This Program contains proprietary information of Oracle Corporation; it is provided under a license agreement containing restrictions on use and disclosure and is also protected by copyright patent and other intellectual property law. Reverse engineering of the software 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.

If this Program is delivered to a U.S. Government Agency of the Department of Defense, then it is delivered with Restricted Rights and the following legend is applicable:

Restricted Rights Legend Programs delivered subject to the DOD FAR Supplement are 'commercial computer software' and use, duplication and disclosure of the Programs 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 52.227-14, Rights in Data -- General, including Alternate III (June 1987). Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065.

Oracle, Oracle Enterprise Manager, Oracle Enterprise Manager Performance Pack, Oracle Server Manager, Oracle Trace, and SQL*Net are registered trademarks of Oracle Corporation. Net8, Oracle Expert, Oracle7, Oracle8, and PL/SQL are trademarks of Oracle Corporation. Windows, Windows NT, Visual C++, and OLE are trademarks of Microsoft Corporation. All trade names referenced are the service mark, trademark, or registered trademark of the respective manufacturer.

All other products or company names are used for identification purposes only, and may be trademarks of their respective owners.

OraTel - Copyright 1993 Tom Poindexter and US West Advanced Technologies, Inc. Permission to use, copy, modify, and distribute the software and its documentation for any purpose and without fee is hereby granted, provided that the copyright notice appear in all copies. Tom Poindexter and U S WEST make no representations about the suitability of this software for any purpose. It is provided as is without express or implied warranty. By use of this software the user agrees to indemnify and hold harmless Tom Poindexter and U S WEST from any claims or liability for loss arising out of such use.

Contents

Send Us Your Comments	xxi
Preface.....	xxiii
Part I Oracle Enterprise Manager and Console	
1 Overview of Oracle Enterprise Manager	
Introducing Enterprise Manager	1-2
Enterprise Manager's Components.....	1-2
Console.....	1-3
Navigator.....	1-4
Map.....	1-4
Job Scheduling	1-4
Event Management	1-5
Repository.....	1-5
Communication Daemon	1-6
Discovery Cache	1-6
Intelligent Agents	1-6
Administrator Toolbar.....	1-7
Database Administration Tools.....	1-7
Server Manager Line Mode.....	1-8
Online Help	1-8
Value Added Products.....	1-9
Console Menus.....	1-10

Context-Sensitive Menus	1-10
File Menu	1-11
View Menu.....	1-11
Tools Menu	1-12
Applications	1-12
Performance Pack	1-12
Daemon Manager	1-12
Help Menu	1-12
Toolbars and Launch Palettes	1-13
Toolbars.....	1-13
Launch Palettes	1-13
Oracle Administrator Toolbar	1-14
Toolbar Menu	1-14
Customize	1-14
Databases	1-14
Database Status	1-15
Application Bar/Float Bar	1-15
Help	1-15
Exit	1-15
ToolBar Customize: General Page.....	1-16
ToolBar Customize: Applications Page.....	1-17
ToolBar Customize: Database Page	1-18
Getting Started	1-19
Launching Enterprise Manager	1-19
Connecting to a Repository	1-19
Privileged Operations	1-20
Repository Manager Wizard.....	1-21
Creating, Validating, or Dropping a Repository.....	1-21
Saving a List of Discovered Nodes	1-21
Logging into a Repository	1-21
Launching Tools.....	1-22
Administration of a Remote Database.....	1-23
Connecting to an Instance	1-23
Multiple Connections.....	1-25
Console User Preferences	1-25

Setting User Preferences	1-26
--------------------------------	------

2 Navigator

Navigator Window	2-2
Expanding Objects in the Navigator.....	2-2
Populating the Navigator Tree.....	2-3
Viewing Specific Tree Objects.....	2-4
Launching Tools	2-4
User-defined Groups.....	2-4
Creating Maps.....	2-5
Navigator Menu	2-5
Navigator Context-Sensitive Menus	2-7
Discovering Services	2-8
Discover New Services Wizard	2-8
Refresh Services Wizard	2-11
Manually Define Services Wizard.....	2-12
Status Option.....	2-13
Update Network Configuration Option	2-14
Destination Validation Dialog	2-14
Manipulating Objects in the Navigator	2-15
Expanding and Collapsing Objects.....	2-15
Administering Objects	2-15
Copying Tree Objects.....	2-16
Example: Using the Navigator	2-17
Filtering Folders.....	2-18
Keystroke Shortcuts	2-19

3 Map

Map Window	3-2
Manipulating Map Views.....	3-2
User-Defined Groups.....	3-3
Status Monitors	3-4
Expanding Objects.....	3-4
Groups.....	3-5
Nodes	3-5

Databases	3-5
Launching Tools from a Map.....	3-5
Map Menu	3-6
Creating, Modifying, and Removing Maps	3-7
Creating a Map.....	3-7
Adding Objects in a Map.....	3-7
Deleting Objects in a Map.....	3-7
Modifying a Map	3-8
Removing a Map.....	3-8
Example: Creating a Map	3-9
Creating, Modifying, and Removing Groups	3-10
Creating a Group in the Current Map or Group.....	3-10
Creating a Group in the Navigator Tree	3-10
Adding Objects in a Group	3-11
Deleting Objects in a Group	3-11
Modifying a Group.....	3-12
Removing Groups.....	3-12
Example: Creating a Group.....	3-13

4 Job Scheduling

Job Scheduling Process	4-2
Job Tasks	4-2
Job Credentials	4-3
Submitting Jobs	4-4
Job Window.....	4-5
Active Jobs	4-6
Job History	4-7
Job Library	4-7
Job Menu	4-8
Creating, Modifying, or Viewing a Job	4-9
Creating a New Job.....	4-9
Modifying a Saved Job.....	4-10
Viewing Job Details	4-10
Job General Page.....	4-10
Job Task Page.....	4-11

Job Parameters Page.....	4-12
Job Schedule Page.....	4-13
Job Progress Page.....	4-13
Output Dialog Box	4-15
Example: Creating a Job	4-15
Oracle Job Tasks	4-17
Oracle Database Tasks	4-18
Operating System or Host Tasks.....	4-21
Tcl Script Examples.....	4-23
Listener Tasks.....	4-24

5 Event Management

Event Management Process	5-2
Determining Administrators	5-2
Using Event Sets	5-3
Creating Events.....	5-3
Registering Events.....	5-3
Event Occurrences.....	5-4
Event Notifications.....	5-4
Interpreting Events.....	5-5
Correcting Problems	5-5
Event Categories and Types	5-6
Fault Management Events.....	5-6
Space Management Events.....	5-7
Resource Management Events.....	5-7
Performance Management Events	5-7
Event Window	5-8
Outstanding Events Page	5-9
Acknowledging Events.....	5-9
Event History Page.....	5-10
Viewing Acknowledgments	5-10
Saving Event History	5-10
Registrations Page	5-10
Modify Registrations Property Sheet	5-11
Event Set Library Page.....	5-11

Event Set Property Sheet	5-11
Oracle Event Sets	5-12
Event Menu	5-12
Context-Sensitive Menus	5-14
Creating or Modifying an Event Set	5-14
Event Set Management General Page	5-15
Event Set Management Events Page	5-15
Event Set Management Parameters Page	5-16
Choose Fixit Job	5-16
Frequency	5-17
Parameters	5-17
Example: Creating an Event Set	5-17
Creating, Modifying, or Viewing an Event Registration	5-19
Registering an Event Set	5-19
Modify or Viewing An Event Set	5-20
Registrations General Page	5-20
Registrations Notification Page	5-21
Registrations Status Page	5-21
Example: Registering an Event Set	5-22
Acknowledging an Event	5-23
Acknowledge Event General Page	5-23
Acknowledge Event Notification Page	5-24
Managing Administrators	5-24
Adding New Administrators	5-24
Modifying Administrators	5-24
Deleting Administrators	5-25
Administrator General Page	5-25
Paging Information:	5-25
EMailing Information:	5-26
Administrator Systems Page	5-26
Administrator Availability Page	5-27
Configuring Mail Services	5-27
MAPI Setup	5-27
Configuring Paging Services	5-28
Adding New Paging Services	5-28

Modifying Paging Services	5-29
Deleting Paging Services	5-29
Paging Services Configuration Dialog	5-29
Oracle UpDown Events	5-31
Database Fault Management Events	5-31
Node Fault Management Events.....	5-32
Listener Fault Management Events	5-32
Numeric Pager Event Ids	5-33
Database events	5-33
Node events.....	5-35
Listener events:	5-35

6 Agents and Communication Daemon

Agents	6-2
Characteristics	6-2
Event and Job Support	6-3
Network Encryption	6-3
SNMP Support	6-3
Communication Daemon	6-4
Message Queues	6-4
Failed Queue	6-5
Notification Queue	6-5
Connection Cache.....	6-5
Daemon Manager	6-6
Agent Pending Operations.....	6-6
Application Pending Notifications	6-6
Monitored Nodes.....	6-7
Daemon Configuration Parameters.....	6-7
View and Updating Parameters	6-9
Daemon Manager Menu.....	6-11

Part II The Database Administration Tools

7 Overview of the Database Tools

Introduction	7-2
Application Objects	7-3
Types of Objects.....	7-3
Manipulating Tree List Objects	7-4
Backup Manager	7-5
Data Manager	7-6
Instance Manager	7-6
Schema Manager	7-7
Security Manager	7-8
Software Manager	7-8
Storage Manager	7-8
SQL Worksheet	7-9
Application Menus	7-10
File Menu	7-10
View Menu.....	7-11
Log Menu.....	7-12
Help Menu.....	7-12
Additional Menus.....	7-13
Context-Sensitive Menus.....	7-13
Multi-Column Lists	7-14
Dialog Boxes	7-15
Property Sheets	7-17
Spreadsheets.....	7-18
Property Sheet Buttons	7-19
Starting an Application	7-20

8 Managing Database Storage

Starting Storage Manager	8-2
Storage Menus	8-4
Tablespace Menu	8-4
Datafile Menu.....	8-5
Rollback Menu	8-6
Tablespaces Folder	8-7
Tablespaces Multi-Column List.....	8-7

Creating A Tablespace	8-7
Showing Tablespace Dependencies	8-11
Editing a Tablespace	8-12
Dropping a Tablespace	8-12
Adding a Datafile to a Tablespace	8-13
Adding a Rollback Segment to a Tablespace	8-13
Changing a Tablespace to Online or Offline Status.....	8-14
Making a Tablespace Read-Only.....	8-14
Making a Tablespace Writeable.....	8-15
Datafiles Folder	8-15
Datafiles Multi-Column List	8-15
Creating a Datafile.....	8-16
Creating a Datafile Like an Existing Datafile	8-18
Editing an Existing Datafile	8-18
Placing a Datafile Online	8-18
Taking a Datafile Offline	8-18
Rollback Segments Folder	8-19
Rollback Multi-Column List	8-19
Creating a Rollback Segment.....	8-19
Altering a Rollback Segment.....	8-22
Dropping a Rollback Segment	8-22
Shrinking a Rollback Segment.....	8-22
Changing a Rollback Segment to Online or Offline Status	8-23

9 Managing Database Security

Starting Security Manager	9-2
Security Manager Menus	9-4
User Menu	9-4
Role Menu.....	9-5
Profile Menu.....	9-6
Security Manager Objects and Folders	9-7
Users Folder	9-7
Users Multi-Column List.....	9-8
Creating a User	9-8
Create Like User	9-14

Altering a User	9-14
Removing a User.....	9-15
Adding Privileges or Roles to Users	9-15
Removing Privileges or Roles from a User	9-17
Privileges Multi-Column Lists.....	9-17
Roles Granted.....	9-17
System Privileges Granted	9-18
Object Privileges Granted.....	9-18
Roles Folder	9-18
Roles Multi-Column List	9-19
Creating a Role	9-19
Create Like Role	9-20
Modifying a Role	9-21
Removing a Role	9-21
Adding Privileges or Roles to Roles.....	9-22
Removing Privileges or Roles from a Role.....	9-23
Profiles Folder	9-24
Profile Multi-Column Lists.....	9-24
Creating a Profile	9-25
Create Like Profile	9-28
Altering a Profile.....	9-28
Showing Profile Dependents and Dependencies	9-29
Removing a Profile	9-29
Assigning a Profile to Users.....	9-29

10 Managing Instances and Sessions

Starting Instance Manager	10-2
Instance Manager Menus	10-3
Database Menu.....	10-4
Sessions Menu.....	10-5
Transactions Menu	10-5
Configuration Menu.....	10-6
Database Object	10-7
Starting Up a Database.....	10-9
Shutting Down a Database	10-10

Mounting or Opening a Database	10-10
Initialization Parameters Folder	10-11
Initialization Parameters Property Sheet	10-11
Editing Initialization Parameters	10-13
Resetting Edits.....	10-13
Applying Edits.....	10-13
Stored Configurations	10-14
Stored Configuration Folder	10-14
Creating a Stored Configuration	10-14
Stored Configuration Property Sheet	10-14
Editing a Stored Configuration	10-15
Deleting Parameters from a Stored Configuration.....	10-15
Adding a Parameter to a Stored Configuration	10-15
Creating a File from a Stored Configuration	10-16
Deleting a Stored Configuration	10-16
Sessions Folder	10-17
Sessions Multi-column List	10-17
Disconnecting a User's Session.....	10-18
Restricting or Allowing Sessions.....	10-19
In-Doubt Transactions Folder	10-19
In-Doubt Transactions Multi-column List	10-19
In-Doubt Transactions Property Sheet	10-20
Forcing a Commit or Rollback.....	10-21

11 Managing Schema Objects

Starting Schema Manager	11-3
Schema Menus	11-4
Schema Folders	11-6
Schema Property Sheets	11-6
Explicit	11-7
Auto Calculation.....	11-8
Array Types Folder	11-10
Array Types Multi-column List.....	11-10
Creating an Array Type.....	11-11
Create Like Array Type	11-12

Editing an Array Type	11-12
Clusters Folder	11-12
Clusters Multi-column List.....	11-13
Creating a Cluster	11-13
Column Details Dialog Box.....	11-15
Create Like Cluster	11-16
Editing a Cluster	11-16
Database Links Folder	11-16
Database Links Multi-column List	11-17
Creating a Database Link.....	11-17
Create Like Database Link.....	11-19
Viewing a Database Link.....	11-19
Functions Folder	11-19
Functions Multi-column List.....	11-20
Creating a Function	11-20
Create Like Function	11-21
Altering a Function.....	11-22
Indexes Folder	11-22
Indexes Multi-column List	11-23
Creating an Index.....	11-24
Create Like Index.....	11-26
Altering an Index.....	11-26
Object Types Folder	11-26
Object Types Multi-column list	11-27
Creating an Object Type	11-27
Adding Attributes to Object Types.....	11-29
Editing or Removing Object Attributes.....	11-30
Adding Methods to Object Types	11-30
Editing or Removing Object Methods.....	11-32
Package Bodies Folder	11-33
Package Bodies Multi-column List.....	11-33
Creating a Package Body	11-34
Create Like Package Body	11-35
Editing a Package Body	11-35
Packages Folder	11-36

Packages Multi-column List.....	11-36
Creating a Package	11-36
Create Like Package	11-37
Editing a Package	11-38
Procedures Folder	11-38
Procedures Multi-column List.....	11-38
Creating a Procedure.....	11-39
Procedures Property Sheet: General Page	11-39
Create Like Procedure.....	11-40
Altering a Procedure	11-40
Queue Tables	11-40
Queues Multi-column List	11-41
Creating a Queue Table	11-41
Create Like Queue Table	11-44
Altering a Queue Table.....	11-44
Creating a Queue.....	11-45
Create Like Queue	11-47
Altering a Queue	11-48
Refresh Groups Folder.....	11-48
Refresh Groups Multi-column List	11-49
Creating a Refresh Group.....	11-49
Create Like Refresh Group.....	11-51
Altering a Refresh Group	11-51
Sequences Folder	11-51
Sequences Multi-column List.....	11-52
Creating a Sequence	11-53
Create Like Sequence	11-54
Altering a Sequence.....	11-55
Snapshot Logs Folder.....	11-55
Snapshot Logs Multi-column List.....	11-56
Creating a Snapshot Log.....	11-56
Create Like Snapshot Log.....	11-57
Altering a Snapshot Log.....	11-58
Snapshots Folder.....	11-59
Snapshots Multi-column List.....	11-59

Creating a Snapshot.....	11-60
Create Like Snapshot.....	11-62
Editing a Snapshot.....	11-63
Synonyms Folder	11-64
Synonyms Multi-column List.....	11-64
Creating a Synonym.....	11-64
Create Like Synonym.....	11-65
Viewing a Synonym	11-66
Table Types Folder	11-66
Table Types Multi-column List.....	11-66
Creating a Table Type	11-66
Create Like Table Type	11-67
Editing a Table Type	11-68
Tables Folder	11-68
Tables Multi-column List.....	11-68
Creating a Table	11-69
Dropping a Constraint on a Table.....	11-74
Removing a Constraint Definition (Create Mode Only).....	11-75
Create Like Table	11-75
Altering a Table.....	11-75
Triggers Folder	11-75
Triggers Multi-column List	11-76
Creating a Trigger.....	11-76
Create Like Trigger.....	11-79
Altering a Trigger	11-79
Views Folder	11-79
Views Multi-column List	11-80
Creating a View.....	11-80
Create Like View.....	11-82
Altering a View	11-82

12 Managing Backup and Recovery

Starting Backup Manager	12-2
Backup Manager: Oracle8 Recovery Manager Subsystem	12-4
Recovery Catalog.....	12-4

Registering a Database with the Recovery Catalog.....	12-5
Oracle8 Recovery Manager Subsystem User Interface	12-6
Oracle8 Recovery Manager Menu Options	12-6
Tree List Objects and Containers	12-10
Tablespace Containers, Multi-column Lists and Property Sheets.....	12-12
Multi-column List for the Tablespaces Container	12-12
Datafiles Container, Multi-column Lists and Property Sheets	12-14
Backups Page	12-15
Redo Log Groups Container, Multi-column Lists, and Property Sheets	12-16
Archived Logs Container, Property Sheets, and Multi-column lists	12-18
Channels Container.....	12-20
Jobs Object	12-22
Jobs Property Sheets.....	12-22
Oracle8 Recovery Manager Wizards	12-24
Backup Wizard	12-24
Restore Wizard	12-24
Report Wizard.....	12-25
Manual Backup Operations	12-26
Managing Restore and Recover Operations.....	12-31
Maintaining the Recovery Catalog.....	12-34
Database Operations	12-35
Creating a Channel.....	12-35
Creating Stored Scripts	12-37
Backup Manager: Operating System Backup Subsystem.....	12-37
Operating System Backup User Interface	12-37
Backup Manager Menu Options	12-38
Controlfile Object and Property Sheet	12-43
Tablespace Container, Multi-column List, and Property Sheets.....	12-43
Redo Log Group Container, Multi-column List, and Property Sheet	12-45
Redo Log Members Multi-column List	12-46
Redo Log Member Property Sheet.....	12-47
Selected Backup Manager Tasks	12-48
Shutting Down the Database	12-48
Starting Up the Database.....	12-49
Backing Up Tablespaces of the Database.....	12-50

Performing Recovery	12-50
Backup Manager: Enterprise Backup Utility Subsystem	12-52
Creating an Enterprise Backup Utility Catalog.....	12-54
Configuring the PC for Enterprise Backup Utility.....	12-55

13 Using the SQL Worksheet

Overview of the SQL Worksheet	13-2
Starting a SQL Worksheet	13-2
Using the SQL Worksheet	13-2
File Menu	13-5
Edit Menu	13-6
Worksheet Menu.....	13-6
Executing Commands in the SQL Worksheet	13-8
Entering Commands in the Input Pane	13-8
Executing Commands in the Input Pane.....	13-8
Executing Server Manager DBA Commands	13-10
Connecting from a Worksheet	13-10
Performing Recovery from the SQL Worksheet	13-11
Using the Command History	13-12
Retrieving a Selection from the Command History	13-13
Retrieving a Command Using the Worksheet Menu	13-13
Saving Your Work	13-14
Running Scripts from the SQL Worksheet	13-15

14 Managing and Moving Data

Starting Data Manager	14-2
Data Manager Window	14-3
Data Manager Menus	14-4
Data Manager Icons	14-5
Exporting Data to a File	14-5
Importing Data from a File	14-6
Loading Data from an External File	14-7
Data Manager Tasks	14-8

15 Managing Software Distribution

Software Manager Introduction	15-2
The Software Manager Interface	15-3
Navigator Window.....	15-4
Creating a Host Folder	15-5
Viewing the Oracle Home Properties.....	15-6
Viewing the Product Properties.....	15-7
Viewing the Release Properties.....	15-8
Catalog Windows	15-9
Choosing the Catalog View	15-10
Setting Preferences	15-10
Toolbar	15-11
Drag and Drop	15-11
The Software Distribution Process	15-12
Specify Distribution Hosts	15-14
Create Releases.....	15-15
Creating a Release	15-15
Displaying the Properties of a Release.....	15-17
Removing a Release from the Releases Catalog Window	15-17
Create Software Packages.....	15-18
Creating a Software Package	15-19
Deleting a Software Package	15-21
Distribute and Install Software Packages	15-21
The Oracle Software Manager Network	15-22
Setting up a Software Distribution Network.....	15-22
Establishing Communication in the Network.....	15-22
Flow of Information	15-23
Software Manager Tasks	15-24
Determining the Job Name, Destination Type, and Job Destination.....	15-24
Choosing a Task.....	15-26
Setting Parameters.....	15-28
Install Package Task.....	15-29
Deinstall Product Task.....	15-31
Delete Package Task.....	15-33
Distribute Package Task	15-34

Saving and Submitting the Job	15-36
Custom Configurations	15-37
Response File Generator	15-37
Modifying a Response File	15-38
Creating Staging Areas with Oracle File Packager	15-38

Part III Reference

A Using Server Manager in Line Mode

Starting Server Manager in Line Mode	A-2
Using Server Manager in Line Mode	A-2
Entering Server Manager Commands	A-2
Entering SQL or PL/SQL Code	A-2
Running Scripts.....	A-2

B DBA Command Reference

DBA Commands	B-2
@ (at symbol)	B-2
ARCHIVE LOG	B-3
CONNECT.....	B-6
DESCRIBE.....	B-9
DISCONNECT	B-11
EXECUTE.....	B-12
EXIT	B-12
PRINT.....	B-13
RECOVER	B-14
REMARK	B-17
SET	B-18
SHOW	B-25
SHUTDOWN	B-29
SPOOL	B-30
STARTUP.....	B-31
VARIABLE.....	B-34

C Compatibility with SQL*DBA

Obsolete Functionality	C-5
Operational Differences	C-6

Send Us Your Comments

Oracle Enterprise Manager Administrator's Guide Release 1.6.0

Part No. A63731-01

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

- electronic mail - infosmp@us.oracle.com
- 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](#)
- [Your Comments Are Welcome](#)

Purpose of this Guide

This guide describes the Oracle Enterprise Manager™, Oracle's system management Console, common services, and integrated platform graphical tool. Oracle Enterprise Manager (OEMGR) includes database administration tools (DBA), plus it allows you to:

- Monitor network and database objects
- Schedule jobs
- Monitor events
- Manage and troubleshoot networks
- Integrate participating third-party tools

This guide also describes Oracle Server Manager™, the line mode component of Oracle Enterprise Manager. Information in this guide applies to Oracle Enterprise Manager running on the Windows NT and Windows 95 32-bit platforms.

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.

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 tools. 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.

This guide also assumes that you are familiar with the operation of Microsoft Windows. Refer to the Windows documentation for your system, if necessary.

How this Guide Is Organized

This guide is divided into parts and chapters, as described below.

Part I: *Oracle Enterprise Manager and Console*

[Chapter 1, "Overview of Oracle Enterprise Manager"](#)

This chapter describes the overall organization and features of Enterprise Manager and the Console.

[Chapter 2, "Navigator"](#)

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

[Chapter 3, "Map"](#)

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

[Chapter 4, "Job Scheduling"](#)

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

[Chapter 5, "Event Management"](#)

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

[Chapter 6, "Agents and Communication Daemon"](#)

This chapter discusses the Intelligent Agents, the Communication Daemon, and the Daemon Manager.

Part II: *The Database Administration Tools*

[Chapter 7, "Overview of the Database Tools"](#)

This chapter introduces the Database Tools. It also describes the user interface elements used in the Database Tools.

[Chapter 8, "Managing Database Storage"](#)

This chapter explains how to use the Storage Manager to manage the datafiles, tablespaces, and rollback segments in your database.

[Chapter 9, "Managing Database Security"](#)

This chapter describes how to use the Security Manager to manage users, roles, and profiles.

Chapter 10, "Managing Instances and Sessions"

This chapter describes how to use the Instance Manager to start up or shut down your database, and examine the values of the initialization parameters used to start up the instance. It also describes how to manage in-doubt transactions and users' sessions.

Chapter 11, "Managing Schema Objects"

This chapter describes how to use the Schema Manager to examine and modify various objects in your database.

Chapter 12, "Managing Backup and Recovery"

This chapter describes the database Backup Manager.

Chapter 13, "Using the SQL Worksheet"

This chapter describes how to use the SQL Worksheet.

Chapter 14, "Managing and Moving Data"

This chapter describes how to manage and move Oracle database data with the Import, Export, and Load functions.

Chapter 15, "Managing Software Distribution"

This chapter describes how to use the Software Manager to distribute software packages in the network system.

Part III: Reference

Appendix A, "Using Server Manager in Line Mode"

This appendix describes how to use Server Manager in line mode.

Appendix B, "DBA Command Reference"

This appendix describes the DBA commands available in the SQL Worksheet and in Server Manager Line Mode.

Appendix C, "Compatibility with SQL*DBA"

This appendix describes the differences in functionality and behavior between Server Manager/Line Mode and SQL*DBA.

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.

Before using Enterprise Manager, you should read [Chapter 1, "Overview of Oracle Enterprise Manager"](#). After reading the chapter, 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 tools, you should read [Chapter 7, "Overview of the Database Tools"](#). This chapter provides an overview of the organization and user interface elements of the tools.

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.

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

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.

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 Installation Guide* provides information about installing Oracle Enterprise Manager components. This guide is in an insert in the Enterprise Manager CDROM.
- 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 Performance Monitoring User's Guide*, *Oracle Expert User's Guide*, *Oracle Trace Developer's Guide*, and *Oracle Trace User's Guide* provide information about performance monitoring applications.
- The *Oracle Enterprise Manager Messages Manual* describes the Oracle Enterprise Manager error messages and methods for diagnosing the messages.
- The *Oracle Enterprise Manager Application Developer's Guide* provides information on programming external interfaces to Oracle Enterprise Manager. The guide includes information on using Tcl and OraTcl to write custom job and event scripts.

The guides are available on the Oracle Enterprise Manager CD in HTML format 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 the related publications. Depending on the version of the Oracle database, you would refer to the Oracle7 or Oracle8 documentation. The related books referred to in this guide are listed below:

- For general information about the Oracle Server and how it works, see the *Oracle Server Concepts Guide*.
- For information about administering the Oracle Server, see the *Oracle Server Administrator's Guide*.
- For information about administering Oracle Parallel Servers, see the *Oracle Parallel Server Support for the Oracle Enterprise Manager Console Guide*.
- For information about developing database applications within the Oracle Server, see the *Oracle Server Application Developer's Guide*.
- For the procedures for migrating a previous version of Oracle to Oracle, see the *Oracle Server Migration*.
- For information on Oracle's SQL commands and functions, see the *Oracle Server SQL Reference*.
- For information about Oracle's procedural language extension to SQL, PL/SQL, see the *PL/SQL User's Guide and Reference*.
- For information about Oracle messages and codes, refer to *Oracle Server Messages*.
- For information about the utilities bundled with the Oracle Server, including Export, Import, and SQL*Loader, refer to the *Oracle Server Utilities*.
- For information about distributing and replicating data, refer to the Oracle server distributed database systems and replication documentation.
- For information about the Oracle networking system for Oracle8, see the Net8 documentation, which includes the *Net8 Administrator's Guide*. For information about SQL*Net, see the SQL*Net documentation, which includes the *Oracle Network Manager Administrator's Guide*.
- For information specific to the Oracle Server working on your host operating system, see your operating system-specific Oracle documentation (specific book titles vary by operating system) and system release bulletins, if available.

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.

Your Comments Are Welcome

We value and appreciate your Comments as an Oracle user and reader of the manuals. As we write, revise, and evaluate our documentation, your opinions are the most important input we receive. Included in our manuals is a Reader's Comment Form, which we encourage you to use to tell us what you like and dislike about this manual or other Oracle manuals. If the form is not available, please use the following address.

Documentation Manager
Systems Management Products
Oracle Corporation
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Part I

Oracle Enterprise Manager and Console

- Chapter 1, “Overview of Oracle Enterprise Manager”
- Chapter 2, “Navigator”
- Chapter 3, “Map”
- Chapter 4, “Job Scheduling”
- Chapter 5, “Event Management”
- Chapter 6, “Agents and Communication Daemon”

Overview of Oracle Enterprise Manager

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

- [Introducing Enterprise Manager](#)
- [Enterprise Manager's Components](#)
- [Console Menus](#)
- [Toolbars and Launch Palettes](#)
- [Oracle Administrator Toolbar](#)
- [Getting Started](#)
- [Connecting to an Instance](#)
- [Console User Preferences](#)

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

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.

Introducing Enterprise Manager

Oracle Enterprise Manager combines a graphical console, agents, common services, and tools to provide an integrated, comprehensive systems management platform for managing Oracle products. From Enterprise Manager's Console, you can:

- Administer, diagnose, and tune multiple databases.
- Distribute software to multiple servers and clients.
- Schedule jobs on multiple nodes at varying time intervals.
- Monitor objects and events throughout the network.
- Customize your display using multiple graphic maps and groups of network objects, such as nodes and databases.
- Administer Oracle Parallel Servers. For information about administering Oracle Parallel Servers, see the *Oracle Parallel Server Support for the Oracle Enterprise Manager Console Guide*.
- Integrate participating Oracle and third-party tools.

Attention: See the Oracle database compatibility matrix in the *Oracle Enterprise Manager Readme* for information on the specific Oracle server releases that are supported by Oracle Enterprise Manager and its components.

Enterprise Manager's Components

Enterprise Manager consists of multiple components that integrate into a powerful and easy-to-use graphical user interface.

- Console
- Intelligent Agents
- Administrator Toolbar
- Database Administration (DBA) Tools and Utilities
- Server Manager Line Mode
- Online Help
- Value-Added Products

Console

The Oracle Enterprise Manager Console is a graphical user interface that provides menus, toolbars, launch palettes, and the framework to access Oracle tools and utilities available through other vendors. The format of the Console and the tools available are determined by the products purchased and user preferences. See [Figure 1-1, "Enterprise Manager Console"](#) for an illustration of a Console screen.

Figure 1-1 Enterprise Manager Console



Console menus, toolbars, and tool palettes provide access to the Console components and database administration (DBA) applications. For information on the menus, see *Console Menus* on page 1-10. The components of the Console include:

- Navigator
- Map
- Job Scheduling
- Event Management
- Repository
- Communication Daemon
- Discovery Cache

Navigator

The Navigator discovers and displays a tree list of all the services and objects in a network, providing a direct view of objects such as databases, groups, listeners, and nodes, plus the objects that they contain. While specific types of objects can be viewed in the tree lists of the DBA tools, the Navigator shows all the network objects with their relationships to other objects.

The Navigator tree is populated by the Navigator Service Discovery feature or by reading an Oracle Parallel Server (OPS) topology file (`topo_ops.ora`) for OPS databases that cannot be added with Service Discovery.

Through the Navigator, you can also launch administration tasks with various objects, such as users and tablespaces. You can also perform some administration tasks in the Navigator, such as creating, editing, or dropping a user. For information, see [Chapter 2, "Navigator"](#).

Map

The Map system allows you to monitor network objects with a graphical display. With the Map system, you can create, save, modify, and recall views of the network. You simply drag and drop objects from the Navigator into the Map view to create groups that you want to monitor. For information, see [Chapter 3, "Map"](#).

Job Scheduling

Job Scheduling allows you to manage job scheduling among the databases, listeners, and nodes that you are administering. Jobs can be scheduled at various times, such as daily or weekly, and at single or multiple destinations. The Job window contains the following pages of information:

- Active Jobs
- Job History
- Job Library

For information, see [Chapter 4, "Job Scheduling"](#).

Event Management

Event Management allows you to track and display the status of events occurring on the databases, listeners, and nodes in your network system. You can also create jobs that can be run to automatically fix the problems.

The Event window contains the following pages of information:

- Outstanding Events
- Event Set Library
- Registrations
- Event History

For information, see [Chapter 5, "Event Management"](#).

Repository

The repository is a set of tables in a user account of an Oracle database. Each administrator is associated with a specific repository. Any information related to the tasks performed by the administrator is stored in that repository. For example, the repository contains information about the jobs submitted and the events monitored by an administrator. The repository also stores information about maps and groups that are created by the administrator.

The repository can be stored in any database accessible to the Console and is automatically created or upgraded when you start the Console. An administrator can log on to the repository database from any machine. The repositories for the administrators do not have to be in the same database.

Note: Only one repository should be set up for each username. Creating identical users with repositories in different databases can cause problems with agent notifications. The intelligent agents use the unique username when tracking jobs and events. The username is appended to the Console location when creating a return address for notifications. When you log in to the Console, you must enter the connect information for the database that contains your repository. The Console looks up the repository associated with your username.

Note: You should not log on to the Console repository multiple times with the same username. A warning displays if you attempt this. You should only ignore the warning if the previous Console session was aborted or a machine was disconnected. When a username is logged in multiple times, agent notifications are sent to most recent login.

The information in the repository can be updated by third-party applications that integrate into the Enterprise Manager Console. Refer to the “Repository Control Interface” in the *Oracle Enterprise Manager Application Developer's Guide*.

Communication Daemon

Enterprise Manager uses a Communication Daemon, which runs on the Console machine, to manage Console communication activities, such as communicating with the intelligent agents running on remote nodes in the system.

The Communication Daemon is also one of the components responsible for discovering services in the network environment that populate the Navigator tree.

For information on the Communication Daemon, see [Chapter 6, "Agents and Communication Daemon"](#). For information on Enterprise Manager configuration files, see the *Oracle Enterprise Manager Configuration Guide*.

Discovery Cache

The discovery cache is an efficient in-memory list of information about services and nodes. While you are logged into the Console, the discovery cache stores the current view of the network, including the user-defined groups and the states of nodes and services. When the you log out of the Console, the contents of the discovery cache are stored in the repository.

The next time you start the Console, the discovery cache is loaded with information from the repository and updated with any new information that has been discovered. The new information can include a new service or object discovered by the daemon, or a state that changed since the last login.

The information in the discovery cache can be accessed by third-party applications that integrate into the Enterprise Manager Console. Refer to the “Discovery Cache Interface” in the *Oracle Enterprise Manager Application Developer's Guide*.

Intelligent Agents

The agents are intelligent processes running on remote nodes in the network. Oracle Enterprise Manager uses Intelligent Agents to run jobs and monitor events on remote sites. The Intelligent Agent can also be used to discover services on the node where it resides. The Console Communication Daemon communicates with the Intelligent Agents on the remote nodes in the system.

For information on the Intelligent Agents, see [Chapter 6, "Agents and Communication Daemon"](#).

Administrator Toolbar

The Administrator Toolbar provides a quick and easy method for launching database administration tools without the Console. The toolbar is installed with several database administration tools loaded in the toolbar. You can customize the toolbar with the tools that you want to display and the databases that you want to administer. For additional information on the toolbar, see *Oracle Administrator Toolbar* on page 1-14 and the online help for the Administrator Toolbar.

Database Administration Tools

Enterprise Manager provides DBA tools for administrative tasks such as:

- Instance Manager for managing instances and sessions
- Schema Manager for managing schema objects
- Security Manager for setting up users and controlling security
- Storage Manager for managing storage and rollback segments
- SQL*Worksheet for entering and executing DBA commands, SQL statements, and PL/SQL commands
- Backup Manager for backing up an Oracle database
- Data Manager for exporting, importing, and loading data in and out of an Oracle database
- Software Manager for distributing, installing, and de-installing software packages on servers and clients throughout the network

The DBA tools provide both tree and multi-column views of objects. Some of the functions of these tools can also be performed from the Navigator; however, the full set of functions is accessed through each application. For example, to run tasks on a specific group of objects, such as users, you would launch the Security application which provides complete task-oriented administration of the user objects. For an overview of the DBA tools, see [Chapter 7, "Overview of the Database Tools"](#).

Server Manager Line Mode

For the times when a command line interface is necessary or desirable, Oracle Server Manager provides a conversational line mode. In line mode, you can explicitly execute database administration (DBA) commands on a command line.

You may want to use Server Manager in line mode when a graphical device is unavailable, such as when dialing-in from a non-GUI terminal, or when performing unattended operations, such as running nightly batch jobs or batch scripts that do not require user intervention.

To start Server Manager, enter the operating-specific command name, such as `svrmgr1` (Unix) or `svrmgr30` (Windows NT), at a system prompt. For more information about using Server Manager in line mode, see [Appendix A, "Using Server Manager in Line Mode"](#). For information on the Server Manager DBA commands, see [Appendix B, "DBA Command Reference"](#).

Online Help

Enterprise Manager uses the Microsoft Windows online help system to provide information for windows and dialog boxes in the Console and database tools. The Help system is context sensitive, but you can also search through the online help contents or index to find a particular topic.

There are several ways of accessing the online Help system. In a main window, such as the Console window, you can access the Help system by pressing F1 or choosing Contents from the Help menu. In a dialog box, select the Help button or press F1 to access the online Help system. For information on the Help menu, see *Help Menu* on page 1-12.

Value Added Products

Oracle produces additional administration utilities that can be integrated into the Enterprise Manager Console. These include:

Oracle Performance Pack Oracle Performance Pack provides easy-to-use tools for monitoring and diagnosing databases. These utilities include Oracle Performance Manager, Expert, Trace, Tablespace Manager, TopSessions, and Lock Manager. These utilities gather and inspect specific performance statistics that are useful for tuning your database.

Oracle Replication Manager Oracle Replication Manager enables you to perform many system management tasks for a replicated environment.

Oracle Rdb for NT Oracle Rdb for NT provides graphical tools to administer an Oracle Rdb database.

Oracle Fail Safe Manager Oracle Fail Safe Manager is the graphical, high-availability option for Oracle solutions on Windows NT clusters.

Oracle Security Server Manager Security Server Manager is a security product based on public-key cryptography that supports authentication and authorization in an Oracle network environment.

Oracle Biometrics Manager Oracle Biometrics Manager administers the biometric credentials (fingerprints) of Oracle database users that use the Oracle Advanced Networking Option.

Console Menus

The Console menu bar provides access to the following menus:

File	View	Navigator	Map	Job	Event	Tools	Help
------	------	-----------	-----	-----	-------	-------	------

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

- For information about the Navigator menu, see *Navigator Menu* on page 2-5.
- For information about the Map menu, see *Map Menu* on page 3-6.
- For information about the Job menu, see *Job Menu* on page 4-8.
- For information about the Event menu, see *Event Menu* on page 5-12.

When using the Console menus, note that:

- Some menus include other menus. For example, when you select the Show Map item from the Map menu, a menu containing a list of saved maps displays.
- Menu options vary depending on the objects selected or the active window in the Console. The Tools menu items vary depending on the setup of the system.
- When a menu item is dimmed, it is not currently available. Menu items that end with an ellipsis (...) indicate that you will be asked to enter more information to complete the command.

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 registered event set in the Event Registrations window pane, a menu appears with the Event menu options.

Note: For information on context-sensitive menu options in the Navigator, see *Navigator Context-Sensitive Menus* on page 2-7.

File Menu

The File menu items allow you to manipulate maps, open or close database connections, and change login connections.

Preferences...

Displays the User Preferences property sheet that contains the user credentials for accessing services through the Console. See *Console User Preferences* on page 1-25.

Print Setup

Determines the print setup for the Oracle Enterprise Manager.

Exit

Exits Enterprise Manager.

View Menu

The View menu customizes the Console layout, allowing you to show or hide toolbars, palettes, or windows.

Show Navigator Pane

Hides or shows the Navigator window. The menu item is checked or unchecked.

Show Map Pane

Hides or shows the Map. The menu item is checked or unchecked.

Show Job Pane

Hides or shows the Job window. The menu item is checked or unchecked.

Show Event Pane

Hides or shows the Event window. The menu item is checked or unchecked.

Launch Palettes menu

Hides or shows each or all of the launch palettes set up for Enterprise Manager.

Toolbar

Hides or shows the Enterprise Manager toolbar.

Status Bar

Hides or shows the Status bar at the bottom of the Console.

Advanced Mode

Displays or hides additional property sheet pages that allow you to access an object's advanced features.

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 system configuration. See [Chapter 7, "Overview of the Database Tools"](#) for a discussion of the database administration tools.

Applications

This menu option lists the database administration tools and utilities that you can execute from the Console. See *Database Administration Tools* on page 1-7.

Performance Pack

This menu option lists the optional Performance Pack applications that you can execute from the Console. See *Value Added Products* on page 1-9.

Daemon Manager

The Daemon Manager executes the Communication Daemon monitoring program. For more information, see *Daemon Manager* on page 6-6.

Help Menu

From the Help menu, you can access the Enterprise Manager online help system. For more information about the online help system, see *Online Help* on page 1-8.

Contents

Displays the overview topic of the Oracle Enterprise Manager help system.

Search for Help On...

Displays a dialog box containing a scrolling list of index entries and keywords that you can search for in the help system.

Using Help

Displays information on using the Microsoft Windows help system.

About Oracle Enterprise Manager

Displays a dialog box containing version information about Enterprise Manager and its components.

Toolbars and Launch Palettes

Enterprise Manager provides toolbars and launch palettes to execute commands and run applications.

Toolbars

The Enterprise Manager provides toolbar options that allow you to:

- Perform various menu options for the Event, Job, and Map components.
- Hide or show the Navigator, Job, Event, or Map panes.

The toolbar icons represent the items in the menus. The toolbar options are enabled depending on the objects viewed or selected in a window. Move the mouse cursor over a toolbar icon to display the description of the icon's function in the status bar at the bottom of the Console.

The toolbar can be:

- Moved to the top or bottom of the Console window, or can float as a separate window.
- Displayed or hidden through the View menu.

Launch Palettes

The launch palettes contain the icons of applications that you can launch from the Enterprise Manager Console.

- Each palette represents a portion of the Tools menu options. The palette is the same as the Tools menu except that the palette provides icons rather than text items to represent the applications.
- The launch palette items vary depending on the applications that have been installed on your system.
- The palettes can be moved to the top, bottom, or sides of the Console window. They can also float as separate windows. The palettes can either be displayed or hidden through the View menu.

- You can move the mouse cursor over a palette icon to display the description of the icon's function in the status bar at the bottom of the Console.

Oracle Administrator Toolbar

The Administrator Toolbar provides a quick method for launching database administration tools installed in the toolbar. The Administrator Toolbar can be run independently of the Console.

You can customize the toolbar with the tools that you want to display and the databases that you want to administer. Clicking a tool name or icon launches the application, depending on the display format you choose for the toolbar. The DBA tool connects to the default database if one has been specified in the Databases Page of the Administrator Toolbar Customize property sheet. You can also use the Databases option of the Toolbar Menu to connect to a database entered in the Databases Page.

Toolbar Menu

To display the Administrator Toolbar menu options, press the right-mouse button when the cursor is on the toolbar. The menu options are:

Customize

Displays the Toolbar Customize property sheet pages: General, Applications, and Databases. With the property sheet pages, you can customize the toolbar layout, determine the applications displayed, and specify settings for the databases you want to access. After you have completed the property sheet pages, select OK to save any updates. The toolbar displays with the modifications. Select Cancel to close the property sheet without saving any modifications.

Databases

Displays the list of databases entered in the Toolbar Customize Databases page. You select one of the databases to administer with the DBA tool. This item is only active when you click the right mouse button on a tool set up as a database application in the Applications page.

Database Status

Displays the status of the default database specified in the Toolbar Customize: Databases Page. The status notification is determined by the format of the toolbar. For the floating toolbar, a message box displays to alert the user whether the database status is up, down, or unknown.

For the Application-style toolbar, the color of the toolbar flag is updated if necessary. The flag reflects the status of the default database as follows:

- Green signifies that the database is up.
- Red signifies that the database is down.
- Yellow signifies an unknown status. This includes cases where the toolbar is unable to connect to the default database because of networking problems such as "TNS: could not resolve service name", "TNS: unable to connect to destination", or "invalid username/password". Yellow also displays when default database is not specified.

The status of the default database is also checked whenever the Administrator Toolbar property sheet is updated.

Application Bar/Float Bar

Toggles the toolbar between the application and float styles. The format of these styles is set up in the General page.

Help

Accesses the online help system menu options. You can also press F1 when the cursor is on the toolbar.

Exit

Exits the toolbar.

ToolBar Customize: General Page

This page allows you to determine the format and location of the Administrator Toolbar.

Floating Bar format:

Always on Top

Select this option if you want the toolbar to remain in the foreground.

Hide Title Bar

Select this option if you do not want the toolbar title to be shown.

Button Size

Select the toolbar icon size in the floating toolbar.

Application Bar format:

Top

Displays the application toolbar at the top of the screen.

Bottom

Displays the application toolbar at the bottom of the screen.

Left

Displays the application toolbar at the left of the screen.

Right

Displays the application toolbar at the right of the screen.

Auto Hide

Automatically hides the toolbar when the Customize property sheet or Database Status dialog box displays.

Always on Top

Select this option if you want the toolbar to remain in the foreground.

ToolBar Customize: Applications Page

This page allows you to add, configure, and remove database administration (DBA) tools and additional utilities that appear in the Administrator Toolbar.

To select an application, click on the button just to the left of the listed application. The checkmark enables or selects the option.

Name column

You can edit the application name that appears in the toolbar by editing the name field in the page.

Show? column

You can click in the column to toggle the field. Rather than delete applications from the page, you may choose to not show the icon in the toolbar.

DB App? column

You can click in the column to toggle the field. This field identifies an application as a database administration tool.

Path column

This cannot be modified.

Add button

Starts the Create Application in Toolbar wizard.

1. On the first page of the wizard, select the Browse button to locate the executable name of the application with the Windows File Selection dialog. The applications are usually located in the ORACLE_HOME\BIN directory.
2. On the next page, enter the application name that you want to appear in the toolbar. Choose Yes if the application administers a database; otherwise choose No. If Yes is selected, you can display the list of databases from the Databases page when you click the right mouse button on the application icon in the toolbar and select the Databases menu option. The application attempts to connect to the selected database in the list.

Delete button

Removes the selected application in the page and removes the icon from the toolbar.

Move Up/Move Down buttons

Moves the selected application up or down in the page. This changes the location of the application icon in the toolbar.

ToolBar Customize: Database Page

With this page you can specify databases you want to administer, plus the credentials for accessing the databases. Databases entered here display in a list when you select Databases from the Toolbar menu.

To select a service, click on the button just to the left of the listed service. The checkmark enables or selects the option.

Service column

You can enter and edit the name of the service that you want to administer.

Default column

Click on this column to toggle the default status. When you select Database Status from the Toolbar menu, the status of the default database is checked. If a database service is set as the default, the database tools in the toolbar attempt to connect to this database when started.

User/Password columns

Enter the username and password that you want to use when you connect to this service.

Connect As column

Choose the role that you want to use when connecting to the database. The available roles are NORMAL, SYSDBA, and SYSOPER.

Delete button

Removes the selected service in the page.

Move Up/Move Down buttons

Moves the selected service up or down in the page. This changes the location of the service in the list that displays when you select Databases from the Toolbar menu.

Getting Started

To get started with Oracle Enterprise Manager:

- Enterprise Manager must be installed on a computer running Windows NT or Windows 95. See the Oracle Enterprise Manager Installation CD insert.
- Before you can use the Enterprise Manager Console, you need to connect to the database that contains the repository. The database can reside on any node that is accessible to the Enterprise Manager Console. The user account that contains the repository should have DBA privileges.
- After the Enterprise Manager Console is running, you need to populate the Navigator tree with the nodes and services in the network that you want to administer. See *Discovering Services* on page 2-8.
- To use the Job and Event systems, an Oracle Intelligent Agent must be running on the node that you plan to run jobs or monitor events.

Note: For instructions on configuring Oracle Enterprise Manager and the Oracle Intelligent Agent, see the *Oracle Enterprise Manager Configuration Guide*.

Suggestion: After you have created a repository and discovered services in the network, back up the repository tables in case the data becomes corrupted.

Launching Enterprise Manager

To launch the Enterprise Manager Console, access the Oracle Enterprise Manager group from the Start->Programs menu and select Enterprise Manager. You can also launch the Enterprise Manager Console (`voc.exe`) with the Start->Run option or from the MS-DOS prompt. You can include a connect string with the Run option or at the command prompt. For example:

```
voc.exe user=scott password=tiger service=mydb
```

Only one instance of the Oracle Enterprise Manager Console can be run on a machine. However, you can run multiple instances of the DBA tools.

Connecting to a Repository

When you start Enterprise Manager, the copyright window displays, then the Login Information dialog box appears. Use the login box to connect to the user account in the Oracle database where your Enterprise Manager repository is stored. For information on connecting to an Oracle database, see *Connecting to an Instance* on page 1-23. For information on repositories, see *Repository* on page 1-5.

If the Enterprise Manager repository does not exist, it is automatically created and a dialog box informs you of the operation. If the repository exists, the version must be compatible with the Enterprise Manager Console version. If the existing Enterprise Manager repository version is older than the Oracle Enterprise Manager Console version, the existing repository is upgraded automatically and a dialog box informs you of the upgrade.

If the Oracle Enterprise Manager Console version is older than the Enterprise Manager repository version, you need to install a more recent, compatible version of Enterprise Manager.

You can also create, validate, or drop a repository with the Repository Manager wizard. See *Repository Manager Wizard* on page 1-21.

Warning: Do not log on to the Oracle Enterprise Manager repository multiple times with the same username. Enterprise Manager requires unique usernames across all repositories whether or not the repositories are in different databases. A warning displays if you attempt this and you are using the same repository. Ignore the warning only if the previous Console session was aborted or a machine was disconnected. Because an Intelligent Agent uses the unique username when tracking jobs and events, identical users cause problems with agent notifications for jobs and events.

Privileged Operations

To perform a task using Enterprise Manager, you must have the appropriate privileges. To fully utilize Enterprise Manager, you must have DBA system privileges when you connect to the Console. For example, to create a tablespace using Enterprise Manager, you must have the CREATE TABLESPACE system privilege. Because many of Enterprise Manager's windows include information selected from data dictionary tables, you need the SELECT ANY TABLE system privilege. These system privilege are part of the SYSDBA/SYSOPER role. For information about privileges and roles, refer to *Oracle Server Concepts* and the *Oracle Server Administrator's Guide*.

Repository Manager Wizard

The Repository Manager Wizard is launched by selecting Repository Manager in the Oracle Enterprise Manager program group. The wizard allows you to:

- Create/validate Enterprise Manager repositories
- Drop Enterprise Manager repositories
- Save a list of the discovered nodes

Creating, Validating, or Dropping a Repository

During the create/validate process, Repository Manager upgrades an existing repository or creates a new one if one does not exist. If you drop a repository, the information contained in the repository is lost. For information on repositories, see *Repository* on page 1-5 and the *Oracle Enterprise Manager Configuration Guide*.

You can choose the repository components that you want to create, validate, or drop. Note that the Performance Pack is dependent on the Enterprise Manager repository. Because of this dependency, a selection may be automatically checked and grayed out when another option is selected. For example, if you choose to create a Performance Pack repository, the Enterprise Manager option is automatically selected because it is required.

When you choose Drop, Enterprise Manager is selected by default. Note that Performance Pack is also selected because that repository cannot exist without the Enterprise Manager repository. If you want to drop only the Performance Pack repository, remove the check on the Enterprise Manager repository then check the Performance Pack option.

Saving a List of Discovered Nodes

You can also choose to save a list of the discovered nodes in an existing repository to a text file. Specify the name, such as `discnode.txt`, and location of the file where you want to save the list of discovered node names. The default location is the `ORACLE_HOME\BIN` directory.

This text file can be loaded into the Add Nodes page of the Discover New Services wizard. See *Add Nodes Page* on page 2-9. Before you drop a repository, you can save the discovered node list so you can quickly discover the nodes in a new repository.

Logging into a Repository

Enter the login information for the account where the repository is located or where you want to create a repository. See *Connecting to an Instance* on page 1-23.

Launching Tools

You can launch an Enterprise Manager database administration tool with the Oracle Administrator Toolbar, as a separate application, or through the Console. However, some tools cannot be launched without the Console and some have limited functionality without the Console.

When a tool is launched, it attempts to connect to a database if one has been selected. If the login is unsuccessful or database has not been selected, the Login Information dialog displays. For information on connecting to an Oracle database, See *Connecting to an Instance* on page 1-23.

- To launch a tool with the Administrator Toolbar, click the tool's icon in the toolbar. The tool connects to default database if one is specified. For more information, see the online help for the Administration Toolbar or *Oracle Administrator Toolbar* on page 1-14.
- To launch a tool separately, select the application in the Oracle Enterprise Manager group. Enter the connect information in the Login Information dialog box. You can also use the Start->Run option or the DOS prompt. You can include a connect string with the Run option or at the command prompt. For example, to launch Security Manager (`vac.exe`):

```
vac.exe user=scott password=tiger service=mydb
```

- To launch a tool through the Console, you can:
 - Select the database you want to administer in the Navigator tree or in the Map window, then choose the application in the Console Tools menu or the launch palette.
 - Select the database you want to administer in the Navigator tree, then choose the application from the Related Tools menu of the context-sensitive menu.
 - Choose the application in the Console Tools menu or in the launch palette, then enter connect information in the Login Information dialog box.

Note: When you select a database in the Console before launching a tool, you are connected to the database according to the preferred credentials that have been set up for the database or the credentials you used to log on to the Console. If connection to the database fails for any reason, the Login Information dialog box displays. See *Console User Preferences* on page 1-25 for more information on connecting to a database.

Administration of a Remote Database

Before submitting a database administration task, such as starting up or shutting down a database, through the Oracle Enterprise Manager Job system, you need to set up the database password file on the node where the database and agent are located. Use the Unix password file creation utility `orapwd`. Enter the following at the Unix operating system prompt:

```
orapwd file=$ORACLE_HOME/dbs/orapw$ORACLE_SID password=manager
```

Also, you need to add the following to the `init.ora` database initialization file:

```
REMOTE_LOGIN_PASSWORDFILE=EXCLUSIVE
```

Note: When you log onto Oracle Enterprise Manager, you need to connect as SYSDBA (or SYSOPER) to have the privileges to startup and shutdown databases.

Connecting to an Instance

You can connect to any Oracle database instance using the Login Information connect dialog box. See [Figure 1-2, "Login Information Dialog Box"](#) for an illustration of the Login Information dialog box.

Figure 1-2 Login Information Dialog Box



You can display the Login Information connect dialog box by selecting a database in the Navigator window and choosing Connect from the Navigator menu. You can also select a database in the Navigator window with the right-mouse to access the Connect menu option.

See *Console User Preferences* on page 1-25 for more information on connecting to a database. See [Chapter 9, "Managing Database Security"](#) for information on managing lock status or password expiration on a user account.

The elements of the Connect dialog box are described below:

Username

Your Oracle username for the database to which you are connecting. For information about connecting to a repository, see *Connecting to a Repository* on page 1-19.

Password

Your Oracle password for the database to which you are connecting.

Service Name

Network service name, such as NY_FINANCE, for the database to which you are connecting. The connect strings for the service names must be set up correctly. If not, the agents used by the Job and Event systems are not able to locate the database. For information on configuration files, see the *Oracle Enterprise Manager Configuration Guide*.

Connect As

Pop-up menu containing NORMAL, SYSOPER, and SYSDBA for Oracle7 release 7.1 or later. Only Normal is allowed for release 7.0. SYSOPER and SYSDBA roles the maximum database administration privileges. For more information, see the *Oracle Server Administrator's Guide*.

OK

Initiates connection.

Cancel

Exits dialog box without connecting.

Help

Displays help information.

Multiple Connections

In Enterprise Manager you can have multiple database connections open simultaneously. For example, you can launch several database applications on different databases. The service name you specify when you connect becomes part of the title for each window associated with that connection.

Note: The number of connections that can be simultaneously opened depends on the system and network you are using.

Console User Preferences

This feature allows you to set up specific usernames, passwords, and roles for different network services, such as databases, nodes, or listeners. You determine the preferences for accessing a service with the information entered in the property sheet. These preferences are used when you access a network object in the Navigator or Map, and when running jobs and registering events. For more information, see *Job Credentials* on page 4-3 and *Registering Events* on page 5-3.

The User Preferences property sheet displays a list of databases, listeners, and nodes in the network, along with the service type and the username for accessing the service. The property sheet is accessed with the Preferences option of the Console File menu. Each row in the list of the property sheet includes:

- Service Name
- Service Type (such as database, listener, or node)
- Username

You can click on a column heading to sort on that column. See [Figure 1-3, "User Preferences Property Sheet"](#) for an illustration of the property sheet.

If you have not set up preferences for a database, listener, or node, the login information that you entered to access the Console is used when accessing that service. If that information is not valid for a selected service, you have to enter login information each time you access the object.

Note: Individual instances of an Oracle Parallel Server are not listed. Oracle recommends that all the instances of a Parallel server use the same preferred credentials as the Oracle Parallel Server.

Setting User Preferences

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

Figure 1–3 *User Preferences Property Sheet*



Username

Enter the username. This field is required if a password has been entered.

Password

Enter the password. You can leave this blank.

Confirm

Confirm the password.

Role

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.

Click the OK button at the bottom of the property sheet to save your updates. For more information about dialog boxes and property sheets, see *Dialog Boxes* on page 7-15 and *Property Sheets* on page 7-17. See [Chapter 9, "Managing Database Security"](#) for information on managing lock status or password expiration on a user account.

Navigator

The Navigator discovers and 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 Window](#)
- [Navigator Menu](#)
- [Navigator Context-Sensitive Menus](#)
- [Discovering Services](#)
- [Manipulating Objects in the Navigator](#)

Navigator Window

The Navigator window provides:

- Identification of the objects or services on nodes in the network 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 to create maps.
- A source of objects 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 window by clicking on the Navigator icon in the Console toolbar.

Some objects in the Navigator tree can be multi-selected using the standard Windows Shift and Ctrl key mouse click combinations. This technique can be useful when creating maps and copying tree objects. For more information on using keystrokes in the tree, see *Keystroke Shortcuts* on page 2-19.

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 contains:

- Databases folder
- Groups folder
- Listeners folder
- Nodes folder
- Folders for additional services that have been integrated into the tree

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. See the specific DBA tools chapter that administers an object for more information on an object's relationships.

Figure 2-1 *Navigator Menu and Window*



Populating the Navigator Tree

The Navigator tree is populated with services that the Communication Daemon retrieves with the Service Discovery feature or by reading an Oracle Parallel Server (OPS) topology file (`topo_ops.ora`). The types of services in the network include:

- Databases, including Parallel Servers
- Groups
- Listeners
- Nodes
- Additional services that have been integrated into the tree

Viewing Specific Tree Objects

In large, complex environments you may not want to view the entire network system in one tree. The Navigator has been designed to be highly flexible for viewing objects. With the Split View menu option, you can split the Navigator tree into views that are displayed in separate pages in the Navigator window. Simply click on the tab of a page to move the page and its contents to the front of the Navigator window.

With the Filter Folder menu option, you can specify the objects in a tree folder that want to manage. In a very large system, this option allows you to display only those objects that you need to access. You can use standard SQL wildcard characters when entering the filtering criteria.

Launching Tools

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

You are connected to the database according to the user preferences that have been set up for the database or the credentials you used to log on to the Console. If connection to the database fails for any reason, the Login Information dialog box displays. For information about connecting to an instance of a database, see *Connecting to an Instance* on page 1-23. For information on user preferences, see *Console User Preferences* on page 1-25.

User-defined Groups

User-defined groups are created with the Create Group option of the Map menu. You can use groups to organize related objects in the tree. For information on groups, see *User-Defined Groups* on page 3-3. The group types are:

- Database Groups
- Listener Groups
- Node Groups
- Parallel Server Groups

Creating Maps

You can drag and drop nodes, databases, groups, and listeners from the Navigator tree to the Map window to add those objects to a map. See *Manipulating Map Views* on page 3-2.

Navigator Menu

The Navigator menu allows you to manage objects in the Navigator window. The menu options are enabled according to the object selected in the Navigator tree. Usually the Create, Create Like, Quick 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.

For information on operations for an object type, see the chapter in this guide that discusses the specific application that manipulates the object type. For example, if you want information on the create or alter property sheet of user, profile, or role object types, see [Chapter 9, "Managing Database Security"](#). For an overview of the database application tools, see [Chapter 7, "Overview of the Database Tools"](#).

Create *Object*

Allows you to create a specified object type. The property sheet for the type displays.

Create *Object Like*

Allows you to create an object similar to the selected object in the tree list.

Quick Edit *Object*

Allows you to alter or edit an object of a specified type. The *Object* is determined by the type of the object selected in the tree list.

Remove *Object*

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

Suggestion: It is often much easier to Filter Folders than remove objects in the Navigator tree. Filtering allows you to quickly include the object in the tree if necessary in the future.

If you remove a node from the Navigator tree that was discovered with the Discover New Services or Refresh Services wizard, the node and all its services are removed from the Navigator tree, and the node is removed from the pages of Refresh Services wizard.

If you remove a database, listener, or other service object from the Navigator tree that was discovered with the Discover New Services or Refresh Services wizard, the object appears again in the Navigator tree when the node that contains the service is rediscovered with the Refresh Services wizard.

If you do not want a service discovered by an Intelligent Agent, do not include an entry for the service in the `tnsnames.ora` file on the node where the service is located. On a Unix platform, remove the service from the `oratab` file also.

Connect

Displays the Login Information dialog box to connect to a database instance. For information about connecting to an instance of a database, see *Connecting to an Instance* on page 1-23.

Disconnect

Disconnects from a database instance.

Startup

Starts up a database. See [Chapter 10, "Managing Instances and Sessions"](#).

Shutdown

Shuts down a database. See [Chapter 10, "Managing Instances and Sessions"](#).

Service Discovery

Discovers services on a node in the network. You can also display the status of network services. See *Discovering Services* on page 2-8.

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 window.
3. Click on the tab of the page to move the page and its contents to the front of the Navigator window.

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 window.

Filter Folder

Filters a folder contents based on the criteria you enter in the filter dialog box. You can use SQL operators in the criteria field. See *Filtering Folders* on page 2-18.

Navigator Context-Sensitive Menus

For example, if you click the right mouse button on a user, such as SCOTT, the following menu options would display:

- Create
- Create Like
- Remove
- Split View
- Quick Edit
- Related Tools

Related Tools in this example are the database tools that administer users. In this case, Security Manager is a related tool.

For information on menu options for an object type, see the chapter in this guide that discusses the specific application that manipulates the object type. For example, if you want information on the menu options for user, profile, or role object types, see [Chapter 9, "Managing Database Security"](#). For an overview of the database application tools, see [Chapter 7, "Overview of the Database Tools"](#).

Discovering Services

The Navigator Service Discovery feature provides wizards for identifying network services and populating the Navigator tree. These services, such as databases and listeners, can be administered with Enterprise Manager components. Navigator Service Discovery functions are:

- [Discover New Services Wizard](#) for discovering new nodes and the services on the nodes
- [Refresh Services Wizard](#) for refreshing and retrying discovery on nodes
- [Manually Define Services Wizard](#) for manually entering information about services on nodes
- [Status Option](#) for viewing the status of the last discovery operation
- [Update Network Configuration Option](#) for updating the network configuration file on the Console machine

Use the Discover New Services and Refresh Services wizards for nodes that have the Intelligent Agent release 7.3.3 or higher running. If you add services to these nodes, you must restart the agent on the nodes before discovering the new services with the Discovery wizard option.

Use the Manual Define Services wizard for nodes that do not have an Oracle Intelligent Agent installed or do not have a release 7.3.3 or higher agent installed.

Discover New Services Wizard

The Discover New Services Wizard option of the Navigator Discovery feature allows you to add nodes where you want to discover services. These nodes must have an Oracle Intelligent Agent release 7.3.3 or higher running. After the nodes and services are discovered, they are added to the Navigator tree. Nodes can be discovered in two modes:

- Immediate executes an immediate, one-time discovery of nodes.
- Interval discovers nodes continuously at a specified interval.

The pages in the wizard are:

- Add Nodes
- Retry Nodes
- Discovery Interval

- Summary

Note: This wizard also re-discovers nodes in the Navigator tree. If you do not want you to refresh previously discovered nodes, use the Refresh Services wizard to stop refreshing the nodes.

Add Nodes Page

Add the nodes where you want to discover services. These are nodes in the network with release 7.3.3 or higher Intelligent Agents. You cannot add nodes that have been previously added and are listed in the Refresh page of the Refresh Services wizard or the Retry page of this wizard. Use the Refresh Services wizard to refresh those nodes. Use the Retry page of this wizard or the Refresh Services wizard to retry nodes where discovery has failed.

Add

Enter the node name and choose the Add button.

Load

Select the load button to load a text file that contains a list of the nodes that you want to discover. The Load dialog box allows you to browse for the text file. In the text file, each node name must be on a separate line followed by a carriage return. To create a file that contains a list of currently discovered nodes, see *Repository Manager Wizard* on page 1-21.

After the list of nodes has been loaded, edit the list if necessary to remove any nodes that you do not want to discover.

Remove

If you do not want to discover services on a node that you have entered, select the node and choose the Remove button.

Retry Nodes Page

The nodes where service discovery has failed are listed. See *Status Option* on page 2-13 for possible discovery problems. The wizard continues to retry discovery on these nodes unless a node is removed from the list.

Undiscovered Nodes

Select a node where discovery has failed and choose the Remove button if you do not want to re-discover the node. If you remove the node from the list, the node and services on the node are removed from the Navigator tree if they existed.

Discovery Interval Page

Nodes can be discovered in Immediate or Interval modes:

Immediate

With Immediate discovery, the nodes are discovered as soon as you select Finish. The nodes are not re-discovered until you initiate Immediate discovery again.

Intervals

With interval discovery, the nodes that you enter are automatically re-discovered on the interval that you specify.

Frequency

Enter the frequency interval in minutes for node re-discovery.

Update

The update option updates the network configuration file (tnsnames.ora) with node information. This should be run after a Discovery operation. If the network configuration is not updated, you may not be able to access a service. If you do not choose automatic updating, you can select the Update Network Configuration menu option to update the network configuration file.

Summary Page

Review the summary page to verify your selections.

Finish

After you select Finish, the status of the current discovery process is displayed as the wizard attempts to discover the services on the nodes in the list. If the service discovery is successful, the service information is added to the Navigator tree. If a discovery operation fails, it is noted in the Status dialog box. For information on possible problems, see *Status Option* on page 2-13.

If a node is successfully discovered, it is set to automatic refresh mode in the Refresh Nodes page of the Refresh Services wizard. All nodes that are not successfully discovered are automatically retried unless they are removed from the Retry Nodes page.

Refresh Services Wizard

The Refresh Services Wizard option of the Navigator Discovery feature allows you to re-discover services on nodes. These nodes must have an Oracle Intelligent Agent release 7.3.3 or higher running. After the nodes and services are discovered, they are added to the Navigator tree. Nodes can be discovered in two modes:

- Immediate executes an immediate, one-time discovery of nodes.
- Interval discovers nodes continuously at a specified interval.

The pages in the wizard are:

- Refresh Nodes
- Retry Nodes. This page is identical to the page of the Discover New Services wizard. See *Retry Nodes Page* on page 2-9.
- Discovery Interval. This page is identical to the page of the Discover New Services wizard. See *Discovery Interval Page* on page 2-10.
- Summary

Refresh Nodes Page

The nodes where services have been successfully discovered are listed. Determine the nodes where you want to refresh service information.

Nodes to Refresh

Check the boxes next to the nodes where you want to re-discover service information.

Summary Page

Review the summary page to verify your selections.

Finish

After you select Finish, the status of the current discovery process is displayed as the wizard attempts to discover the services on the nodes in the list. If the service discovery is successful, the service information is added to the Navigator tree. If a discovery operation fails, that is noted in the Status dialog box. For information on possible problems, see *Status Option* on page 2-13.

If a node is successfully discovered, it is set to automatic refresh mode in the Refresh Nodes page of the Refresh Services wizard. All nodes that are not successfully discovered are automatically retried unless removed from the Retry Nodes page.

Manually Define Services Wizard

Use this wizard to manually enter and update services on nodes in the network. The pages in the wizard are:

- Add Nodes and Services
- Summary

Note: The services on these nodes cannot be discovered with the Discover New Services wizard because a 7.3.3 or higher Oracle Intelligent Agent is not running on the nodes.

After you have defined the nodes and services, they are added to the Navigator tree. To access the nodes and services, you must add valid entries for the databases and agents in the `tnsnames.ora` file on the Console machine. To submit jobs or events to these nodes, an entry for the agent on the node must be in the `tnsnames.ora` file in order for the destination to be validated. For information on the `tnsnames.ora` file, see the *Oracle Enterprise Manager Configuration Guide*. Also, see the Oracle networking documentation for the specific platform and release of your Oracle database.

Add Nodes and Services Page

Add nodes or services on the nodes with this page. Nodes and services previously discovered are listed here and can be updated.

Node Name

Enter a new node name to add the node or any services on the node to the Navigator tree. Select Add after you have entered the name.

Agent?

Check the box if the node contains a pre-7.3.3 Oracle Intelligent Agent. This box must be checked and the node must have an Intelligent Agent running to use the Jobs or Events system. See *Destination Validation Dialog* on page 2-14.

Databases

After the node name has been added, add the names of the databases on the node. If you plan to use an Intelligent Agent to manage jobs or events on the database, the database name must match the name in the `SNMP.VISIBLESERVICES` parameter in the `snmp.ora` file on the agent machine.

Add

After a new node has been entered, select the Add button.

Remove

Select a node or service and choose the Remove button to remove it from the list and the Navigator tree.

Summary Page

Review the summary page to verify your selections.

Finish

After you select Finish, the nodes and services are updated in the Navigator tree. The services, such as databases and agents, must have valid entries in the `tnsnames.ora` file on the Console machine.

Status Option

The status of the most recent discovery attempt is displayed in the Status dialog box after a discovery operation and when the Status menu option is selected. Status is not displayed after the Manually Define Services menu option. Any error conditions for nodes are listed.

The following are common problems to check if service discovery has failed:

- The node must have an Oracle Intelligent Agent release 7.3.3 or higher running.
- TCP/IP network protocol must be used.
- The Intelligent Agent must be listening on port 1748. Check the `snmp_rw.ora` file in the `ORACLE_HOME/network/admin` directory (Unix example) to see if that port has been changed with the `DBSNMP.ADDRESS` parameter.

The following are common problems to check if a service has not been discovered:

- 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 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 agent if the database was not available when the agent was last started. When the agent starts, database entries are written to the `services.ora` file in the `ORACLE_HOME/network/agent` directory (Unix example).

After you correct any problems, you can execute the Discover New Services wizard to retry discovery on nodes where discovery has failed. If the problem cannot be fixed, you can remove the node from discovery operations with the Retry Nodes Page. If discovery succeeds but a service on a node is missing, use the Refresh Services Wizard to rediscover the node after you have resolved the problem.

Update Network Configuration Option

This option updates the network configuration file (`tnsnames.ora`) with current service discovery information from the most recent successful discovery. This function should be run after a Discover New Services or Refresh Services discovery operation to ensure access to services in the Navigator tree.

When the Discover New Services and Refresh Services wizards are executed, service discovery information is updated. If you use choose automatic updating, or explicitly choose the Update Network Configuration option, new addresses are added and previously discovered addresses are updated in the `tnsnames.ora` file.

Note: This function does not update services entered with the Manually Define Service wizard. You must manually add entries in the `tnsnames.ora` file for services updated with that wizard.

Destination Validation Dialog

The agents at the specified destinations must be validated before the destinations can be sent jobs and events. Any problems are displayed in the Destination Validation dialog and must be corrected before a job or event can be submitted to the destination.

To use a destination entered with the Navigator Manually Define Services wizard, the Agent box must be checked and there must be valid entry for the agent in the `tnsnames.ora` file on the Console machine. This entry must include the agent name listed in the dialog box and must match the `DBSNMP.ADDRESS` parameter in the `snmp.ora` file on the machine where the agent is located. If there is a `sqlnet.ora` file that specifies a default domain, such as `world`, that domain must be added to the agent name. For example, `AGENT_name.world`.

Manipulating Objects in the Navigator

The Navigator interface provides easy manipulation of the services and objects in the tree list. For manipulating objects with keystrokes, see *Keystroke Shortcuts* on page 2-19.

Expanding and Collapsing Objects

The Navigator tree provides various methods for expanding and collapsing objects.

To expand and display the contents of a container, click the '+' to the left of a container's icon and name. The contents of the container varies depending on the type of object container. To collapse a container, click the '-' to the left of a container's icon and name. You can also double-click on containers to collapse and expand them.

Note: The tree list display is often the result of queries executed on a database over a network. Because of this, there may be a delay when expanding a container.

When you expand a database, you implicitly connect to the database. The Console uses the preferred credentials that have been set up or the credentials used when you logged on to the Console. If the connection to the database fails for any reason, the Login Information dialog box displays.

For information on the user preferences, see *Console User Preferences* on page 1-25. For information about connecting to an instance of a database, see *Connecting to an Instance* on page 1-23.

Administering Objects

You can perform some of the administrative tasks that are in the DBA tool applications with the Navigator. The Navigator provides a subset of the options for administering an object.

To administer an object, select the object in the Navigator tree and choose the administration task from the Navigator menu. The menu options available vary according to the object selected. See *Navigator Menu* on page 2-5. 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.

When you choose to edit a database, the database property sheet displays. If the connection to the database fails for any reason, the Login Information dialog box displays.

For information on the user preferences, see *Console User Preferences* on page 1-25. For information about connecting to an instance of a database, see *Connecting to an Instance* on page 1-23. For information on the database instance property sheet, see [Chapter 10, "Managing Instances and Sessions"](#).

When you select tree objects with the right mouse button, a context-sensitive menu displays with a subset of the Navigator menu options and a Related Tools menu. The Related Tools menu lists the database tools that administer the object. For information on the context-sensitive (right mouse) menu options, see *Navigator Context-Sensitive Menus* on page 2-7. For information on the Navigator menu options, see *Navigator Menu* on page 2-5.

Copying Tree Objects

You can drag and drop some objects in the tree list to make copies of the object in different locations. For example, you can drag and drop a user or role from one database to another to create an identical user in a different database. If you attempt a drag and drop operation that cannot be done, a circle with a slash appears to alert you. You can select multiple objects using the standard Windows selection keys.

To make an identical copy of a user in several databases, select the user you want to copy from one database and drag and drop the user to the additional databases 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 or to the appropriate folder under the user.

Note: If you drag and drop a user to a database where the user's Default or System tablespace does not exist, you are prompted to reassign the tablespace. If you choose to reassign the tablespace, the User property sheet displays. If you do not choose to reassign the tablespace, you cannot add the user.

Example: Using the Navigator

This example shows you how to manipulate various objects in the Navigator tree.

1. Click the '+' to the left of the Network folder to expand the tree. If there is a '-' instead of a '+', the folder is already expanded. You can also use the right arrow to expand the tree.
2. Select the Databases folder from the Navigator tree and choose Split View from the Navigator menu to create a new page in the Navigator window. Repeat the process to create a new page for the Nodes folder.
3. In the Databases page, click the '+' sign to the left of the Databases folder to expand the branch and display the databases.
4. Click the '-' sign to the left of the Databases folder to collapse the branch. You can also use the left arrow to collapse the tree.
5. In the Nodes page, click on the '+' sign to the left of the Nodes folder to expand the branch.
6. Click the '+' sign to the left of any node to expand the node. Each node contains a Databases and Listeners folder. However, those folders may be empty.
7. Click the '+' sign to the left of the Databases and Listeners folders to expand the folders and view the contents. The folders may be empty.
8. Continue to expand the objects in the Databases and Listeners folders until all containers are expanded. When you expand a database, you may need to enter login information. See *Connecting to an Instance* on page 1-23.
9. Select any object in the tree and pull down the Navigator menu. Note the options that are available. These are the administrative tasks that can be performed on this object in the Navigator.
10. Select the same object in the tree and hold down the right mouse button to display a context-sensitive menu. Note the options that are available.
11. Select a Users folder in an expanded database. Expand the Users folder and select a user with the right mouse button. Choose Quick Edit from the context-sensitive menu to display the property sheet. If Advanced Mode is not selected in the View menu, select that option to display all the property sheets. Click the OK button when you are finished viewing the pages of the property sheet.

12. Select the Databases page in the Navigator window and choose Remove Split View from the Navigator menu. Repeat the process to remove the Nodes page from the Navigator window.

Filtering Folders

With the Filter Folder menu item you can filter folders to display specific sets of objects. When you select the Filter Folder menu item, the Filter Folder dialog box displays.

To filter a folder:

1. Select the folder you want to filter. For example, select the Databases folder.
2. Choose Filter Folder from the Navigator menu.
3. In the Filter Text field, enter the criteria that you want to filter on. You can use SQL operators in the criteria field. For example, if you selected the Databases folder and entered `db%` as the criteria, only those databases that begin with "db" are displayed. The filter is case-sensitive.
4. Make sure the Disable Filter button is not checked.
5. Click the OK button. Note that a filtered folder has a '?' on the folder in the tree.
6. If the folder is not already open, expand the folder to view the filtered contents.

To turn off filtering:

1. Select the folder where you want to remove the filter. The folder has a '?' on it.
2. Choose Filter Folder from the Navigator menu to display the Filter Folder dialog box.
3. Check the Disable Filter box to turn off filtering.
4. Click the OK button.

Keystroke Shortcuts

The following keystrokes can be used in the Navigator tree to perform various actions. Functionality is dependent on the object selected.

Table 2-1 Navigator Keystrokes

Keystroke	Function
Insert	Create object
Delete	Remove object
Enter	Quick Edit object or right-mouse menu
Ctrl-Enter	Right-mouse menu
Up arrow	Move up the tree
Down arrow	Move down the tree
Right arrow	Expand the container or move right
Left arrow	Collapse the container or move left
Escape	Cancels a dialog box, property sheet, or action.

The Map system allows you to create customized graphical views of network objects. You can modify and save these maps of your network environment. You can also create user-defined groups of network objects to help you organize the objects in the maps. With the Map customized views, you can manage and monitor a subset or group of objects in your system. You simply create a specific map comprised of databases, groups, listeners, and nodes that you need to monitor, then recall the map whenever you need to view it. The topics discussed in this chapter include:

- [Map Window](#)
- [Map Menu](#)
- [Creating, Modifying, and Removing Maps](#)
- [Creating, Modifying, and Removing Groups](#)

Map Window

The Map window provides an area for constructing a graphical custom view of the network environment. See [Figure 3–1, "Map Window"](#) for an illustration of a Map window.

Figure 3–1 *Map Window*



The Map window can be hidden or shown by selecting Map Window in the Console View menu. You can also hide or show the window by clicking on the world icon in the Console toolbar.

Manipulating Map Views

Maps are created by dragging and dropping databases, groups, nodes, and listeners from the Navigator tree into the Map window. You can create, update, and save any number of maps. In addition, you can add a bitmap as a background to the map for visual identification or to graphically show the location of nodes. For example, you can use a map drawing of a city or state for the background of your map.

Maps can be quickly displayed by selecting the map name from pull-down list in the Console toolbar or from the Show Map item in the Map menu.

You can create map views of subsets of nodes and services to be managed. Map views might be based on:

- Functionality
- Organizations
- Geographic location

For example, the administrator who is responsible for the databases in the San Francisco office can construct a map view containing only those databases. The administrator responsible for the Human Resources department in Chicago can create another map view containing the subset of nodes in Chicago. The background of the maps can be set to a bitmap of either San Francisco or Chicago to help identify the map view.

User-Defined Groups

You can create user defined groups of network objects to further organize your network view.

- A group is a collection of similar objects such as databases, listeners, or nodes that share a common location or function.
- A group can be added to a map view or another user-defined group.
- A group is represented by a group icon and name. The icon represents the type of group, such as databases.
- You can double-click on the group's icon to expand the group, and view and update the group's members and sub-groups.
- Groups appear in the Navigator tree after they are created.

For example, you could create a group of databases called CHI_MAIL that contains the mail databases in Chicago. You could add this group to a map view that contains all the databases in the U.S. or you could add this group to a user-defined group of all databases in Chicago.

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.

Status Monitors

You can easily determine the status of an object or user-defined group on a map. If the object or group has been registered as an event destination, the state of the event condition is graphically displayed. Each object, such as a node or database, that is a registered destination displays a signal flag that graphically depicts the status of the event on the object.

- If the flag signal is green, there are no problems on the objects that are monitored.
- 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 object or group is down or unavailable, the object displays a circle with a slash through it. The icon only displays if the UpDown event registered to the object or group.
- If there is no icon, then the object is not being monitored. Note that groups do not show a status monitor when the event is clear (green).

User-defined groups inherit the worst state of any of the members. If one service in a group is down, the group displays a circle with a slash through it.

Note: For status monitors to display on a map object, an agent must be running on the node where the object is located.

For more information on events, see [Chapter 5, "Event Management"](#).

Expanding Objects

Map objects can be expanded 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 definitions of the objects.

- You can double-click on databases to display property sheets.
- You can double-click on a node or user-defined group to expand the object and display the contents of the node or group.

See [Figure 3-1, "Map Window"](#) for an illustration of an expanded node in the Map window. For more information about dialog boxes and property sheets, see *Dialog Boxes* on page 7-15 and *Property Sheets* on page 7-17.

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 Map, the Navigator tree, and other groups to add them to the group. You can delete objects from the group by selecting them and pressing the Delete key.

Nodes

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

Note: The objects on a node cannot be altered in the expanded node window.

Databases

When you double-click on a database, you connect to the database and the instance property sheet displays. If the connection to the database fails for any reason, the Login Information dialog box displays. For information about connecting to an instance of a database, see *Connecting to an Instance* on page 1-23. For information on the user preferences, see *Console User Preferences* on page 1-25.

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.

For information on the database instance property sheet, see [Chapter 10, "Managing Instances and Sessions"](#).

Launching Tools from a Map

You can launch a database tool using an object in the map. Select a database icon in the Map window, 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. For more information on user credentials, see *Console User Preferences* on page 1-25.

Map Menu

The Map menu allows you to customize and manage network views. See [Figure 3-1, "Map Window"](#) for an illustration of the Map menu.

Create Map

Creates a new map with the name you enter. See *Creating, Modifying, and Removing Groups* on page 3-10.

Modify Map

Changes the bitmap used as the background of a map. See *Creating, Modifying, and Removing Groups* on page 3-10.

Remove Map

Deletes an existing map from the repository. See *Creating, Modifying, and Removing Groups* on page 3-10.

Show Map

Displays a list of saved maps for your selection. Select a map from the list to display it in the Map window. You can also select a map from the Map pull-down list in the console toolbar.

Save Map As

Saves the current map with the name you enter. Enter a name in the Map Name field of the Save Map As dialog box and click OK. You can also select a map name in the Existing Maps list and click OK. Maps are stored in the repository.

Create Group

Adds a new group to the current map or a selected group. See *Creating, Modifying, and Removing Groups* on page 3-10.

Modify Group

Changes the bitmap used as the background of a group. See *Creating, Modifying, and Removing Groups* on page 3-10.

Remove Groups

Removes an existing group from the current map or group, or deletes the group from the entire system. See *Creating, Modifying, and Removing Groups* on page 3-10.

Remove Selection

Deletes the selected item from the map window. You can also delete an object from the current map or group by selecting the object and pressing the Delete key.

Creating, Modifying, and Removing Maps

You can create, modify, and remove maps with the Map menu options. A map is automatically saved when you close the map. You can also select the Save Map As menu option to save a new map

Creating a Map

1. Select the Create Map menu item to display the Create a New Map dialog box.
2. Enter a name in the Map Name field.
3. Check the Use Background Bitmap box if you want to add a bitmap to the background of the map. If you check the box, enter the name of the bitmap file in the Bitmap File field or click the Browse button to locate the file with the standard Windows file selection dialog box. When a bitmap is added to a map, the bitmap is stored in the database repository.

Note: If you check the box and do not enter the name of the bitmap file in the Bitmap File field, no bitmap is used. If you enter a file name and remove the check from the box, the bitmap file will not be used.

4. Click the OK button when you have finished entering a filename or an optional bitmap filename. The new map is listed in the pull-down map list located in the toolbar and in the Show Map option of the Map menu.

Adding Objects in a Map

Drag and drop databases, groups, listeners, nodes, and objects in expanded groups from the Navigator tree to add the object in the map window. Only one instance of an item can exist on a map or group.

Deleting Objects in a Map

Select objects in the map and press the Delete key to remove objects from the map. You can also select an item and choose the Remove Selection menu option.

Modifying a Map

Modifying a map allows you to change the bitmap background of a map.

1. Select the Modify Map menu item to display the Modify Map dialog box.
2. The name of the map is automatically entered in the Map Name field. The name cannot be modified.
3. Determine whether you want a bitmap background for the map.
 - a. Click on the empty Use Background Bitmap box if you want to add a bitmap to the background of the map. After you check the box, enter the name of an existing bitmap file in the Bitmap File field or click the Browse button to locate the file with the standard Windows file selection dialog box.
 - b. If the box already has a check and you want to change the bitmap, enter a new name in the Bitmap File field or click the Browse button to locate the file with the standard Windows file selection dialog box.
 - c. If the box already has a check and you want to remove the bitmap background, click on the box to remove the check the box and the bitmap from the map.
4. Click the OK button when you have finished entering an optional bitmap filename.

Removing a Map

1. Select the Remove Map menu item to display the Remove Map dialog box.
2. Select the map you want to remove from the maps in the Existing Maps list.
3. Click the OK button. The map is deleted from the Enterprise Manager repository.

Note: Groups that are displayed on the removed map are not dropped from the system. The groups remain in the Navigator tree.

Example: Creating a Map

The following example illustrates how to create a new map.

1. Select Create Map from the Map menu to create a new map.
2. Enter a name for the new map, such as MyMap.
3. Check the Use Background Bitmap box to add a bitmap to the background. Enter a bitmap name, such as world.bmp, in the Bitmap File field or click on the Browse button to locate the bitmap filename.
4. Click on the OK button to create the new map with a bitmap background.
5. Expand the Databases branch of the Navigator tree.
6. Drag and drop several databases from the tree into the Map window.
7. Expand the Nodes branch of the Navigator tree.
8. Drag and drop several nodes from the tree into the Map window.
9. Double-click on a node icon in the map to open the node window and view the contents. A node may contain databases and listeners.
10. Close (Alt-F4) the node window when you are finished viewing the contents.
11. Any updates made to a map are saved automatically. If you want to make a revised map but keep the original:
 - a. Select Save Map As from the Map menu.
 - b. Enter a new name for the updated map, such as MyMapToo.
 - c. Make updates to the map.

Maps are listed in the pull-down map list and in the Show Map option of the Map menu.

Creating, Modifying, and Removing Groups

You can create, modify, and remove groups with the Map menu items.

Creating a Group in the Current Map or Group

1. Select the Create Group menu item to display the Create a Group dialog box.
2. Enter a name for the new group in the Group Name field of the dialog box.
3. Select the group type from the pull-down list. The group types are Database, Listeners, Node, and Parallel Server.
4. Check the Use Background Bitmap box if you want to add a bitmap to the background of the group. If you check the box, enter the name of the bitmap file in the Bitmap File field or click the Browse button to locate the file with the standard Windows file selection dialog box. When a bitmap is added to a group, the bitmap is stored in the database repository.

Note: If you check the box and do not enter the name of the bitmap file in the Bitmap File field, no bitmap is used.

5. Click the OK button to save the group.

If an existing group is selected, the new group is added to that group. Otherwise, the new group is added to the current map. The new group is empty and you need to add objects to the group. The new group can be added to other maps and groups. To add an existing group to the current map or another group, drag and drop the group from the Navigator tree or a different group into the map or group window to add the object.

Creating a Group in the Navigator Tree

1. Select an existing group or one of the Group type folders.
2. Display the Navigator menu or display the context-sensitive menu with the right mouse button.
3. Choose the Create menu option.
4. Follow the steps listed for creating a group. Note that the group type has been determined by the group type selected in the Navigator and cannot be changed.

Adding Objects in a Group

To add objects in a group window, drag and drop groups, nodes, listeners, or databases from the Navigator tree, the current map, or a different group into the group window to add the objects to the group.

To add objects to a group in the Navigator tree, drag and drop groups, nodes, listeners, or databases from the Navigator tree, or a different group on the group folder in the tree to add the objects to the group.

Note: Only one instance of an item can exist in the group. The objects must be of the group type.

Note: 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 window:

1. Double-click on a group icon in a map 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 tree:

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. You can also choose the Remove from Group option in the Navigator menu.

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

Modifying a group allows you to modify the bitmap background of a group.

1. Select a group in the map window.
2. Select the Modify Group menu item to display the Modify Group dialog box. The name of the group is automatically entered in the name field. The name cannot be modified.
3. Determine whether you want a bitmap background for the group.
 - a. Click on the empty Use Background Bitmap box if you want to add a bitmap to the background of the group. After you check the box, enter the name of an existing bitmap file in the Bitmap File field or click the Browse button to locate the file with the standard Windows file selection dialog.
 - b. If the Use Background Bitmap box already has a check and you want to change the bitmap, enter a new name in the Bitmap File field or click the Browse button to locate the file with the standard Windows file selection dialog box.
 - c. If the Use Background Bitmap box already has a check and you want to remove the bitmap background, click on the box to remove the check the box and the bitmap from the group.
4. Click the OK button when you have finished.

Removing Groups

1. Select the Remove Groups menu item to display the Remove Groups dialog.
2. Select the radio button for deleting the group from the current map or group, or from the system.

Note: If you remove a group from the current map or group, you can add the group again at a later time. If you remove the group from the system, the group is deleted from database repository and must be created again.

3. Select the group names from the list and click the OK button. You can select multiple groups.

Additional methods for removing groups are:

- Select the group in the map window and press the Delete key, or choose Remove Selection from the Map menu. The group is removed from the group or map but still exists in the system.

- Select the group in the Navigator tree using the right mouse button and choose the Remove menu option. When you delete a group from another group, the group is only removed from the parent group and is not removed from the system. When you delete a group from a group folder, the group is removed from the entire system.

Example: Creating a Group

The following example illustrates how to create a new database group in a map and add databases to the group.

1. Use the Show Map menu item or the pull-down list of maps to select the map where you want to add the new group.
2. Select Create Group from the Map menu to display the Create a Group dialog.
3. Enter a name for new group in the Group Name field.
4. Select Database Group from the Group Type pull-down list. This is the default and should already be selected.
5. Check the Use Background Bitmap box to add a bitmap to the background. Enter a bitmap name, such as europe.bmp, in the Bitmap File field or click on the Browse button to locate the bitmap filename.
6. Click on the OK button to create the new group. An icon for the new group appears in the current map.
7. Double-click on the new group icon in the map to open the group window.
8. Expand the Databases branch of the Navigator tree.
9. Drag and drop several databases into the group window.
10. Double-click on one of the database icons to display the property sheet. You may have to enter login information to access the database. For information about connecting to an instance of a database, see Entering Login Information.
11. Click OK to close the property sheet when you are finished viewing the pages.
12. Close (Alt-F4) the group window. The updates to the group are saved.

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

Job Scheduling

The Job Scheduling services allow 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, 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 Scheduling Process](#)
- [Job Window](#)
- [Job Menu](#)
- [Creating, Modifying, or Viewing a Job](#)
- [Oracle Job Tasks](#)

Job Scheduling Process

The Job scheduling process includes:

1. Creating a job. This involves:
 - a. Determining the type and destination of the job.
 - b. Determining the tasks of the job.
 - c. Determining the parameters for each task.
 - d. Scheduling the times that the job executes.
2. Saving and modifying a job, if desired.
3. Submitting a job to the Job system.
4. 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 Enterprise Manager Application Developer's Guide*.

The tasks are grouped by the destination of the task:

- Databases, which also include Node tasks
- Name Servers, which also include Node tasks
- Nodes
- Listeners, which also include Node tasks

The tasks allow you to:

- Execute operating system commands or shell scripts.
- Execute SQL and DBA scripts.
- Perform database administration tasks such as starting up and shutting down Oracle databases.
- Start up and shut down Listeners.
- Distribute and install software.

For information on Oracle predefined job tasks and their parameters, see the online help for Oracle job tasks and *Oracle Job Tasks* on page 4-17.

You can combine two or more tasks into one job, called a composite job. Composite jobs can contain test conditions based on the success of a 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.

You can create jobs that can be used as *fixit* jobs for event conditions that are monitored with the Event Management System. Fixit jobs cannot be scheduled. See [Chapter 5, "Event Management"](#) 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 *Administration of a Remote Database* on page 1-23.

Job Credentials

Jobs are normally run with the user 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 agent. If the job runs on a node, the job system passes either the user preferences for the node or, if none are specified, your Console username and password. If the job runs on a service, such as a database, the job system also passes the user preferences for the service. For information on user preferences, see *Console User Preferences* on page 1-25.

Note: 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 node credentials are not set up correctly for an NT node, you may get the "Failed to Authenticate User" error message. 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 or the number of times the job will be run.

When you submit a job, the Console's daemon process sends the job information to the appropriate intelligent agents on the destinations you selected. The agents are responsible for executing the job on the specified schedule and returning job status messages to the Console through the daemon.

Note: There is usually a slight delay between the submitting the job and the notification by the 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 to the Console and specify the destinations on which it should run. The destinations can include nodes, databases, listeners, and user-defined groups that have been created with the Map system.

Note: 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 agent is running.

The Job system of Enterprise Manager allows you to efficiently run jobs on multiple remote nodes by transferring job information to the agents servicing the nodes. When a job is executed, it is run by the agent on that node, thus minimizing network traffic between the remote node and the Console and communication daemon. In addition, jobs can be run on multiple nodes simultaneously because there is an intelligent agent residing on each node.

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 agent is down, the Communication Daemon queues any job requests that could not be delivered to the site. Once the site can be contacted, the daemon submits the queued job to the agent. When a job is submitted to an agent, the Communication Daemon attempts to validate the specified destinations. See *Destination Validation Dialog* on page 2-14.

Job Window

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

- Active Jobs
- Job History
- Job Library

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 window can be hidden or shown by selecting Job in the Console View menu. You can also hide or show the window by clicking the Job system icon in the Console toolbar. See [Figure 4-1, "Job Window in the Console"](#) for an illustration of the Job Scheduling window.

Figure 4-1 *Job Window in the Console*



Active Jobs

The Active Jobs 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 Jobs page is the one that is currently scheduled or running. You can use the Show Details 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 Jobs page to view the job details.

- Job type is determined by an icon at the left of each row.
 - A database icon signifies a database.
 - A computer icon signifies a node.
 - An ear icon signifies a listener.
- Job Name is the name of the job.
- Status of job is one of the following:
 - Submitted: The job has been submitted to the agent at the job destination.
 - Scheduled: The job has been successfully delivered to the agent and is scheduled for execution.
 - Started: The job execution has started. After the job executes, the job execution is listed 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.
 - 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.
- When. This is the time the agent returns after the job has been scheduled by the agent. Note that a fixit job cannot be scheduled and displays "On event occurrence" for this field.
- Destination of the job.

Job History

Job History contains a list of previous job executions. These are jobs that have been submitted to an agent and have executed successfully or unsuccessfully. This page also lists deleted jobs. You cannot modify these jobs.

- Job type is determined by an icon at the left of each row. These icons are the same as the Active Jobs page. See *Active Jobs* on page 4-6.
- Job Name is the name of the job.
- Destination 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.
- Execution Start Time is the time when the job started or was deleted. Because the order of agent notifications may vary, it is possible that the Console receives a completed or failed notification before a running notification. If this happens, the start time displays Unknown.
- Execution End Time is the time when the job finished, failed, or was deleted.

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.

You can save the jobs from the Job History page to a text file, then clear the jobs from the Job History window. This prevents the Job History page from being overloaded with obsolete jobs that occurred in previous days.

Job Library

Job Library contains a list of the jobs that you have created and saved. These jobs can be submitted to the job system at a later date. This is useful if you want to submit the same job at different times. You can use the Modify Saved Job menu option to modify a job selected in this page. You can also double-click on a job listed in the Job Library page to modify the job.

- Job Name is the name of the job.
- Job Description is the user-supplied description of the job.

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 window. See [Figure 4-1, "Job Window in the Console"](#) for an illustration of the Job menu.

Note: When you submit or delete a job, there is usually a slight delay while the agent processes the request.

Create Job

Allows you to create a new job. See *Creating, Modifying, or Viewing a Job* on page 4-9.

Show Details

Displays the property sheet for the selected job in the Active Jobs page of the Job window. 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 4-9.

Remove Job

Removes the selected job from the Active Jobs page of the Job window. If you experience problems removing a job, see the Delete option of the *Daemon Manager Menu* on page 6-11.

Save 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.

Print History

Prints the jobs listed in the Job History page.

Clear History

Clears the jobs listed in the Job History page.

Modify Saved Job

Allows you to modify the job selected in the Job Library page of the Job window. The property sheet is the same as the property sheet for creating a new job. See *Creating, Modifying, or Viewing a Job* on page 4-9.

Creating, Modifying, or Viewing a Job

When you create, modify, or view details of a job, similar property sheets display. The contents of the Create Job and Job Details property sheets are nearly identical. See [Figure 4-2, "Job Property Sheet"](#) for an illustration of a Job property sheet. The property sheets contains:

- General Page
- Task Page
- Parameters Page
- Schedule Page
- Progress Page (Show Details only)

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. Click on the Submit button to submit the job to the agents at the selected destinations. The job appears in the Active Jobs window.
 - b. Click on the Save button to save the job. The job appears in the Job Library window. You can modify or submit a saved job at a later time.

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

Modifying a Saved Job

Note: Only saved jobs that are in the Job Library page of the Job window can be modified. Submitted jobs in the Active Jobs window cannot be modified.

1. Select a saved job in the Job Library page of the Job window.
2. Select Modify Saved Job from the Job menu to display the Job Details property sheet.
3. Update the pages of the Job Details property sheet.
 - a. Determine whether the job is ready to submit.
 - b. Click on the Submit button to submit the job to the agents at the selected destinations. The job appears in the Active Jobs window.
 - c. Click on the Save button to save the job. The job appears in the Job Library window. You can modify or submit a saved job at a later time.

Viewing Job Details

1. Select a job in the Active Jobs page of the Job window.
2. Select Show Details from the Job menu to display the Job Details property sheet.
3. View the pages of the Job Details property sheet. This property sheet cannot be modified.

Job General Page

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

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 Set Management Parameters page after it has been successfully submitted to an agent. See *Creating or Modifying an Event Set* on page 5-14.

Destination Type

Select the destination type from the pull-down list: Database, Listener, Node, 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 Job 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 2-8.

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. For information on Oracle predefined job tasks and their parameters, see the online help for Oracle job tasks and *Oracle Job Tasks* on page 4-17.

Selected Tasks

You can remove tasks from this list. Select the task and click on the >> (Remove) button.

Figure 4–2 Job Property Sheet



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. For information on Oracle predefined job tasks and their parameters, see the online help for Oracle job tasks and *Oracle Job Tasks* on page 4-17.

For some jobs, you can override the preferred credentials for connecting to the service. This allows you to enter a username and password. For information on user preferences, see *Job Credentials* on page 4-3 and *Console User Preferences* on page 1-25.

Job Schedule Page

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

Execute

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.

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.

Time Zone

Select the time zone from the pull-down list. The choices are Agent, Console, and GMT. For this release, only the Agent time zone is available.

Job Progress Page

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.

Note: The Progress page is only available with the Show Details option of the Job menu.

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.

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:

Job Type

The job type is depicted with an icon at the left of the row. These icons are the same as the Active Jobs page. See *Active Jobs* on page 4-6.

Status

For information on the status of a job, see *Active Jobs* on page 4-6 and *Job History* on page 4-7.

Destination

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

Notification Time

This is the time the Console 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.

Note: 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:

Save As

Select the Save As button to save the output as a local file using the standard Windows file dialog box.

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.

Job Definition

Select the Job Definition button to display more details.

Example: Creating a Job

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

1. Select Create from the Job menu to create a new job. The Create Job property sheet displays.
2. Enter a name for the new job in the Job Name field of the General Page. You may also enter a description for the job in the Description field.
3. 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. For information on fixit jobs with events, see *Event Set Management Parameters Page* on page 5-16.
4. Select Database from the Destination Type pull-down list.
5. 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 4-3.

6. Repeat the previous step for another destination. You are choosing the destinations where the job will be run.
7. Click on the Task Page tab of the Create Job property sheet.
8. Select Run SQL*Plus 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.
9. Click on the Parameters Page tab of the Create Job property sheet to set the parameters:
 - If there is an existing SQL*Plus file that you want to run, click on the Browse button of the SQL File Name field to locate the existing SQL*Plus script file. After you locate the file in the file selection box, click the OK button. The contents of the SQL*Plus script appears in the Script Text box. Edit the contents of the existing SQL*Plus script if necessary. Click on the Save button to save your updates.
 - If you want to create a new script, enter SQL*Plus commands in the Script Text box before entering a name in the SQL*Plus File Name field. For example, you could enter “SELECT * FROM dba_users;” in the text box. When you are finished, click the Save button and provide a name for the new SQL*Plus script file.
 - Enter any command-line parameters that you want to run with the SQL*Plus script in the SQL Parameters field.
 - Do not check the Override Preferred Credentials box. If you do not override the credentials, the information that is set up with the User Preferred Credentials property sheet is used. See *Job Credentials* on page 4-3 for permissions needed to run job tasks.
10. Click on the Schedule Page tab of the Create Job property sheet to schedule the execution of the job.
11. Select the Interval under Execute to execute the job on a specific interval.
12. Set the Start Execution Date of the job to 7/1/98. Select the month value in the Start Execution Date field and enter 7. You can also click the up and down arrows to change the value when a number is selected.

13. Repeat the process in the previous step for the day and year values of the Start Execution Date.
14. Change the Start Execution Time to 12:00 AM. Use the same procedure that you used to set the date.
15. Check the End Execution box to set a final execution date for this job.
16. Set the End Execution Date of the job to 12/30/98. Set the End Execution Time of the job to 12:00 AM.
17. The Time Zone should be set to Agent. This job will execute at the local time zone where the agent is located. This is the time zone of the destination.
18. Select the Every ... Days button to set the job frequency for day interval. Click on the up arrow to change the value to 3 days.
19. Click the Save button to save the new job in the Job Library window. You may want to modify this job later.
20. Select the job in the Job Library window and choose the Modify Saved Job menu option from the Job menu. You can also double-click on the job.
21. Click on the Submit button to submit the job to the Intelligent agents at the selected destinations. After the destination is validated and the 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 7/1/98 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 Progress Page.

Oracle Job Tasks

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

- Oracle databases
- Operating systems or hosts
- 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.

For information on the import, export, and load tasks, see *Data Manager Tasks* on page 14-8. For information on Software Manager job tasks, see *Distribute and Install Software Packages* on page 15-21. For information on backup tasks, see *Selected Backup Manager Tasks* on page 12-48. Other tasks; such as Parallel Server job tasks; are described in the online help and documentation for the specific product.

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
- Run DBA Script
- Shutdown Database
- Startup Database
- Backup Tablespace, see *Selected Backup Manager Tasks* on page 12-48

Note: You need to set up a password file to perform administration tasks on a remote database. See *Administration of a Remote Database* on page 1-23.

Run SQL*Plus

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 load a previously created script or simply type SQL commands in the Script Text box. You can record and save SQL commands with the Log menu options and use those saved commands as a script. When you create a SQL script, save it in the `ORACLE_HOME\sysman\scripts\sql` directory to make the script easily accessible. SQL*Plus must be installed in the `ORACLE_HOME` of the target database.

Parameters:

1. **SQL File Name.** The file must be accessible from the Console. Click on the Browse button to display the standard Windows File selection box to locate the name of an existing file. When you click the OK button after locating a SQL*Plus file, the contents of the file are displayed in the Script Text box.

Note: If you want to create a new SQL*Plus script file, first click the Save button and enter a file name in the File Name field of the Save As dialog box.

2. **SQL Parameters.** Enter one or more arguments that you want the script to use.

3. **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 *Console User Preferences* on page 1-25 for more information.

4. **Script Text.** You can edit the commands in the Script Text box. Click the Save button to save the updates. If you want to save the changes to a different file, you can change the name of the SQL file in the File Name field of the Save As dialog 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. You can record and save SQL commands with the Log menu options and use those saved commands as a script. When you create a DBA script, save it in the `ORACLE_HOME\sysman\scripts\dba` directory to make the script easily accessible.

Parameters:

1. **DBA Script Name.** The file must be accessible from the Console. Click on the Browse button to display the standard Windows File selection box to locate the name of an existing file. When you click the OK button after locating a DBA script file, the contents of the file are displayed in the Script Text box.

Note: If you want to create a new script file, first click the Save button and enter a file name in the File Name field of the Save As dialog box.

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.
 - c. **Connect As.** Select the role you want to connect as from the pull-down list.

Note: See *Console User Preferences* on page 1-25 for more information.

3. Script Text. You can edit the commands in the DBA Script box. Click the Save button to save the updates.

Note: If you want to save the changes to a different file, you can change the name of the DBA script file in the File Name field of the Save As dialog box.

Shutdown Database

This job task shuts down an Oracle database instance.

Parameters:

1. Mode:
 - Immediate
 - Abort
2. Role:
 - 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 *Console User Preferences* on page 1-25 for more information.

Startup Database

This job task starts up an Oracle database instance.

Parameters:

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

2. **Parameter File.** Enter the initialization parameter filename you want to use for the database. This file is located on the node where the 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 *Console User Preferences* on page 1-25 for more information.

4. **Mount Mode.** Select the mount option from the pull-down list:
 - Exclusive
 - Normal
 - Parallel
5. **Connect As.** Select the connecting role from the pull-down list:
 - SYSDBA
 - SYSOPER
6. **Restrict Connections.** Check this box if you want to 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.

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 an Windows platform, this task sends the message to ALL users on the network. To send a message to specific users/domain, 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 agent and database reside. You may have to include the path for the 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 *Console User Preferences* on page 1-25 for more information.

When you create a Tcl script, save it in the `ORACLE_HOME\sysman\scripts\tcl` directory on the Console machine to make the script easily accessible.

Parameters:

1. TCL File Name. The file must be accessible from the Console. Click on the Browse button to display the standard Windows File selection box to locate the name of an existing file. When you click the OK button after locating a Tcl script file, the contents of the file are displayed in the Script Text box.

Note: If you want to create a new Tcl script file, first click the Save button and enter a file name in the File Name field of the Save As dialog box.

2. 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]}
```

3. Input Files. Enter the filenames for any input files that are used by the Tcl script.
4. Script Text. You can edit the commands in the Script Text box. Click the Save button to save the updates.

Note: If you want to save the changes to a different file, you can change the name of the Tcl script file in the File Name field of the Save As dialog box.

Tcl Script Examples

For information on writing Tcl job tasks, see the Oracle Enterprise Manager Application Developer'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 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 [oraopen $connect_str]
set curl [oraopen $lda]
orasql $curl $sql_statement
set result_row [orafetch $curl]
puts $result_row
oraclose $curl
oracloseoff $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
    orareporevent /user/host/file/alert $oramsgr(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 set that has the "Accept third party events" box checked. See *Event Set Management General Page* on page 5-15.

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. For information on user preferences, see *Console User Preferences* on page 1-25.

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. For information on user preferences, see *Console User Preferences* on page 1-25.

Parameters:

None

Event Management

The Event Management System (EMS) allows you to monitor specific event conditions, such as loss of service or lack of storage, that occur in your network environment. You choose events on databases, listeners, or nodes, then select the threshold parameters for which you want to be notified. You can notify specific system administrators when an event condition occurs. For some events, you can also choose to execute a *fixit* job that automatically corrects the problem.

The following topics are discussed in this chapter:

- [Event Management Process](#)
- [Event Categories and Types](#)
- [Event Window](#)
- [Event Menu](#)
- [Creating or Modifying an Event Set](#)
- [Creating, Modifying, or Viewing an Event Registration](#)
- [Managing Administrators](#)
- [Configuring Mail Services](#)
- [Configuring Paging Services](#)
- [Oracle UpDown Events](#)

Event Management Process

The Event Management System allows you to efficiently monitor a large system. Using EMS and Intelligent Agents, you can effectively monitor any number of databases 24 hours a day, and be alerted when a problem is detected. You can also pinpoint only the services you wish to monitor. EMS can be extended to include other third-party applications that detect events independent of the Intelligent Agents. These applications can be integrated into EMS and communicate directly with the Intelligent Agents.

In the Event Management 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 only those messages related to the events that they have submitted.

The Event Management System includes the following processes:

- Setting up administrators, mail and paging services, and the system modem.
- Creating and registering an event set.
- Interpreting and correcting an event occurrence.

Determining Administrators

You set up the notification services and the administrators that need to be notified for the events. If you plan to notify administrators other than the current Console operator, you need to determine:

- 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.
- The system responsibility and availability of each administrator.

Using Event Sets

You need to create and register event sets, which are simply a group of event conditions that you want to monitor. Oracle Enterprise Manager includes a variety of predefined events that you can use when creating event sets. The events are grouped by services:

- Database
- Listener
- Node

Creating Events

You can use several Oracle predefined event sets that have been installed with Enterprise Manager. See *Event Categories and Types* on page 5-6. You can also create your own event scripts. EMS events are implemented as Tcl scripts with Oracle extensions (OraTcl) to include database specific commands. For information on writing event scripts, see the *Oracle Enterprise Manager Application Developer's Guide*.

Some events have parameters with threshold values that you can customize for your system. The parameter settings, fixit job, and frequency are entered in the Event Information section of the Parameters page of the Event Set Management property sheet. See *Event Set Management Parameters Page* on page 5-16.

Registering Events

Event sets are registered to specific destinations, such as nodes, listeners, or databases. The status of a registered set is viewed in the Status page of the Modify Registration property sheet. This property sheet is displayed with the Modify Registration option of the Event menu. See *Registrations Status Page* on page 5-21.

When an event is submitted to an agent, the Communication Daemon attempts to validate the specified destinations. See *Destination Validation Dialog* on page 2-14.

The event scripts are stored, registered, and executed on a node with the permissions of the Intelligent agent. However, some of the database events, 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 a user that has system privileges. See *Console User Preferences* on page 1-25.

The Intelligent agent is responsible for detecting when a specific event condition has occurred, notifying the Console, and running a fixit job to correct the problem if specified. You need to create fixit jobs with the Job system before you can associate them with an event.

Event scripts can save a state, such as values of global variables, between executions of an event script. This allows the agent to remember if it has already detected a certain event and eliminates redundant event messages to the Console.

The Event Management system uses the Console's daemon process to register event set information with the appropriate Intelligent Agents on nodes in the network. You determine the frequency that an agent checks an event. See *Frequency* on page 5-17 for details on setting the frequency interval for an event.

Event Occurrences

When an event condition occurs, the agent is responsible for alerting the Console and administrators. Each event is logged in the repository and can be viewed and acknowledged in the Outstanding Events window of the Console. See [Figure 5-1, "Event Management Menu and Window"](#) for an illustration of the Event window.

Note: If an event has been registered on a node, then the UpDown Node event is implicitly registered and is triggered when the Intelligent Agent is down.

Event Notifications

The administrators can be notified in various ways, such as electronic mail or paging. Notification occurs as follows:

- A notification is sent when the threshold of the event exceeds the level specified by parameter values. If the event does not have parameters, a notification is sent when the event occurs.
- If the threshold of the event remains above the level specified by parameter values longer than the specified interval, a new notification is not sent.
- If the event condition changes from a warning to an alert or alert to warning, a new notification is sent to the event window.
- If you acknowledge and move an outstanding event to history, a new notification is not sent to the Outstanding Events window.

Interpreting Events

All events return values and some events produce output messages. The events return color values depending on the severity of the event. The colors are displayed on the event severity flag that is located:

- Next to the event name listed in the Outstanding Events window.
- On the object in the map window. For information on the Map window, see [Chapter 3, "Map"](#).

The colors of the event severity flag are:

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

Note: Some events, such as UpDown, 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.

Some events produce output messages that are displayed in the Error Messages for Event list in the Status page of the Modify Registration property sheet. See *Registrations Status Page* on page 5-21.

Correcting Problems

When an event occurs, you need to correct the problem. In some cases, you can create a fixit job specifically for the event. 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 predefined events for the database, listener, and node service types are grouped into categories:

- Fault Management events
- Space Management events
- Resource Management events
- Performance Management events

Only the UpDown events are included with Oracle Enterprise Manager. Additional advanced events are available with the optional Oracle Performance Pack. For information on events and their parameters, see the online help for Oracle predefined events and *Oracle UpDown Events* on page 5-31. All the Node events are supported on Unix and Windows NT platforms. For other platforms, see your platform-specific documentation.

Fault Management Events

This category of events 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 events available in this category include:

- Alert
- UpDown

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

The UpDown events are provided with the Enterprise Manager base product. These events check whether a database, listener, or node is available. With the UpDown event 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 Events

This category of events tracks possible space problems, such as running out of space on a disk or archive device. Examples of space management events 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 Events

This category of events tracks possible resource problems, such as exceeding datafile or lock limits. Examples of resource management events 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 Events

This category of event 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.

Event Window

The Event Management window contains the following pages:

- [Outstanding Events Page](#)
- [Event History Page](#)
- [Registrations Page](#)
- [Event Set Library Page](#)

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. See [Figure 5-1, "Event Management Menu and Window"](#) for an illustration of the Event window.

Figure 5-1 *Event Management Menu and Window*



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

Outstanding Events Page

The Outstanding Events page displays event sets that have been applied and have occurred.

Event

Name of the event and severity flag. Severity of the event can be alert, warning, and clear. These are represented by red, yellow, and green.

Node

Node where the event was monitored.

Type

Type of service that the event monitors, such as Database, Listener, or Node.

Destination

The service that is being monitored.

Date

Date of the event.

Event Message

Message associated with the event, such as “Database went down” or “Tablespace SYSTEM reached 90% capacity.”

Ack. Text

The Comment entered when the event was acknowledged.

Acknowledging Events

To acknowledge an event that has occurred, double-click on the event to display the Acknowledge Event property sheet. You can also select an event and choose the Acknowledge option in the Event menu. When you acknowledge an event, you can enter notes on the nature and progress of the event condition. After you have reviewed an event, you can move it to the History window. See *Acknowledging an Event* on page 5-23.

Event History Page

The Event History page displays a history of events that have occurred and have been acknowledged by an administrator or cleared by an agent. The Event History Page contains the same columns as the Outstanding Events page, plus the following:

Cleared By

Name of user that cleared the event.

Ack. Date

Date that the event was acknowledged.

Viewing Acknowledgments

To view any Comments made by the administrator that acknowledged the event, double-click on the event to display the Acknowledge Event dialog box.

Saving Event History

You can save the events from the Event History page to a text file and clear the events from the Event History window. This prevents the Event History page from being overloaded with obsolete events that occurred in previous days.

Registrations Page

The Registrations page displays the event sets that have been applied and registered to monitor events on any network objects. The Registrations page contains the following information:

Type

Type of event set, such as Database, Node, or Listener is signified by an icon.

Destination

The service that event is monitoring.

Set Name

Name of the event set that is registered.

Event Status

The number of the pending, and successful, and failed events in the event set.

Modify Registrations Property Sheet

Select Modify Registration from the Event menu or double-click on the selected registered event set to display the property sheet for the event set. The property sheet allows you to view and modify registration information. Select the status page tab to display a detailed view of the status of each event. Some events produce output messages that are displayed in the Error Messages for Event list in the Status page of the Modify Registration property sheet. See *Registrations Status Page* on page 5-21.

Event Set Library Page

The Event Set Library page displays the event sets that have been created. The Event Set Library page contains the following information:

Type

Type of event set, such as Database, Node, or Listener is signified by an icon.

Set Name

Name of the event set.

Description

Description of the event set.

of Events

Number of events in the set.

Event Set Property Sheet

Double-click on an event set to display the property sheet for the event set. The property sheet allows you to view and modify the event sets.

Oracle Event Sets

Several predefined event sets have been installed with Enterprise Manager. These appear in the Event Set Library page of the Event window. You can apply these event sets to destinations on your network. The event sets 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.
- Oracle DB Fault: checks whether a database fault event has occurred.
- Oracle DB Resource: checks whether any of the resource event limits have been exceeded for a database.

Note: Only the UpDown event sets are included with Oracle Enterprise Manager. Additional advanced event sets are available with the optional Oracle Performance Pack.

To view the specific events assigned to an event set, double-click on the event set in the Event Set Library page and view the Event page of the Event Set Management property sheet. For information on Oracle events and their parameters, see the online help for Oracle events and *Oracle UpDown Events* on page 5-31.

Event Menu

The Event menu allows you to set up event set 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 window. See [Figure 5-1, "Event Management Menu and Window"](#) for an illustration of the Event menu.

Note: When you register or remove an event, there is usually a slight delay while the agent processes the request.

Acknowledge Event

Acknowledges the selected event in the Outstanding Events page.

Enter any optional Comments on the nature or progress of the event condition in the General page of the Acknowledge Event property sheet, then click on the Move to History button to move the event to the Event History page or click on the Keep button to leave the event in the Outstanding Events page. See *Acknowledging an Event* on page 5-23.

Register Event Set

Registers, or applies, the selected event set to a destination. This action monitors the destination, or network object. See *Creating, Modifying, or Viewing an Event Registration* on page 5-19.

Modify Registration

Modifies the selected registered event set. See *Creating, Modifying, or Viewing an Event Registration* on page 5-19.

Remove Registration

Removes the selected registered event set. If you experience problems removing an event, see the Delete option of the *Daemon Manager Menu* on page 6-11.

Create Event Set

Displays the Event Set Management property sheet and allows you to create the definition of a new event set. See *Creating or Modifying an Event Set* on page 5-14.

Modify Event Set

Modifies the definition of an existing event set. See *Creating or Modifying an Event Set* on page 5-14.

Remove Event Set

Removes the selected event set from the Event window.

Save History

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

Print History

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

Clear History

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

Administrator List

Displays a list of administrators in the Administrators dialog box. Each row in the list contains the Administrator's Name, Paging Service, Email Service, and Email ID.

You can add, delete, and modify administrator information in the Administrators dialog box. When add or modify an administrator, the Administrator property sheet appears. The property sheet contains the General, Systems, and Availability pages. See *Managing Administrators* on page 5-24.

Configure Services

Displays the Services Configuration menu. You can configure Mail, Paging, or Modem information.

To configure mail services, see *Configuring Mail Services* on page 5-27.

To configure paging services, see *Configuring Paging Services* on page 5-28.

To configure modem settings, dialing information, and trace log, see *Oracle UpDown Events* on page 5-31.

Context-Sensitive Menus

If you select an item in the Event window with the right mouse button, the context-sensitive menu for that item appears. This menu is a subset of the Event menu.

Creating or Modifying an Event Set

Event sets include the service type and the event information that you want to monitor. Event sets can consist of multiple events. To create or modify an event set:

1. Choose the Create Event Set option from the Event menu to display the Event Set Management property sheet
2. Complete or modify the fields in the General, Events, and Parameter pages of the property sheet to create a new event set. If you modify an event set that has been registered, those changes are not used by the registered event set.
3. Click the OK button at the bottom of the property sheet to save and exit. The new event set appears in the Event Set Library page of the Event window.
4. You can now register this event set to the destinations you want to monitor.

The property sheet for creating a new event set is the same as the property sheet for modifying an event set, except that the event set name field is read-only. See [Figure 5-2, "Event Set Property Sheet"](#) for an illustration of the Event Set Management property sheet.

You can also use event sets that have been installed with Oracle Enterprise Manager. See *Event Categories and Types* on page 5-6.

Event Set Management General Page

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

Name:

Enter an event set name.

Service Type:

Select the service 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.

Description:

Enter a description or comment for the event set.

Accept third-party events

Check this box to allow the Event system to monitor events that have been initiated outside the Event Management system. You do not need to enter information in the Events or Parameters pages. After completing the Event Set General page, save the event set and register the set at selected destinations. For information on event scripts, see the *Oracle Enterprise Manager Application Developer's Guide*.

Event Set Management Events Page

On the Events page, you determine the events that you want to monitor.

Available Events:

Select the events in the list you want to monitor in this event set, then click on the << (Add) button to move the events to the Selected Events list.

Selected Events:

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

Figure 5–2 Event Set Property Sheet



Event Set Management Parameters Page

The fixit job, frequency, and parameter settings are entered in the Event Information section of the Parameters page of the Event Set Management property sheet. The settings and types of parameters vary according to the event selected. Some events do not have parameters. For information on events and their parameters, see the online help for Oracle events and *Oracle UpDown Events* on page 5-31.

Choose Fixit Job

A fixit job is designed to correct a problem. For example, you may want the 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. See *Job General Page* on page 4-10.

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 set has been registered, you must remove the event registration, select <None> for the fixit job, and then register the event set.

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

Frequency

Determine the frequency that you want the event to monitor 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.

Parameters

The parameters for an event are displayed when the event is selected in the Selected Events 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 event sets 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.

Example: Creating an Event Set

In the following example, a new event set for monitoring extents in a database is created and the General, Events, and Parameters Pages are completed.

1. To create a new event set, select Create Event Set from the Event Menu to display the General Page of the Event Set Management property sheet.

2. Enter an event set name in the Name field, such as MyEventSet.
3. Select Database from the Service Type pull-down list.
4. Enter a brief description of and comments about the event set in the Description field. The event in this example will monitor Maximum Extents. This event is only available with the optional Performance Pack product.
5. Do not check the Accept third party events box.
6. Click on the Events page tab to display the Events Page.
7. Select the Maximum Extents event. This event monitors the number of available extents for table, index, cluster, or rollback segments. Because this event accesses system tables, you must set up User Preferred Credentials with a user that has system privileges. See *Console User Preferences* on page 1-25.
8. Click on the << (Add) button. The event is moved to the Selected Events list and will be added to the event set.
9. Click on the Parameters page tab to display that page. If you had selected several events, you would need to select a specific event in the Selected Events list to display the Event Information and Parameters for the event.
10. 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.
11. Do not change the frequency that the event monitors the destination. The default frequency interval is acceptable.
12. Double-click on Warning threshold in the scrolling list, to change the value of the parameter. Enter 3 in the New Value field of the Enter Parameter Value dialog box, then click the OK button.
13. Double-click on Alert threshold in the scrolling list, to change the value of the parameter. Enter 2 in the New Value field of the Enter Parameter Value dialog box, then click the OK button.
14. 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.
15. Click the OK button at the bottom of the property sheet to save and exit. The new event set appears in the Event Set Library page of the Event window.

After you have created an event set, you can now register the event set to the destinations you want to monitor.

Creating, Modifying, or Viewing an Event Registration

You need to register, or apply, an event set to monitor events on specific destinations in the network environment. The Register Event Set and Modify Registration property sheets are very similar. The property sheets contain:

- General page (read only when modifying)
- Notification page
- Status page (Modify Registration only)

Registering an Event Set

1. Select an event set in the Event Set Library page of the Event window.
2. Choose the Register Event Set option of the Event menu to display the Register Event Set property sheet.

Note: You can first select the Register Event Set option of the Event menu and then select an event set from the set name drop-down list in the General page.

3. Complete the fields in the General and Notification pages of the property sheet to register the event set.
4. Click the OK button at the bottom of the property sheet to save and exit. The event set appears in the Registrations page of the Event window. Each destination database is listed separately with the event set.
5. After the agent for a destination processes the event set, the event set appears in the Registrations page of the Event window. If an event condition occurs, the event set appears in the Outstanding Events page of the Event window.

When threshold values are exceeded for the events, a notification is sent to Console window. This notification changes the color of the severity flag for the event in the Outstanding Events page. If the destination database icon is displayed in the Map window, the flag on the icon changes color. The colors and their meaning are:

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

Note: There is usually a slight delay between the registering the event set and the notification by the agent.

Modify or Viewing An Event Set

1. Select a registered event set in the Registrations page of the Event window.
2. Choose the Modify Registration option of the Event menu to display the Modify Registration property sheet.
3. You can modify fields in the Notification page but not the General page. Display the Status page to view of the status of each event at every destination.

Registrations General Page

On the General page, you identify the event set and destinations that you want the event set to monitor.

Warning: Do not register an UpDown event against the database or node where the repository table is stored. Note that the UpDown event is included in the Oracle DB Fault event set. The communication between the Intelligent Agent and Communication Daemon can cause problems with the repository database connection. If the database or node is down, the event cannot notify the Console because the repository is not available.

Service Type:

Select a service type from the pull-down list. The types include Database, Node, or Listener.

Set Name:

Select an existing event set name from the pull-down list.

Description:

Displays the description of the event set.

Available Destinations:

Select the destinations in the list you want to monitor in this event set, then click on the << (Add) button. The Available destinations vary according to the service type and event set selected. You can register an event only once on any destination; the destination will not display for the event again. 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 2-8.

Selected Destinations:

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

Registrations Notification Page

On the Notification page, you determine who is notified and how they are notified for each destination.

Note: You need to set up administrators and notification methods before you complete this page. See *Managing Administrators* on page 5-24.

Destination

Select a destination from the pull-down list.

Notify Operator on Duty

Check this option to notify the operator that is on duty when the event occurs.

Cause SNMP Trap At Agent

Check this box to trap the notification at the destination where the agent is located.

Operators

Select an operator (administrator) in the list, then select the notification method for this operator. Notification can also be specified when setting up administrators. See *Managing Administrators* on page 5-24.

Notification Method

Choose the method for notifying the operator selected in the Operators list. You can choose paging or mail if those methods have been set up for the operator.

Registrations Status Page

On the Status page, you can view the status and output messages (if any) of each event at every destination where the event has been registered. The Status page contains the following:

Destination

Select an event destination from the pull-down list. The events registered at this destination display in the Status of Registered Event list.

Status of Registered Event at Destination

Select an event in the list to view the error messages, if present.

Error Message for Event

Scroll through the list of messages to view the output of the event. The specific output of each event is described in the online help for Oracle events.

Example: Registering an Event Set

In the following example, an event set for monitoring a database is registered at several databases.

1. Select the event set name from the Event Set Library page in the Event window.
2. Select Register Event Set from the Event menu to display the Registration Property sheet.
3. If necessary, select Database from the Service Type pull-down list in the General page.
4. 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 set.
5. Repeat the previous step for several additional databases.
6. Click on the Notification page tab to display the Notification page.
7. Verify that the Notify Operator on Duty box is checked.
8. Do not check the Cause SNMP Trap At Agent box.
9. If you have set up email services and an administrator with email, select that administrator in the list and check the Email box.
10. Click the OK button at the bottom of the property sheet to save and exit. After the destination is validated and the agent for a destination processes the event set, the set appears in the Registrations page of the Event window. Each destination database is listed separately with the event set.
11. If an event condition occurs, the event set appears in the Outstanding Events page of the Event window. If an administrator has been selected to be notified by email or paging, an email or page is sent from the Console.

When threshold values are exceeded for the events, a notification is sent to Console window. This notification will change the color of the severity flag for the event in the Outstanding Events page. If the destination database icon is displayed in the Map window, the flag on the icon changes color. The colors and their meaning are:

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

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

Acknowledging an Event

Acknowledging events allows you to add Comments to an event and move the event from the Outstanding Events to the Event History page.

To acknowledge an event that has occurred, double-click on the event in the Outstanding Events page to display the Acknowledge Event property sheet. You can also select an event and choose the Acknowledge option in the Event menu. The Acknowledge Event property sheet contains the General and Notification pages. After you have viewed the information on the pages and entered any optional Comments in the Comments field of the General page, you can:

- Click on the Move to History button to move the event to the History page.
- Click on the Keep button to save any Comments and leave the event in the Outstanding Events page.

Acknowledge Event General Page

The General page contains the following fields. Only the Comments field can be updated.

Event

Name of the event.

Date

Date that the event occurred.

System

System where the event occurred.

Node

Destination node.

Event Message

A message describing the event that took place.

Comments

Enter an optional Comment for this event. Comments are useful for explaining how the problem was corrected.

Acknowledge Event Notification Page

The Notification page contains the Notification Status field. This field is for display only and cannot be updated.

Managing Administrators

The Event Management system allows you to set up the administrators that are notified when an event occurs. To manage administrators, choose Administrator List in the Event menu to display the Administrator List dialog box. From the dialog box, you can add, modify, and delete administrators.

You can select specific administrators to be notified when applying an event set. When using a paging service or mail notification, each administrator can be assigned responsibility for specific systems at specific days and times.

Each administrator can be associated with both an email ID and a pager number. The administrator's schedule can be set up for any hours on any days of the week by specific system. For example, an administrator can be scheduled from 7:00 AM through 12:00 PM from Monday to Friday, 10:00AM through 2:00 PM on Saturday, and not available on Sunday.

Note: If you plan to use a mail or paging service for notification, you need to configure the mail, paging, and modem settings for your system before you set up administrators. See *Configuring Mail Services* on page 5-27 and *Configuring Paging Services* on page 5-28. To configure the modem for the system, access the Modems option in the Control Panel.

Adding New Administrators

1. Click on the Add button in the Administrator dialog box to display the Administrator property sheet.
2. Complete the General, Systems, and Availability pages in the Administrator property sheet to set up a new administrator.

Modifying Administrators

1. Select an administrator in the list and click on the Modify button in the Administrator dialog box to display the Administrator property sheet.
2. Modify the General, Systems, and Availability pages in the Administrator property sheet.

Deleting Administrators

1. Select an administrator in the list.
2. Click on the Delete button in the Administrator dialog box to delete the administrator.

Administrator General Page

The General page contains the Administrator's Name, Comments, Paging Information, and EMailing Information. Complete this page to set up the paging service and email address for the administrator.

Name:

Enter the name of the administrator.

Comments:

Enter any Comments.

Paging Information:**Carrier:**

Select the name of the paging carrier service from the pull-down list. The carriers are set up with the Configure Paging Services option in the Event menu.

Pin:

Enter the paging PIN, if required. This entry is required for the GSM and TAP protocols for alphanumeric pagers. For GSM, the PIN is the actual phone number of the administrator.

Test:

Click the Test button to check the accuracy of the Paging information. In the Test Pager dialog box, select the paging service from the pull-down list. For alphanumeric pagers, enter the PIN number and a text message. For numeric pagers, 12345 and a prefix, if specified, is sent. Click on the Send button to send a test page to the specified pager. A message displays to inform you of the status of the test page.

If the test fails, check the paging services configuration. You can also view the paging trace log file in the `ORACLE_HOME\sysman\output` directory if the trace file is enabled in the paging configuration Trace Setup dialog. Contact Oracle Worldwide Customer Support for an explanation of the messages in the log file.

Prefix:

Enter an optional prefix that is added to the numeric pager event Ids that are sent to a numeric pager. This allows you to avoid conflicts with other numeric Ids that are sent to the pager. It also allows you to set up a prefix for specific destinations where events are monitored. You can set up identical administrators except for the administrator's name and prefix assigned to the numeric event Ids. For a list of pager event Ids, see *Numeric Pager Event Ids* on page 5-33.

EMailing Information:

MAPI Mail ID:

Enter the MAPI mail address of the administrator.

SMTP Mail ID:

Enter the SMTP mail address of the administrator, such as `jayr@smp.com`.

Test:

Click the Test button to check the accuracy of the mail information. In the Test Email dialog box, you can enter the subject of the email and the text of the message. Click on the Send button to send a test email to the specified mail Id. A message displays to inform you of the status of the test email.

If the test fails, check the mail configuration. You can also view the mail log file in the `ORACLE_HOME\sysman\output` directory. Contact Oracle Worldwide Customer Support for an explanation of the messages in the log file.

Email Subject

Enter up to 200 characters in the subject field of the email message. The default is OEM Notification.

Administrator Systems Page

The Systems page defines the systems that the administrator is responsible for.

Available Systems:

Select the names of the systems that the administrator is responsible for, then click the << (Add) button.

Systems responsible for:

Select the names of the systems that the administrator is no longer responsible for, then click the >> (Remove) button.

Administrator Availability Page

The Availability page defines the Notification method and the Availability of the administrator for each system that the administrator is responsible for.

Note: The Console must be running for the notification to be sent.

Notification:

Select the notification method, All, Email, or Paging, from the pull-down list. Paging is recommended for urgent events or critical systems.

Systems responsible for:

Select the system in the tree list for which the administrator is responsible.

Availability

Select the hours and days of the week in the table when the administrator is available for the system selected in the list. This day/hour selection determines when the administrator is notified for an event on the selected system.

Configuring Mail Services

The Mail Configuration dialog box allows you to determine the mail service information for your system. You must set up this information correctly to notify administrators by mail. To configure mail services, select the Configure Services Mail option from the Event menu. In the Mail Configuration dialog box, define the information in the following fields:

Current Mail Engine

Set up the mail engine to match the mail engine of your system. Select from the following options in the pull-down list:

MAPI Setup

See the Windows documentation and online help for information on setting up and using the Microsoft Mail Messaging System. You can use this setup with Microsoft Mail or any mail system that is compatible with MAPI.

Send Mail As

Enter the name of a MAPI profile that you have previously set up with the Mail program in the Control Panel. Mail notification from the Console will be sent from this profile.

Password

Enter the password for the mail system.

SMTP Setup

SMTP is the Simple Mail Transport Protocol that uses a TCP/IP-based mail protocol and requires that TCP/IP services are set up on the Console machine. To set up TCP/IP services, access Network settings in the Control Panel of the Windows system.

Send Mail As

Enter an optional name or label that you want to attach to the address of the email notification. For example, `event_system`.

SMTP Server:

Enter the location of the SMTP server, such as `mailserver.company.com`.

Configuring Paging Services

The Event Management paging services notify an administrator with a page when an event has occurred. Alphanumeric pagers provide a brief text message identifying the event and numeric pagers provide the numeric pager event Ids. To configure paging services, select the Configure Services Paging option in the Event menu to display the Paging Services Configuration dialog.

In the Paging Services Configuration dialog box, you can add, modify, or delete paging services. You determine specific Service Settings, pager type and access number, for each service. You also specify the system modem, Telephony Dialing properties, and Trace log setup for all services. After a service is configured, it can be assigned to system administrators with the General Page of the Administrator Paging Property Sheet. See *Administrator General Page* on page 5-25.

Note: Verify that the system modem has been set up. To configure the modem for the system, access the Modems option in the Control Panel. If a modem has not been set up or the selected modem is not available, a warning message displays.

Adding New Paging Services

1. Select the Add button in the Paging Services Configuration dialog box and enter a name for the paging service in the Service Name field of the New Service Name dialog box. Select the OK button.

2. Determine the settings in the Paging Services Configuration dialog box and select the OK button.

Modifying Paging Services

1. Select a paging service name in the Paging Services Names list of the Paging Services Configuration dialog box.
2. Modify the settings in the Paging Services Configuration dialog box and select the OK button.

Deleting Paging Services

1. Select a paging service name in the Paging Services Names list of the Paging Services Configuration dialog box.
2. Select the Delete button. Select Yes to confirm the deletion.

Paging Services Configuration Dialog

Complete the fields to set up the paging service.

Paging Services Names

Add, modify, or delete a paging service name.

- To add a new paging service, select the Add button. Enter a name for the paging service in the Service Name field of the New Service Name dialog box. This name appears in the General page of the Administrator property sheet. Select the OK button. Determine the Service Settings for this paging service.
- To modify a paging service, select a paging service name in the Paging Services Names list and modify the Service Settings for this paging service.
- To delete a paging service, select a paging service name in the Paging Services Names list and select the Delete button. Select Yes to confirm the deletion.

Pager Type

Select Alphanumeric or Numeric depending on the pager type.

- For alphanumeric modems, select either the TAP (Telocator Alphanumeric Paging/input) or GSM (Global System for Mobile communications) protocol to match the protocol of the paging service. Make sure you use a pager phone number that is answered by a software program, not an operator. Contact the paging service for more information.

- For numeric pagers, enter the number of seconds to delay after a connection is made. The default and minimum is 3 seconds. This delay allows the system to wait for the beeps or messages from the paging provider before sending a page. The delay may have to be set on a trial and error basis.

Where to Dial

Specify the paging phone number.

- Select the Country from the list that is provided.
- Enter the Area Code of the paging service.
- Enter the Phone Number. Use commas to force a one second delay. If the paging service requires that you make a selection after a connection is made, enter the necessary characters such as, 5551212,,,,,1#.

Modem

Select the modem for use with the current Console machine. This modem is used for all the paging services. To configure the modem for the system, access the Modems option in the Control Panel. If you are having problems with the paging system, try reducing the baud rate to 2400 or less. If the modem disconnects prematurely from the paging service, check the property setting that determines when the modem stops trying to connect. For the TAP and GSM protocols, the data bits, parity, and stop bits settings are provided and override the modem settings in the Control Panel.

How to Dial

Determine whether you want to use the Telephony Dialing Properties for all paging services on the current Console machine. You can access the Telephony properties with Dialing Properties button if you have elected to use Telephony dialing properties. You can also access Telephony properties with the Telephony option in the Windows Control Panel. You would not want use the Telephony properties if they conflicted with a paging service access number. This could occur when a telephone country code has been changed and does not match the Telephony setting.

Trace Setup

The Paging Trace Configuration dialog allows you to turn Log Paging Trace Information on and off for the current Console machine. By default, the tracing is on and a log is written to a machine-specific location. Each time the Console is started, the existing trace log is replaced. The trace information is very important for debugging problems with the paging system. Oracle recommends that tracing is on.

If you experience problems with the paging system, contact Oracle Worldwide Customer Support to arrange for the trace log to be sent to support.

Oracle UpDown Events

This section lists the Event Management System UpDown events with their parameters and return values. For information on entering parameter values, see *Event Set Management Parameters Page* on page 5-16. A list of events with numeric pager event Ids is also provided. See *Numeric Pager Event Ids* on page 5-33.

UpDown events are specified for database, listener, and node services. The events are also divided into fault, space, resource, and performance management categories.

Note: Only the UpDown events are included with Oracle Enterprise Manager. Additional advanced events are available with the optional Oracle Diagnostics Pack. For more information on advanced events, see the *Oracle Enterprise Manager Performance Monitoring and Planning Guide*.

The event 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 events require special tables in the database. For example, the `catblock.sql` script needs to be run to use the User Blocks event. 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 events require access to system tables and require additional permissions. You need to set up preferred credentials for the monitored database with a user that has system privileges. See *Console User Preferences* on page 1-25.

Database Fault Management Events

This category of events monitors for catastrophic conditions on the system. The administrator should take immediate action when these conditions arise.

UpDown (Database)

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

Parameters

none

User Action

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

Node Fault Management Events

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

Up/Down (Node)

This event checks whether the Data Gatherer being monitored is running. If the Data Gatherer is down, this event is triggered.

Parameters

none

Output

none

Default Frequency

30 seconds

User Action

Restart Data Gatherer.

Listener Fault Management Events

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

UpDown (Listener)

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

Parameters

none

User Action

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

Numeric Pager Event Ids

The Event Management System provides paging services which 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.

The following is a list of Oracle events and the assigned numeric paging Ids.

Database events

- Fault Management events
 - UpDown (Database) 101
 - Alert 107
 - User Blocks 108
 - Probe 109
 - Archiver Hung 129
 - Session Terminated 131
 - Broken DBMS Jobs 132
 - Failed DBMS Jobs 133
 - Unscheduled DBMS Jobs 134
 - Deferred Transactions 135
 - Error Transactions 136
 - Block Corrupted 148
- Space Management events
 - Archive Full 123
 - Chunk Small 124
 - Dump Full 125
 - Maximum Extents 126

■	SnapShot Log Full	137
■	Tablespace Full	142
■	ALERT File Too Large	143
■	Multiple Extents	144
■	Fast Segment Growth	149
■	Resource Management events	
■	Datafile Limit	118
■	Lock Limit	119
■	Process Limit	120
■	Session Limit	121
■	User Limit	122
■	Performance Management events	
■	Buffer Cache Hit Ratio	110
■	Continued Row	111
■	Data Dictionary Cache Miss Ratio	112
■	Disk I/O	113
■	Library Cache Miss Ratio	114
■	Net I/O	115
■	SysStat Table	116
■	SysStat Table Delta	117
■	In Memory Sort	138
■	Free Buffer Wait	139
■	Rollback Segment Contention	140
■	Index Rebuild	141
■	Redo Log Allocation Hit	145
■	Audit Management events	
■	User Audit	147

Node events

- Fault Management events
 - UpDown (Node) 100
 - Data Gatherer Alert 127
 - Data Gatherer Updown 128
- Space Management events
 - Disk Full 105
 - Swap Full 106
- Performance Management events
 - CPU Utilization 103
 - Paging Rate 104

Listener events:

- Fault Management events
 - UpDown (Listener) 102

Agents and Communication Daemon

The Oracle Enterprise Manager Console works with Intelligent Agents and a Communication Daemon to gather information about the network environment, communicate with network objects, and manage jobs and events. Topics discussed in this chapter include:

- [Agents](#)
- [Communication Daemon](#)
- [Daemon Manager](#)

Agents

The agents are intelligent processes running on remote nodes 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:

- Accepting jobs or events from the Console or other third-party applications.
- Running jobs, collecting their results and output, and queueing them for the Communication Daemon.
- Checking for events, and queueing the resulting event reports for the Communication Daemon.
- Returning job and event reports to the appropriate Communication Daemon.
- 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.

Characteristics

Intelligent Agents are autonomous because they function without requiring that the Console or daemon 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 of the Console and are able to execute jobs and monitor events when the administrator has logged out of the Console. 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. These files have a ".q" extension and are stored in the `$ORACLE_HOME/network/agent` directory (Unix platform).

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 Communication Daemon. 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.

When the daemon sends a message to an agent on behalf of an administrator logged into the Console, the daemon sends the agent information about the administrator's language and character set environment. The agent uses the NLS environment information when it performs database administration tasks on behalf of the administrator. This allows administrators to manage databases in their native languages. For example, an administrator in France can administer a database in Germany and receive messages in French.

Network Encryption

Network encryption 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. When the agent supports direct TCP/IP connections, as an option instead of TNS, ANO features are still accessible. 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 Communication Daemon 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*.

Communication Daemon

The Communication Daemon is a process that runs with the Console on the client machine. There is one Communication Daemon for each Console. The Communication Daemon is responsible for:

- 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.
- Contacting nodes periodically to determine if they are up.
- Maintaining a cache of connections to agents on nodes.

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

Communication between the daemon and the agents is vital to the Job and Event systems. The daemon 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 daemon to report results and status messages for the jobs and events.

The daemon and agents communicate using Oracle Remote Operations, which is a remote procedure call mechanism based on the Transparent Network Substrate (TNS). The daemon and agents can also use Oracle Secure Network Services (SNS) to maintain the security of their network transmissions. The daemon can communicate with any agent in the system, regardless of the different protocols used in the distributed environment.

Message Queues

The daemon and agents use message queues for the messages they send. Using queues ensures that no messages are lost even when the Communication Daemon or agent is down. The daemon 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 daemon's operations queue.

Failed Queue

When the daemon 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 daemon places the request in its failed queue. Periodically the daemon tries to contact the agent which is responsible for the operation request in the failed queue. If the daemon successfully contacts the agent, the operation request is removed from the failed queue and sent to the responsible agent.

Notification Queue

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

For example, when the daemon 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 a daemon and an agent to pass messages, they must establish a connection. Rather than requiring that the daemon or agent create a new connection each time it wants to send a message, the daemon 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.

Daemon Manager

The Oracle Daemon Manager allows you to manage communication between the Console's Communication Daemon and agents. The [Daemon Manager Menu](#) provides options for performing administration tasks. The Daemon Manager window provides a tree structure for viewing:

- [Agent Pending Operations](#)
- [Application Pending Notifications](#)
- [Monitored Nodes](#)
- [Daemon Configuration Parameters](#)

Note: If you launch the Daemon Manager when Oracle Enterprise Manager is not running, you can only configure the parameter settings.

Agent Pending Operations

The Agent Pending Operations folder lists the nodes that have pending job or event operations that have not been delivered to the agent. The nodes are listed in the following folders:

- Job Operations
- Event Operations

When viewing this folder, a multi-column list displays on the right side of the screen. The list identifies the node name, the last contact, the last connect attempted, and the number of job and event operations for each node.

Application Pending Notifications

The Application Pending Notifications folder lists the third-party applications and users that have pending job or event notifications that have not been delivered to the agent. The applications and users are listed in the following folders:

- Job Notifications
- Event Notifications

When viewing this folder, a multi-column list displays on the right side of the screen. The list identifies the username, application, and number of job and event notifications for each application.

Monitored Nodes

The Monitored Nodes folder lists the nodes that are being monitored for the UpDown event. When viewing this folder, a multi-column list displays on the right side of the screen. The list identifies the node name, the last contact, and the last connect attempted for each node. If you select a specific node in the tree, additional information identifies the application and user for that node.

Daemon Configuration Parameters

The Communication Daemon parameters and settings are viewed and updated with Oracle Daemon Manager. The defaults are usually sufficient to run Enterprise Manager. See [Figure 6-1, "Parameter Settings for the Communication Daemon"](#) for an illustration of the daemon parameters.

Figure 6-1 *Parameter Settings for the Communication Daemon*



The parameters for the Communication Daemon include:

Listening Address

The listening address of the daemon. The default (Not Found) is actually:

```
(ADDRESS=(HOST=console_hostname)(PROTOCOL=tcp)(PORT=7770)
```

If this address is changed, the setting must be a valid TNS address. If this address is set, the daemon uses this address and the TCP/IP Port setting is ignored.

The Listening Address parameter should be changed when:

- Using a protocol other than TCP/IP, such as the SPX protocol.
- Using the IP address rather than the Console hostname. An IP address is necessary when using a dial-in Console or any other network configuration with dynamically assigned IP addresses.
- The Console hostname is not contained in the Domain Name Server (DNS) of the agent machine.

Note: When entering an address, make sure you are using a valid TNS address. For more information on TNS addresses, see the Oracle networking documentation, such as the Net8 Administrator's Guide or the Network Manager Administrator's Guide.

TCP/IP Port

The TCP/IP port on which the daemon listens on. If the Listening Address parameter is not set, the daemon listens using TCP/IP with this port setting. The default is 7770. If an error message displays regarding "failed to listen for incoming requests", you may need to change the default value.

Number of Cached Connections

The number of open connections to agents. Default is 5 connections.

Networking Polling Timer

The frequency that the daemon checks for incoming data from an agent. Default is 3 seconds. This parameter is not used in this release.

Service Discovery Timer

Frequency that the daemon queries for additional services that have been added to the network. Default is 1800 seconds.

Node Heartbeat Timer

Frequency that the daemon checks to see if a node is up. This applies to nodes registered with the UpDown event. Default is 60 seconds.

Node Heartbeat Interval

Frequency that the daemon checks if a node is up after the node has been determined to be up and working. This can be set to a different value than the Node Heartbeat Timer so that a working node is checked at a less frequent interval. Default is 60 seconds.

Operation Retry Timer

Frequency that the daemon retries any failed operation. Default is 60 seconds.

Number of Worker Threads

The number of threads available. This setting should match the size of the connection cache. Default is 5.

Register User at Startup

Determines whether the user is registered at the computer where Enterprise Manager is started. If set to 1, the user is registered at the machine where Enterprise Manager is started. If set to 0, the registration is ignored. Default is 1.

Use PPP

Instructs the daemon to use PPP to establish communication with the Agent when connecting to the network by modem. If set to 0, the daemon will not use PPP. If set to 1, the daemon will use PPP to connect to the Agent. The parameter must be set to 1 for both hosts with multiple homes and machines that use only PPP. If PPP is not available and this parameter is set to 1, an error message will result when the daemon attempts to connect. The default value is 0.

Note: The PPP support works only with Windows Socket (WinSock) version 2.0 and above. If PPP is enabled, you must use the Windows Dial-Up Networking application when dialing in to the network.

View and Updating Parameters

1. Select Configuration Parameters in the Daemon Manager tree on the left side of the screen to view the parameters and settings. The parameters are displayed on the right side of the screen. Usually, the default settings are sufficient.
2. Double-click on a parameter listed on the right side of the screen to display a dialog box that allows you to update the parameter setting.

3. Enter a new setting value in the dialog box. You can also:
 - a. Click the Default button to enter the default value.
 - b. If a Listening Address has been entered, click the Remove button to remove the address.
4. Click on the Set button to save the value or click the Cancel button to exit the dialog box without changing the parameter setting.

A new parameter setting is used the next time Enterprise Manager is started. You must have permission to update the NT registry to change the parameters.

Daemon Manager Menu

The Daemon Manager menu options allow you to view and manage event, job, and daemon operations.

File

The File menu provides the following option:

Exit

Exits the Daemon Manager.

Edit

The Edit menu provides the following option:

Delete

Deletes the selected job operation from the tree, repository, and the daemon queue. Also notifies the Console of the change.

Note: When problems persist with a job or event operation, the Console and agent may be out of sync. This can happen due to data corruption or the accidental deletion of files. You may also need to delete the operation in the files that the agent maintains on the node where the agent is running. Those files are stored in the `$ORACLE_HOME/network/agent` directory (Unix platform) and have a “.q” extension. The agent must be shut down and all the “.q” files must be deleted. This removes all the jobs and events from that node.

Control Oracle Daemon

The Control Oracle Daemon menu provides the following options:

Force Service Discovery

Forces the daemon to check for the current network objects.

Force Operation Retry

Forces the retry of operations for all nodes. This is useful if the Console machine has been disconnected or if you are dialing into the network at intervals.

Force Node Heartbeating

Forces the daemon to monitor (heartbeat) the nodes to see if they are up.

Node

The Node menu provides the following option:

Ping

Pings the monitored node to check whether the agent on the node is running.

Help

The Help menu provides the following options:

Contents

Displays the overview topic of help.

About...

Displays version information about the program.

Part II

The Database Administration Tools

- Chapter 7, “Overview of the Database Tools”
- Chapter 8, “Managing Database Storage”
- Chapter 9, “Managing Database Security”
- Chapter 10, “Managing Instances and Sessions”
- Chapter 11, “Managing Schema Objects”
- Chapter 12, “Managing Backup and Recovery”
- Chapter 13, “Using the SQL Worksheet”
- Chapter 14, “Managing and Moving Data”
- Chapter 15, “Managing Software Distribution”

Overview of the Database Tools

This chapter provides an introduction to the Oracle Enterprise Manager database administration (DBA) tools and utilities and describes their organization and basic elements. This chapter covers the following topics:

- [Introduction](#)
- [Application Objects](#)
- [Application Menus](#)
- [Multi-Column Lists](#)
- [Dialog Boxes](#)
- [Property Sheets](#)
- [Starting an Application](#)

Introduction

The database tool and utility applications are the primary administrative components of Oracle Enterprise Manager. You can use these applications to perform most of your DBA administration tasks.

The design of the applications allows you to:

- Use the graphical interface to execute the commands quickly and conveniently by pointing and clicking with the mouse rather than manually entering the SQL commands to perform these tasks.
- Concurrently perform multiple tasks in multiple Oracle Enterprise Manager windows.
- Administer multiple databases simultaneously in multiple windows.
- Centralize database administration tasks for both local and remote databases running on any Oracle platform in any location worldwide. In addition, these Oracle platforms can be connected by any network protocols supported by SQL*Net and the MultiProtocol Interchange.
- Perform DBA administrative tasks using the Server Manager line mode interface when a graphical user interface is unavailable or undesirable. The line mode option is available on all platforms.

The applications allow you to perform the following database administration tasks:

- Manage instances, initialization parameters, and sessions
- Manage schema objects
- Manage users, privileges, and profiles
- Manage software distribution
- Manage storage
- Enter and execute SQL statements, PL/SQL code, and Server Manager DBA commands
- Back up, recover, and restore a database
- Export/Import data to/from an Oracle database

Additional applications can be purchased to allow you to:

- Diagnose your database system
- Manage a replication environment

Application Objects

This section provides an overview of the objects administered by the database tools and utilities and briefly describes the objects that each application manages. For detailed descriptions of the objects and commands available in each application, see the specific chapter on the application.

Types of Objects

The database tools and utilities, like the Oracle Enterprise Manager Console, provide a tree list of groups and instances to organize the objects associated with the application. The tree list displays on the left side of an application window after the application has successfully connected to a database. See [Figure 7-1, "An Application Window"](#) for an example of a tree list in an application window.

The groups and instances are usually contained in a database folder. The database folder is the root of the tree list and displays the name of the database that the application is connected to.

Folders hold other objects and display a '+' or '-' to the left of the object's icon and name. A folder can be expanded or collapsed to show or hide its contents. Objects that do not contain other objects do not display a '+' or '-' to the left of the object's icon or name.

- A group is a folder that holds a group of specific object types and is identified with a folder icon. For example, a Users group contains only users.
- A composite or expandable instance is a folder that consists of various objects. For example, the user SCOTT is a composite instance that consists of the system privileges, object privileges, and roles that have been granted to SCOTT. Composite instances have property sheets that can be edited.
- A simple instance is a single object that does not contain other objects. For example, a system privilege, such as ALTER TABLESPACE, is an instance and does not contain any objects.

Figure 7–1 An Application Window



Manipulating Tree List Objects

With the tree list, you can:

- Click on the '+' to the left of an icon to expand and display the contents of a folder. You can also double-click on a collapsed folder icon or name.
- Click on the '-' to the left of an icon to collapse a folder. You can also double-click on an expanded folder icon or name.

Suggestion: For the fastest performance when expanding or collapsing folders in the tree list, click only on the '+' or '-' to the left of the folder rather than double-clicking an object name. If you select an object in the tree list, the application pauses to refresh the contents of the window on the right side before expanding or collapsing the folder.

- Select the Collapse and Expand menu options in the View menu to collapse and expand folders.
- Click on a group name, such as Users or Tablespaces, to display a multi-column scrolling list of summary information on the objects in the folder in the right side of the application window.

An example of a list is shown in [Figure 7-1, "An Application Window"](#) . For information about multi-column scrolling lists, see *Multi-Column Lists* on page 7-14 .

- Click on a composite instance, such as the SYSTEM tablespace or the user SCOTT, to display the associated property sheet for the object. Menu commands are also enabled for the group you selected.

For information about property sheets, see *Property Sheets* on page 7-17.

Attention: If you alter an instance, such as a datafile named DATA1 or a role named CLERK, in any location of the tree list, all occurrences of the object in the tree are changed.

- Drag and drop an object onto a folder to assign the object to the folder. For example, you can drag and drop a role named CLERK onto the user named SCOTT to grant the CLERK role to SCOTT.

Backup Manager

Backup Manager allows you to administer your backup and recovery environment. You can now select from three backup subsystems:

- *Oracle8 Recovery Manager:* Provides an intuitive graphical user interface for Recovery Manager, an extremely powerful and flexible command line backup and recover utility designed expressly for the Oracle8 database. Wizards simplify backup, recovery, and report operations.
- *Operating System Backup:* Provides automated tablespace backup and recovery using the Oracle Enterprise Manager job system (Oracle7 databases). Operating System Backup also provides limited backup and recover capability for Oracle8 databases.
- *Enterprise Backup Utility: Installation and Configuration:* Provides backup and recovery functionality for Oracle7 databases only. See your Enterprise Backup Utility documentation for further information. **Note:** Enterprise Backup Utility is not compatible with Oracle running on Windows NT systems.

Data Manager

Using an intuitive wizard to guide you, Oracle Data Manager supports the Export, Import, and Load utilities.

The Export utility allows you to export data from an Oracle database. The exported data is written to an operating system file in Oracle-binary format and can only be read by the Import utility. Export files can be used to transfer data between databases or used as backups. The Export utility provides dialog boxes that allow you to define how the data is exported.

The Import utility allows you to import data to an Oracle database. Import can only read data that has been created with the Export utility. The Import utility provides dialog boxes that allow you to define how the data is imported.

The Load utility allows you to transfer data to an Oracle database from files, such as text files, that have not been exported from an Oracle database. By using a control file you specify how the data is stored externally and how it is to be loaded into the database.

Instance Manager

With the Instance application you can start up or shut down a database, place the database in ARCHIVELOG mode or NOARCHIVELOG mode, examine and edit database initialization parameters, manage in-doubt transactions, and manage users' sessions. The Instance application manages the following groups:

- Initialization Parameters
- In-doubt Transactions
- Stored Configurations
- Sessions

Note: If you have SYSOPER privileges, you can only startup and shutdown the database from the Console using the context-sensitive menu.

Schema Manager

With the Schema application you can create, edit, and examine schema objects. The Schema application manages the following groups:

- Array Types
- Clusters
- Database Links
- Functions and Procedures
- Indexes
- Object Types
- Packages and Package Bodies
- Refresh Groups
- Queue Tables
- Sequences
- Snapshots and Snapshot Logs
- Synonyms
- Tables
- Table Types
- Triggers
- Views

Security Manager

With the Security application you can create, alter, and drop users, roles, and profiles. The Security application manages the following groups:

- Users
- Roles
- Profiles

Software Manager

Oracle Software Manager provides key services that allow all software to be managed from a central location. This product addresses the problems faced by application software administrators by offering:

- An integrated toolset for managing software in a distributed fashion
- A next generation architecture that offers network scalability through an advanced client-agent-server paradigm
- An easy-to-use graphical user interface that allows administrators to easily run software configuration jobs and track assets across the entire network

Storage Manager

With the Storage Manager you can perform administrative tasks associated with managing database storage. These tasks include managing tablespaces and rollback segments, and adding and renaming datafiles. The Storage Manager manages the following groups:

- Tablespaces
- Datafiles
- Rollback Segments

SQL Worksheet

A SQL Worksheet allows you to dynamically enter SQL statements, PL/SQL code, and Server Manager DBA commands. You can also run scripts from a SQL Worksheet.

A SQL Worksheet maintains a history of the commands you have entered. Thus, you can use the SQL Worksheet to edit and re-execute an earlier command without having to retype it.

You can have multiple SQL Worksheets open at a time, each of which is separate from the others. So, you can commit or roll back work in each worksheet independently.

Application Menus

Database tools and utilities applications include four main pull-down menus in the menu bar:

- File
- View
- Log
- Help

The applications include other menus that are specific to the groups that the application manages. For example, the Security Manager provides the User, Profile, and Role menus in addition to the standard menus.

When a menu item is dimmed, it is not currently available. Menu items that end with an ellipsis (...) indicate that you will be asked to enter more information to complete the command.

File Menu

The File menu items allow you to open or close application windows and create new connections. The File menu includes:

Change Database Connection

Displays the Connect dialog box to connect to an instance and open a window for that database. See *Connecting to an Instance* on page 1-23. To connect to multiple databases concurrently, you need to launch the application multiple times from the console.

Enable Roles

Allows you to select a role from the multi-column list of available roles.

Exit

Exits the application.

View Menu

The View menu allows you to show or hide objects, such as the toolbar or status line, and expand, collapse, or refresh the tree list.

Refresh

Refreshes the tree list. This command is useful when other database administrators are making changes to the objects you are viewing, or when you are making changes in a different window, such as a SQL Worksheet.

Filter Folder

Filters the tree list based on criteria you choose. This allows you to display only the objects that match the criteria.

Expand One Level

Opens and displays the next level of items contained in the selected objects in the tree.

Expanding a folder can be delayed because the objects displayed in the listing are often queries on remote databases.

Collapse Branch

Closes the selected folder in the tree list.

Collapse All

Closes all open folders in the tree list.

Toolbar

Shows or hides the toolbar.

Status bar

Shows or hides the status bar.

Basic/Advanced Mode

Allows you to choose between basic property sheet functionality (Basic Mode) or complete property sheet functionality (Advanced Mode). By default, the last chosen mode is selected.

Log Menu

The Log menu options allow you to record and save the SQL statements that you have executed. When you save the statements to a file, you can also choose to schedule a job using the file on multiple databases. This feature provides the ability to perform identical actions on multiple databases without having to enter each SQL statement over again. If the databases are not identical, you can simply edit the SQL statements in the file before running the file on a database.

Record

Start recording the SQL statements that are executed through the application.

Stop

Stop recording SQL statements. A dialog displays for you to save the recorded statements in a file. Enter a filename. Schedule a job that would run the file of SQL statements.

Pause

Pause the recording of statements while you execute actions that you do not want recorded.

Help Menu

From the Help menu, you can access the Oracle Enterprise Manager online Help system. For more information about Oracle Enterprise Manager online Help system, see *Online Help* on page 1-8.

Contents

Displays the overview page for Oracle Enterprise Manager help.

Search for help on

Displays a search dialog box. You can either search by index or key word.

Using Help

Displays a dialog box containing information about the Help system. See *Online Help* on page 1-8 for information about the Help dialog box.

About Application

Displays a dialog box containing version information about the application.

Additional Menus

Most database tool and utility applications include pull-down menus that contain the commands available for specific groups. The menu titles usually match the names of the groups. For example, the Security application provides the User, Profile, and Role menus.

Some of the menu commands, such as Drop, require that you select an individual object first. These commands are unavailable until you select an appropriate object. Unavailable commands are represented by dimmed menu items. Other menu items, such as Create, do not require an object selection and are never dimmed.

A menu item can be dimmed even when an object is selected because the command does not apply to that particular object. For example, you cannot drop an object type group, such as the Users group.

Some menu items are followed by an ellipsis (...). This indicates that to complete the command, the application prompts you for more information with a dialog box or property sheet.

Context-Sensitive Menus

Context-sensitive menus may also be provided when you select a tree list object or item in the multi-column list with the right mouse button. This feature provides quick access to a subset of the menu options provided in the menu bar. Normally, objects in an application tree list (or the console Navigator) and objects in a multi-column list have similar context-sensitive menu options.

For detailed information about the application menus, see the chapters on the specific applications.

Multi-Column Lists

A multi-column list displays when a group folder, such as a Users group, is selected. Each row in the list provides a quick view of summary information about an object in the selected folder. Additionally, you can select and alter items through their respective property sheets,

Note: If an object has a property sheet, display the property sheet if you want more detailed information. See *Property Sheets* on page 7-17 for information.

The location of a group folder determines the columns that are displayed in the list. For example, the Roles group that is a main branch of the database folder displays different columns than a Roles group under a user.

By default, the objects are sorted on the first column. You can sort a list on any column by clicking on the column heading. The column heading appears in bold to indicate that the list is sorted on that column.

If the list is longer than the its window display area, you can use the vertical scroll bar at the side of list to scroll to the first or last row. You can also resize the application window to make the window display area taller.

If the list is wider than its window display area, you can use the horizontal scroll bar at the bottom of the list to scroll to the first or last column. You can also resize the application window or drag the splitter bar in the window to widen the right side of the window that contains the multi-column list.

An example of a multi-column list is shown in [Figure 7-2, "Multi-column List"](#) .

Figure 7-2 *Multi-column List*




Dialog Boxes

Dialog boxes allow you to make selections, enter information, or confirm actions.

You exit a dialog box by clicking a button. In many cases, you first enter field values, or choose options from pop-up menus or buttons, before you exit. An example of a dialog box is shown in [Figure 7-3, "Dialog Box"](#).

Figure 7-3 Dialog Box



The boxes generally contain one or more of the following command buttons:

OK

Executes the action button that was specified and closes the dialog box. The name of this button may vary.

Cancel

Cancels any changes made and closes the dialog box.

Apply

Implements any changes made to the dialog box.

SQL

Opens or closes the SQL window to show or hide the SQL statements that are used to execute commands.

Help

Displays online Help about the dialog box.


Property Sheets

A property sheet is a dialog box used to specify options when you are creating or altering a composite instance, such as a user or tablespace. A property sheet consists of one or more pages that are labeled with page tabs. By default, one page is open when the sheet is entered, but all pages show page tabs. To access another page, click its page tab.

Property sheets can be displayed on one of two modes: Basic and Advanced. When in Basic mode, a property sheet displays only those pages that are required for basic functionality. Advanced mode displays additional pages that go beyond basic functionality for a particular function. When you close an application, the current mode is saved.

When a composite instance is selected in the tree list on the left side of an application window, the Alter property sheet that is associated with the object is usually displayed on the right side of the window. The properties available for each sheet are specific to the group type. An example of an Alter property sheet for a user is shown in [Figure 7-4, "Property Sheet"](#).

Figure 7-4 *Property Sheet*



Spreadsheets

Some property sheets, particularly those requiring large amounts of data, use spreadsheets for simplified data entry. Spreadsheets consist of a table in which data is typed directly into rows and categorized by columns. Some entry fields in the spreadsheet have predefined content and are toggled by clicking on the table cell with the left mouse button (e.g. setting the Admin Option on or off by clicking directly on the cell).

For most spreadsheets, context-sensitive menus are available. You first select a single row or spreadsheet cell with the left mouse button and then click the right mouse button to call up a context-sensitive menu that is appropriate for a particular application. In some cases, a spreadsheet context-sensitive menu, unlike those available from the tree list or multi-column list, may not be a subset of the menu options available in the application's menu bar. In such cases, the context-sensitive menu may consist of operations that are specific to a particular spreadsheet or spreadsheet cell.

Newly added rows are indicated by a plus '+' sign in the far left column of the spreadsheet until you click Apply.

Property Sheet Buttons

The property sheets generally contain the following command buttons:

Apply/Create

Implements an action specified in the property sheet. Apply would save any changes made to the Quick Edit property sheet and Create would create a new the object specified by the Create property sheet.

Revert

Cancels any changes made to the property sheet and resets the fields to previously saved values.

SQL

Opens or closes the SQL window to show or hide the SQL statements that are used to execute commands.

Help

Displays online Help about the property sheet.

Starting an Application

To start an application, click its icon in the Launch Palette or select its name from the Launch menu in the Oracle Enterprise Manager Console.

- If you have already connected to a database in the Oracle Enterprise Manager Console, the application window displays.
- If a valid database connection has not already been made, the Connect dialog box displays. See *Connecting to an Instance* on page 1-23.

Note: You can change the database connection with the Change Database Connection option in the application File menu.

After a successful connection to a database, the tree list displays a tree list of objects managed by the application. The display on the right side of the window is determined by the object selected in the tree list. A multi-column list, property sheet, or other information displays.

For example, after the Security application completes a connection to a database, the Users, Profiles, and Roles groups display under the database folder in a tree list on the left side of the Security window. The Users, Profiles, and Roles groups can be expanded to show the objects they contain. An example of an application window with an expanded tree list is shown in [Figure 7-1, "An Application Window"](#).

Managing Database Storage

This chapter describes how to use Storage Manager to administer tablespace, rollback segment, and datafile storage in a database. This chapter assumes that you have read [Chapter 7, "Overview of the Database Tools"](#) and are familiar with the interface elements of the database tools. The topics included in this chapter are:

- [Starting Storage Manager](#)
- [Storage Menus](#)
- [Tablespaces Folder](#)
- [Datafiles Folder](#)
- [Rollback Segments Folder](#)

Starting Storage Manager

To start Storage Manager, select a database object from the tree list and click on the Storage icon in the Launch Palette or choose Storage Manager from the Console Tools menu.

- If you have already connected to a database, Storage Manager appears.
- If a valid database connection has not been made, the Login information dialog box displays.

Note: You can change the database connection with the Change Database Connection option in the File menu of the console.

After Storage Manager connects to a database, the Tablespaces, Rollback Segments, and Datafiles folders display in a tree list on the left side of the Storage window. These folders are contained in the database folder which displays the name of the database that the application is connected to.

Figure 8–1 Storage Manager



The display on the right side of the window is determined by the objects selected on the left side of the screen. The right side may contain a multi-column scrolling list or a property sheet.

For more information, see:

- *Application Objects* on page 7-3
- *Multi-Column Lists* on page 7-14
- *Dialog Boxes* on page 7-15
- *Property Sheets* on page 7-17

Storage Menus

The Storage application has four standard menus, File, View, Log, and Help, plus Tablespace, Datafile, and Rollback. The options for the specific Storage menus are described in this chapter. For information on the standard menus, see *Application Menus* on page 7-10.

Tablespace Menu

The Tablespace menu allows you to perform the following operations:

Create

Creates a new tablespace.

Remove

Removes an existing tablespace.

Show Dependencies

Displays database object(s) that rely on a selected tablespace and object(s) that selected tablespace relies on.

Add Datafile

Adds a datafile to a tablespace.

Add Rollback Segment

Adds a rollback segment to a tablespace.

Place Online

Places a tablespace online.

Take Offline (Normal, Temporary, Immediate)

Make Read-Only

Makes a tablespace read-only.

Make Writeable

Makes a tablespace read/write.

These menu options are enabled depending on the object selected in the tree list.

Datafile Menu

When you select the Datafiles container or a datafile in the container, various menu options in the Datafile menu are enabled. Depending on what objects are selected in the tree list, the Datafile menu allows you to perform the following operations:

Create

Creates a new datafile.

Create Like

Creates a new datafile based on parameter settings of an existing datafile.

Show Dependencies

Displays database objects that rely on a selected rollback segment and any objects that the selected rollback segment relies on.

Place Online

Places a datafile online.

Take Offline

Takes a datafile Offline.

Rollback Menu

The Rollback Segment menu allows you to perform the following operations:

Create

Creates a new rollback segment.

Create Like

Creates a new rollback segment based on an existing rollback segment.

Show Dependencies

Displays database object(s) that rely on a selected rollback segment and any object(s) that selected rollback segment relies on.

Remove

Removes an existing rollback segment.

Shrink

Shrinks an existing rollback segment (Oracle 7.2 or later)

Place Online

Places a rollback segment online.

Take Offline

Takes a rollback segment offline.

These menu options are enabled depending on the object selected in the tree list.

Tablespaces Folder

The Tablespaces object type folder contains all the tablespaces in the database arranged alphabetically in the tree list. An individual tablespace can be expanded to show the datafiles and rollback segments in the tablespace.

When you select:

- The Tablespaces object type folder, a multi-column list displays a row of summary information for each of the tablespaces in the database.
- An individual tablespace, a property sheet displays the detailed settings for the tablespace.

For information about managing tablespaces, see the *Oracle Server Concepts*, the *Oracle Server Administrator's Guide*, and the *Oracle Server SQL Reference*.

Tablespaces Multi-Column List

The Tablespace multi-column list displays when the Tablespace folder is selected in tree list. The columns of the list are:

Name

Name of the tablespace.

Status

Status of the tablespace: Online, Offline, or Read-only.

Size (M)

Total size in megabytes of the datafiles that comprise the tablespace.

Used (M)

Amount of space in megabytes used in the tablespace.

Creating A Tablespace

To create a new tablespace, choose Create from the Tablespace menu. The Create Tablespace property sheet appears. The Create Tablespace property sheet consists of the following pages:

- General (Basic UI mode)
- Extents (Advanced UI mode)

Tablespace Property Sheet: General Page

The General page of the Create Tablespace property sheet contains fields that define the tablespace. These fields are described below:

Name

Name of the tablespace to be created.

Enter the name of the new tablespace. The tablespace name can only contain characters from your database character set and can be at most 30 characters long.

Status

Status of the tablespace to be created. To select the status, click on the desired option.

- If the status is Offline, click the Online button to change the status to online.
- If the status is Online, click the Offline button and select one Normal, Immediate, or Temporary options from the list of values.
- If the status is Online and the database is version 7.1 or later, a Read-Only option is also available. If a tablespace has read-only status, the Make Writeable menu option is enabled. When you select the Make Writeable menu option, the tablespace becomes both readable and writeable and the status is changed to online.

Note: The Read-Only option is disabled if there any open transactions in the database or any active rollback segments in the tablespace.

If the tablespace status is Online, you can click the Offline button to change the status to offline with Offline Normal, Offline Immediate, or Offline Temporary options. Select one of the Offline options.

Datafiles

Multi-column scrolling list of the datafiles belonging to the tablespace. The columns include Name, Status, Size (M), and Used (M). These are the same columns as those in the Datafiles multi-column list. See *Datafiles Multi-Column List* on page 8-15.

Add

Displays the Create Datafile property sheet, which allows you to specify each new datafile belonging to the new tablespace. When you create a new datafile, the status column in the datafile list reads NEW. For a description of the Create Datafile property sheet, see *Creating a Datafile* on page 8-16.

Edit

Displays the Quick Edit Datafile property sheet, which allows you to edit the file specification for the datafile selected in the Datafiles scrolling list. You can also double-click on a datafile in the Datafiles scrolling list to display the Quick Edit Datafile property sheet.

This property sheet is the same as the Create Datafile property sheet. For a description of the Create Datafile property sheet, see *Creating a Datafile* on page 8-16.

Remove

Removes the datafile selected in the Datafiles scrolling list. You can only remove a datafile that has been newly created and has not been committed to a tablespace.

Type

Permanent: Specifies that the tablespace will be used to hold permanent database objects. This option is selected by default.

Temporary: Specifies that the tablespace will only be used to hold temporary objects (sort segments). No permanent object can reside in a temporary tablespace.

Attention: This option is only available for Oracle version 7.3 or greater. For more information on temporary tablespaces, see *Oracle Server Concepts*.

Tablespace Property Sheet: Extents Page

On the Extents page, you can specify the default storage parameters for all objects created in the tablespace. The Extents page contains the following fields:

Override Default Values

If this box is checked, you can edit all of the fields on the Extents page. If the box is not checked, the default value for each field displays and the values cannot be modified. This checkbox is only enabled during tablespace creation.

Note: If you do not enter an integral multiple of the operating system data block size when defining the size of extents, your entry is increased to the next multiple. If you do not enter a value, the default value remains. If a default value has not been explicitly specified, the field contains “Default.”

Initial Size

Size of the object’s first extent.

Enter the size of the initial extent. Use the unit button to specify either kilobytes or megabytes. If you do not specify a size, the default is the size of 5 data blocks.

Next Size

Size of the next extent to be allocated to the object.

Enter the size of the next extent. Use the unit button to specify either kilobytes or megabytes. If you do not specify a size, the default is the size of 5 data blocks.

Increase Size By

Percent by which each extent after the second grows over the previous extent.

Enter a value for percent increase. If you do not specify a value, the default is 50.

Note: It is recommended that 0 be entered to obtain extents of uniform size and to avoid fragmentation.

Minimum Number

Total number of extents to be allocated when an object is created in the tablespace.

Enter the minimum number of extents. If you do not specify a number, the default value is 1.

Maximum Number

Maximum number of extents that can be allocated to an object created in the tablespace.

Unlimited: When selected, allows you to create a number of extents that is only limited by the amount of contiguous free space in a tablespace.

Value: When selected, allows you to specify the maximum number of extents. If you do not specify a value, the default value applies. The default and maximum values depend on the data block size.

Showing Tablespace Dependencies

To display tablespace dependencies and dependents:

1. Select a tablespace from the tree list.
2. Choose Show Dependencies from the Tablespace menu. The Dependencies Viewer displays.

Editing a Tablespace

To alter an existing tablespace:

1. Select a tablespace in the Tablespace multi-column list using the right mouse button. This displays the context-sensitive menu.
2. Choose the Quick Edit menu option. The Quick Edit Tablespace property sheet appears.
3. Modify the property sheet parameters as necessary.
4. Click Apply.

You can also display tablespace information by clicking on a tablespace in the Storage Manager tree list.

The Quick Edit Tablespace property sheet contains the same elements as the Create Tablespace property sheet except as noted above. See *Creating A Tablespace* on page 8-7.

Attention: If you alter an object, such as a datafile named DATA1, at any location in the tree list, all instances of the object in the tree are changed.

Dropping a Tablespace

To drop an existing tablespace:

1. Select the tablespace to be dropped from the Tablespace tree list.
2. Choose Remove from the Tablespace menu. The Remove Tablespace alert box appears.
3. Click Yes.

When you drop a tablespace, all objects in the tablespace are dropped as well. Storage Manager also drops all referential integrity constraints from tables outside the tablespace that refer to primary or unique keys in the tables stored in the dropped tablespace.

Suggestion: Before dropping a tablespace, take it offline. This ensures that SQL statements that are currently running transactions do not access objects in the tablespace.

Note: You need to remove any OS files manually.

Adding a Datafile to a Tablespace

To add a datafile to an existing tablespace:

1. Select the tablespace from the Tablespace object list..
2. Select the desired tablespace from the drop-down menu and fill out the rest of the property sheet parameters.
3. Click the Create button.

You can also add a datafile to a tablespace through the Quick Edit/Create Tablespace property sheet or the Create Datafile property sheet.

Attention: When adding a datafile to either a new or existing tablespace, the Online and Offline option will be disabled.

Adding a Rollback Segment to a Tablespace

To add a rollback segment to an existing tablespace:

1. Select the tablespace from the Tablespace object list.
2. Choose Add Rollback Segment from the Tablespace menu. The Create Rollback Segment property sheet appears. See *Creating a Rollback Segment* on page 8-19.
3. Click Create.

You can also add a rollback segment to a tablespace through the Create Rollback Segment property sheet.

Changing a Tablespace to Online or Offline Status

To place a tablespace online, select the tablespace from the Tablespace tree list and choose Place Online from the Tablespace menu. The tablespace is placed online.

To take a tablespace offline, select the tablespace from the Tablespace object list and choose one of the Take Offline cascading menu options: Normal, Temporary, or Immediate.

You can also perform these tasks from the General page of the Tablespace property sheet.

Take Offline

Normal: Takes the tablespace offline in normal mode.

A checkpoint is performed for all datafiles in the tablespace (all of these datafiles must be available). You need not perform media recovery on this tablespace before placing it back online. You must use this option if the database is in NOARCHIVELOG mode.

Temporary: Takes the tablespace offline in temporary mode.

A checkpoint is performed for all online datafiles in the tablespace but does not ensure that all files can be written. Any offline files may require media recovery before you place the tablespace back online.

Immediate: Takes the tablespace offline in immediate mode.

Oracle does not ensure that the datafiles are available, and no checkpoint is performed. You must perform media recovery on the tablespace before placing it back online.

Making a Tablespace Read-Only

To change a writeable tablespace to read-only status:

1. Select the tablespace from the tree list.
2. Choose Make Read-Only from the Tablespace menu.

You can also change the tablespace status from the Quick Edit Tablespace property sheet.

Making a Tablespace Writeable

To change a read-only tablespace to writeable status:

1. Select the tablespace from the tree list
2. Choose Make Writeable from the Tablespace menu.

You can also change the tablespace status from the Quick Edit Tablespace property sheet that appears when you select a tablespace from the tree list.

Datafiles Folder

The Datafiles folder contains information about the datafiles in the database. The listing of the datafiles is arranged in a tree structure that can be expanded to show individual datafiles along with their respective path information.

For information about datafiles, see the *Oracle Server Concepts* and the *Oracle Server Administrator's Guide*.

Datafiles Multi-Column List

The columns of the Datafile multi-column list are described below:

Name

Name of the datafile and the path (truncated by ellipses if the path is too long).

Tablespace

Tablespace to which the datafile belongs.

Status

Status of the datafile: Online or Offline.

Size (M)

Size of the datafile in megabytes.

Used (M)

Amount (in megabytes) of data filling the datafile. Displayed as a bar chart.

Creating a Datafile

To create a new datafile, choose Create from the Datafile menu. The Create Datafile property sheet appears.

The Create Datafile property sheet contains the following pages:

- General (Basic UI mode)
- Auto Extend (Advanced UI mode)

Note: The Auto Extend page only appears if the database is version 7.2 or later.

Datafile Property Sheet: General Page

The General page of the Datafile property sheet allows you to modify existing datafile parameters or create a new datafile when the Create menu option is selected from the Datafile menu. The General page of the Create Datafile property sheet is described below:

Name

Name of the datafile to be created.

Enter the file and path name of the new datafile. The filename must be specified according to the conventions of your operating system.

When altering an existing datafile, you can rename the datafile by typing in a new name.

Attention: You must rename a datafile if you have changed the name of the corresponding operating system file or if you have moved the file to a new location. To rename a datafile, its tablespace must be offline.

Attention: When you rename a datafile using the Storage application, the name of the operating system file is not changed. The new filename is only associated with the tablespace. Before renaming the datafile using the Storage application, you must change the name of the file through your operating system.

Tablespace

Name of the tablespace to which the new datafile belongs. Use the drop-down list to choose the tablespace. You can only choose the tablespace when the datafile is created.

Status

Online: Specifies that the datafile be placed online.

Offline: Specifies that the datafile be place offline.

Note: When adding a datafile to either a new or existing tablespace, the Online and Offline option will be disabled.

Note: When altering a tablespace, the fields on this page cannot be modified if the datafile belongs to a read-only tablespace.

To determine the file size, select either the Use Existing File or New File Size option.

Size

File Size: Allows you to designate the file size of a new or existing file. Enter the size of the new or existing datafile. Use the unit buttons to specify either kilobytes or megabytes.

Reuse Existing File: Designates that the datafile already exists and should be reused.

Note: When altering a tablespace, the fields on this page cannot be modified if the datafile belongs to a read-only tablespace.

Datafile Property Sheet: Auto Extend Page

The AutoExtend page (available in Advanced UI mode) sets the Auto Extend feature for a datafile used in a database that is version 7.2 or later. The page consists of the following:

Enable Auto Extend

Determines whether the Auto Extend feature is enabled or disabled. Disabled is the default setting. If the box is not checked, the other fields on the page are disabled.

Increment

Determines the size and units of the increment size.

Maximum Extent

Determines the size of the maximum extent. You can select the Unlimited button to set the maximum extent size to Unlimited or select Value to specify the size and units. Unlimited is selected by default.

Note: If the datafile belongs to a read-only tablespace, all the fields on this page are disabled.

Creating a Datafile Like an Existing Datafile

To create a new datafile with parameters set like an existing datafile:

1. Select a datafile from the tree list.
2. Choose Create Like from the Datafile menu.

The Create Datafile property sheet appears with all parameters set except the name. See *Creating a Datafile* on page 8-16.

Editing an Existing Datafile

To edit an existing datafile:

1. Select the datafile from the tree list. The Datafile property sheet displays with all parameters set except the name. The property sheet is the same as the Create Datafile property sheet. See *Creating a Datafile* on page 8-16.
2. Modify the Datafile property sheet as desired.

You can also alter an existing datafile by selecting a datafile from the multi-column list using the right mouse button and then choosing Quick Edit from the context-sensitive menu.

Placing a Datafile Online

To place an existing datafile online:

1. Select a datafile from the tree list.
2. Select the Place Online menu option from the Datafile menu.

You can also change the online/offline status of the datafile by selecting the desired datafile and applying the change from the Datafile property sheet.

Taking a Datafile Offline

To take an existing datafile offline:

1. Select the desired datafile from the tree list
2. Select the Take Offline menu option from Datafile menu.

Attention: When adding a datafile to either a new or existing tablespace, the Online and Offline option will be disabled.

Rollback Segments Folder

The Rollback Segments folder contains information about the rollback segments in the database. The listing of the rollback segments is arranged in a tree structure that can be expanded to show individual rollback segments.

For information about managing rollback segments, see the *Oracle Server Concepts*, the *Oracle Server Administrator's Guide*, and the *Oracle Server SQL Reference*.

Rollback Multi-Column List

The columns of the Rollback multi-column list are described below:

Name

Name of the rollback segment.

Tablespace

Tablespace that contains the rollback segment.

Status

Status of the rollback segment: ONLINE, Offline, Needs Recovery, or Partly Available.

Size (M)

Space allocated in megabytes to the rollback segment.

High Water Mark (M)

The percentage of the datafile that has ever been filled with data. Displayed as a notch on a bar chart.

Creating a Rollback Segment

To create a new rollback segment, choose Create from the Rollback menu. The Create Rollback Segment property sheet appears.

The Create Rollback Segment property sheet consists of the following pages:

- General (Basic UI mode)
- Extents (Advance UI mode)

Create Rollback Segment Property Sheet: General Page

The General page of the Create Rollback Segment property sheet is described below:

Name

Name of the rollback segment to be created.

Enter the name of the new rollback segment. The rollback segment name can only contain characters from your database character set and can be at most 30 characters long.

Tablespace

Name of the tablespace in which to create the rollback segment.

Choose the tablespace from the drop-down list.

Public

Toggles between public and private rollback segments. A private rollback segment is acquired explicitly by an instance when the instance opens the database (Parallel Server option). A public rollback segment forms a pool of rollback segments that any instance requiring rollback segments can use.

Status

Online: Specifies that the rollback segment be placed online.

Offline: Specifies that the rollback segment be place offline.

Rollback Segment Property Sheet: Extents Page

On the Extents page you can specify the storage characteristics of the rollback segment. The Extents page contains the following fields:

Override Default Values

If this box is checked, you can edit all of the fields on the Extents page. If the box is not checked, the default value for each field displays and the values cannot be modified. This checkbox is only enabled during rollback Segment creation.

Initial Size

Size of the rollback segment's first extent.

Enter the size of the initial extent. Use the unit button to specify either kilobytes or megabytes. The default is the size of 5 data blocks.

Next Size

Size of the next extent allocated to the rollback segment.

Enter the size of the next extent. Use the unit button to specify either kilobytes or megabytes. The default is the size of 5 data blocks.

Optimal Size

Optimal size for the rollback segment. Optimal is not displayed for offline rollback segments.

Enter the value for Optimal. Use the unit button to specify either kilobytes or megabytes. Oracle tries to maintain the optimal size of the rollback segment by dynamically deallocating extents when their data is no longer needed for active transactions. A blank field assumes the default value.

The default value of Optimal is null. If Optimal is null, Oracle never deallocates the rollback segment's unused extents.

The value of Optimal can never be less than the space initially allocated to the rollback segment, as specified by the values of Initial Extent, Next Extent, and Minimum Extents.

Minimum Number

Total number of extents to be allocated when the rollback segment is created. This field is only enabled when creating a rollback segment.

Enter the minimum number of extents. The default and minimum value is 2.

Maximum

Unlimited: Specifies that an unlimited number of extents can be allocated to the rollback segment.

Value: Maximum number of extents that can be allocated to the rollback segment. The default is operating system dependent.

Note: The default and maximum values depend on the data block size.

Altering a Rollback Segment

To alter an existing rollback segment:

1. Select the rollback segment from the tree list. The Rollback property sheet appears.
2. Modify the property sheet settings as desired.
3. Click Apply to implement the changes.

You can also perform this operation using the Quick Edit popup menu option by selecting the rollback segment from the multi-column list using the right mouse button. The Quick Edit Rollback Segment property sheet is the same as the Create Rollback Segment property sheet. See *Creating a Rollback Segment* on page 8-19.

Note: Any changes you make on the Storage page apply to any subsequent extent allocations to the rollback segment, not existing extents.

Dropping a Rollback Segment

To drop an existing rollback segment:

1. Select the rollback segment to be dropped from the multi-column list.
2. Choose Remove from the Rollback menu. The Remove Rollback Segment alert box appears.
3. Click Yes.

Attention: You can only drop a rollback segment that is offline.

Shrinking a Rollback Segment

To shrink an existing rollback segment:

1. Select the rollback segment from the tree list.
2. Choose Shrink from the Rollback menu. The Shrink Rollback Segment dialog box appears.
3. Click OK.

You can also display the Shrink Rollback Segment property sheet by selecting a rollback segment from the multi-column list and choosing Shrink from the context-sensitive menu.

The Shrink Rollback Segment dialog box contains the following information:

Shrink to

Optimal Size: Shrink the rollback segment to an optimal size. The optimal value is determined by the value of the STORAGE parameter set when you originally created the rollback segment.

Size: Specify the number of bytes (K or M) in active extents in the rollback segment.

Attention: You can only shrink a rollback segment that is online.

The amount of rollback segment shrinkage depends on the following factors:

- available free space in the rollback segment.
- how active transactions are holding space in the rollback segment.

Changing a Rollback Segment to Online or Offline Status

To place a rollback segment online:

1. Select the rollback segment from the tree list.
2. Choose Place Online from the Rollback menu. The rollback segment is placed online.

To take a rollback segment offline:

1. Select the rollback segment from the tree list.
2. Choose Take Offline from the Rollback menu. The rollback segment is taken offline.

When you change the status of a rollback segment to offline, Oracle takes the rollback segment offline immediately if the rollback segment does not contain information necessary to roll back any active transactions. If the rollback segment *does* contain information for active transactions, Oracle makes the rollback segment unavailable for future transactions and takes it offline after all the active transactions are committed or rolled back.

Attention: Because Oracle does not take a rollback segment offline until all its active transactions have completed, there may be some delay before the status of the rollback segment is changed to Offline in the Rollback object list. Pending Offline displays in the object list to indicate that the rollback segment was taken offline while it was busy.

Managing Database Security

This chapter describes how to use Security Manager to control database security. With Security Manager, you can manage users, roles, and profiles. This chapter assumes that you have read [Chapter 7, "Overview of the Database Tools"](#) and are familiar with the interface elements of the database tools. The topics in this chapter are:

- [Starting Security Manager](#)
- [Security Manager Objects and Folders](#)
- [Security Manager Menus](#)
- [Users Folder](#)
- [Roles Folder](#)
- [Profiles Folder](#)

Starting Security Manager

To start Security Manager, click the Security icon in the Launch Palette or choose Security Manager from the Console Tools menu.

- If a connection to a database has already been made through the Console or an application, the Security window displays.
- If a valid database connection has not been made, the Login information dialog box displays. See *Connecting to an Instance* on page 1-23.

Note: You can change the database connection with the Change Database Connection option in the File menu. For more information, see *Application Menus* on page 7-10.

After Security Manager has successfully connected to a database, the Users, Roles, and Profiles folders display in a tree list on the left side of the Security window. These folders are located in the database folder which displays the name of the database that the application is connected to.

Figure 9–1 Security Manager



The display on the right side of the window is determined by the object selected on the left side of the screen. The right side may contain a multi-column list, property

sheet, or other information. An example of a Security Manager window is shown in [Figure 9-1, "Security Manager"](#).

Refer to the following sections:

- *Application Objects* on page 7-3.
- *Multi-Column Lists* on page 7-14.
- *Dialog Boxes* on page 7-15.
- *Property Sheets* on page 7-17.

Security Manager Menus

Security Manager includes the standard menus, File, View, Log, and Help, plus the User, Profile, and Role menus. The options for each of these menus are described in this chapter. For information on the standard menus, see *Application Menus* on page 7-10.

Context-sensitive menus may also be active when you press the right mouse button to select a specific object from the tree list or the multi-column list. This feature provides quick access to a subset of the menu options provided in the menu bars.

User Menu

The User menu contains the following menu options:

Create

Creates a new user.

Create Like

Creates a new user based on the selected user in the tree list.

Remove

Deletes the selected user from the tree list.

Revoke Privilege

Removes a selected privilege or role.

Show Dependencies

Displays database objects that rely on a selected user and any objects that the selected user relies on.

Add Privileges to User

Adds multiple privileges to one or more users.

Change Account Status (Oracle 8)

Unlock: Unlocks the user's account and enables access to the account.

Lock: Locks the user's account and disables access to the account.

Expire: Expires the user password.

Role Menu

The Role menu contains the following menu options:

Create

Creates a new role.

Create Like

Creates a new role based on the selected role.

Remove

Deletes the selected role.

Revoke Privilege

Removes a privilege or role from a role.

Show Dependencies

Displays database objects that rely on a selected role and any objects that the selected role relies on.

Add Privileges to Roles

Adds privileges or roles to roles.

Note: These menu options are enabled depending on the object selected.

Profile Menu

The Profile menu contains the following menu options:

Create

Creates a new profile.

Create Like

Creates a new profile that is based on the selected profile.

Remove

Deletes the selected profile.

Show Dependencies

Displays database objects that rely on a selected profile and any objects that the selected profile relies on.

Assign Profile to Users

Assigns a profile to a specific user.

Note: These menu options are enabled depending on the object selected in the tree list.

Security Manager Objects and Folders

The objects in the tree list are identified by various icons. In the listing:

- A folder icon identifies an object type folder.
- A single person icon depicts a user.
- A mask icon signifies a role or subrole.
- A scroll icon depicts a profile.
- An open door icon identifies a system privilege.
- A suitcase shows and identifies an object privilege.

Attention: Roles, Object Privileges, and System Privileges icons appear with a key overlay if these objects have been granted using the Admin option/Grant option.

Users Folder

The User object type folder contains information about the users in the database arranged alphabetically in a tree structure. An individual user can be expanded to show the roles, system privileges, and object privileges granted to the user.

When you select:

- The Users group folder, a multi-column list displays a row of summary information for each of the users in the database.
- An individual user, the detailed settings for the user are displayed in a property sheet.
- The Roles Granted, System Privileges Granted, or Object Privileges Granted folder under an individual user, a multi-column list of the privileges or roles granted to the user displays. See *Privileges Multi-Column Lists* on page 9-17.

For more information about users, see the *Oracle Server Concepts*, the *Oracle Server Administrator's Guide*, and the *Oracle Server SQL Reference*.

Users Multi-Column List

A Users multi-column list displays when a User folder is selected in the tree list. The list contains a row of summary information for each of the users in the Users folder.

If you select an individual User icon, and that icon is also on the main branch of the Database folder, the columns of the multi-column list summarize all information from the General page of the Create User property sheet. For more information on these columns, see the description of the Create User property sheet in *Creating a User* on page 9-8.

Suggestion: If a multi-column list is wider than the its window display area, you can increase the viewing area by resizing the application window or dragging the splitter between the left and right sections of the window.

Creating a User

To create a new user:

1. Choose Create from the User menu. The Create User property sheet appears.
2. Fill in the appropriate property sheet information.
3. Click the Create button after specifying the requisite parameters.

The Create User property sheet consists of the following pages when in Advanced mode:

- General (Basic UI mode)
- Roles/Privileges (Advanced UI mode)
- Object Privileges (Advanced UI mode)
- Quotas (Advanced UI mode)

Create User Property Sheet: General Page

The General page allows you to specify a user's name (when creating a new user), their default profile, authentication method, and the default tablespace. The Create User property sheet contains the following:

Name

The name of the user to be created. Enter the name of the new user. The username can only contain characters from your database character set and can be at most 30 bytes long.

Profile

The profile assigned to the user. Use the drop-down list to choose the profile you want to assign to the user. The DEFAULT profile is assigned if you do not make a selection.

Authentication

The method Oracle uses to authenticate the user.

Global: Specifies that the user be identified globally amongst multiple databases. The global authorization option is only available with Oracle 8 databases.

External: Specifies that the operating system verify the user.

Password: Specifies that a password be required for login. Enter the password in the adjacent text entry field. Enter the password again in the Confirm text entry field for verification.

Expire Now: Forces the user's password to expire immediately. If you create a new user with this option selected, the user's password must be changed during the first attempted login. This feature is available for Oracle8 databases only.

Tablespaces

The user's default and temporary tablespaces.

Default: Use the drop-down list to choose the default tablespace for user-created objects.

Temporary: Use the drop-down list to choose the tablespace for the user's temporary segments.

Status (Oracle 8 only)

Active status of the user's account.

Lock: Locks the user's account and prevents further access.

Unlock: Unlocks the user's account and enables access to the account.

Create User Property Sheet: Roles/Privileges Page

On the Privileges page of the Create User property sheet, you can specify the system privileges and roles assigned to the user. The Privileges page contains the following:

Privilege Type

A drop-down list containing System Privileges and Roles. Your selection in the drop-down list determines what is displayed on the rest of the page.

Available

List of available roles or system privileges available for assignment to a user.

Roles: If you selected Roles as the privilege type, the roles that you are allowed grant to a user display in a scrolling list. These are roles you have created and roles you have been granted with the Admin Option.

Roles that you add to the user are assigned as default roles unless you change the specification by clicking on the role's entry in the Default column.

You must add the roles with the Admin Option in a separate operation from the roles you want to add without the Admin Option.

Note: When you grant the DBA and RESOURCE roles to a user or role with Oracle7 release 7.2.2 or later, the user or role is also granted the UNLIMITED TABLESPACE system privilege. When you revoke either role from a user or role, the UNLIMITED TABLESPACE system privilege is also revoked. The UNLIMITED TABLESPACE can also be revoked independent of the DBA and RESOURCE roles.

System Privileges: If you select System Privileges as the privilege type, system privileges that you are able to grant to a user display in a scrolling list. These are the system privileges the you have been granted with the Admin Option. If you have the GRANT ANY PRIVILEGE system privilege, all privileges are listed. Select the privileges that you want to add to the user.

Attention: If you want to grant the Admin or Grant option of a current privilege or role, you must add the privilege or role with Admin or Grant option specified as you want.

Up and Down Arrows

Up Arrow adds roles or system privileges that are selected in the Available list to the Granted spreadsheet.

Down Arrow removes roles or system privileges that are selected in the Granted spreadsheet.

Granted

An editable spreadsheet displaying roles or system privileges assigned to a user. New additions (selected but not yet applied) are identified by a hand with a plus sign. When creating a user, the spreadsheet consists of three columns:

System Privilege or Role: Name of the role or system privilege.

Admin Option: When checked, allows the user to grant the system privileges to other users or roles. By default, Admin Option is disabled. You enable the Admin Option by clicking on the spreadsheet entry. In this case, the "X" becomes a check.

Default: (Users property sheet only): When checked, establishes the role as a default for the user upon system logon.

When creating a Role, this spreadsheet consists of two columns: System Privilege or Role and Admin Option.

Create User Property Sheet: Object Privileges Page

The Object Privileges page, available in Advance UI mode, allows you to grant or revoke privileges for a specific user on schema objects. This page contains the following:

Objects

A tree listing of schemas in the database and objects in the schemas displays in the Object window. Click on the '+' to the left of a folder icon to display the object types contained in the schema, then click on the '+' to the left of the object type to display the actual objects.

Select the object from the tree list that you want to grant privileges for. After the object is selected, the available privileges for the object are displayed to the right in the Available Privileges scrolling list.

You can grant an object privilege that you have been granted with the Grant Option. If you are the owner of the object, you can grant all privileges on the object. Select the privileges you want to grant for the selected object. The scrolling list includes the privileges you can grant on this object.

Grant Option box to allow the user to grant the object privilege to other users and roles.

Click the Add button to add the selected object privileges to the user.

Available Privileges

Displays privileges available for the schema object selected in the tree list.

Up and Down Arrows

Down Arrow adds privileges that are selected in the Available Privileges list to the Granted Object Privileges spreadsheet.

Up Arrow removes privileges that are selected in the Granted Object Privileges spreadsheet.

Granted Object Privileges

An editable spreadsheet displaying object privileges to be made available to a user. New additions (selected but not yet applied) are identified by a hand with a plus sign.

When creating a new user, the spreadsheet consists of two columns indicating the name of the object privilege and whether or not the Grant Option is specified for that privilege.

When enabled, the Grant Option allows the user to grant the specific object privilege to other users and roles. By default, this option is disabled. To enable the grant option, click on the specific spreadsheet entry. The "X" is replaced with a check.

When creating a role, the spreadsheet consists of a single Object Privilege column.

Create User Property Sheet: Quotas Page

On the Quotas page of the Create User property sheet, you can specify the tablespaces in which the user can allocate space and the maximum amount of space the user can allocate within each tablespace. This page is available in Advanced mode. The Quotas page contains the following items:

Quota Details

Scrolling list of the tablespaces in the database and the maximum amount of space the user has been allowed in each tablespace. The list can be sorted on the Tablespace or Quota Size column.

To specify a quota size for a tablespace, select the tablespace in the scrolling list and specify a quota size by clicking on the None, Unlimited, or Value button.

None

Click None if you do not want the user to have any quota on the selected tablespace.

Unlimited

To specify an unlimited quota for the tablespace, click the Unlimited button. With an unlimited quota, the user can allocate an unbounded amount of space in the tablespace.

Value

To specify a specific quota, click Value and enter a quota value in the adjacent text entry field. Select the K or M button to specify Kilobytes or Megabytes.

Note: If the user has been granted the UNLIMITED TABLESPACE System Privilege, the Quota Details option is disabled.

Create Like User

To create a new user with the same attributes as an existing user:

1. Choose Create Like from the User menu. The Create User property sheet appears with all parameters specified except the Name.
2. Modify any property sheet parameter for the new user as necessary.
3. Click Create.

You can also perform this operation by selecting a user from the tree list and then choosing the Create Like menu option. You must enter the name of the new user and enter a new password if the Password button is selected.

The format and content of the Create Like property sheet is identical to the Create User property sheet. Refer to *Creating a User* on page 9-8 for information about the property sheet.

Altering a User

To change the characteristics of a user:

1. Select the user from the tree list to display the user details property sheet.
2. Modify the property sheet parameters as necessary.
3. Click Apply.

You can also display the Quick Edit property sheet by selecting a user from the multi-column list and choosing Quick Edit from the context sensitive menu. The Quick Edit property sheet is identical to the details property sheet.

The details/Quick Edit property sheet is identical in format and content to the Create User property sheet except that the name field is read-only. See *Creating a User* on page 9-8 for information about the property sheet.

Suggestion: If you want to add privileges or roles to multiple users, use the Add Privileges and Roles to Users menu item. See *Adding Privileges or Roles to Users* on page 9-15.

Attention: If you alter an object, such as a user named DAVE or a role named CLERK, in any location of the tree list, all instances of the object in the tree are changed.

Removing a User

If you no longer need a particular user in your database, you can remove the user by selecting the user to be dropped from the Users folder in the tree list and choosing Remove from the User menu. The Remove User alert box appears.

The Remove User alert box indicates if the user still owns any objects. If you remove a user who owns objects, Security Manager:

- Drops all of the objects in the user's schema.
- Drops any referential integrity constraints in other schemas that refer to the dropped user's tables.
- Invalidates any views or synonyms for objects in the dropped user's schema.
- Invalidates any stored procedures, functions, or packages that query objects in the dropped user's schema.
- Does not drop snapshots on tables or views in the dropped user's schema.
- Does not drop any roles created by the user.

Adding Privileges or Roles to Users

To add multiple roles and grant multiple system or object privileges to users:

1. Choose Add Privileges to Users from the User menu or one of the context-sensitive menus. The Add Privileges to Users dialog box displays with a scrolling list of users is displayed in the top half of the dialog box.
2. Select the users in the list that you want to add privileges or roles to.
3. Select Roles, System Privileges, or Object Privileges from the Privilege Type drop-down list. The display in the bottom half of the dialog box varies according to your selection.
4. Click OK to commit the changes you have made. For details on the dialog box command buttons, see *Dialog Boxes* on page 7-15.

Roles

If you selected Roles, the roles that you can grant to a user display in a scrolling list. These are roles you have created and roles you have been granted with the Admin Option. If you have the GRANT ANY ROLE system privilege, all roles are listed.

Select the roles that you want to add to the selected users.

Attention: The roles that you add to the users are assigned as default roles unless you change the specification on the Default Roles page of the Alter property sheet of each user.

Click the With Admin Option box to allow the user to grant the role to other users or roles. If you grant a role with the Admin Option, the user can also alter or drop the role.

You must add the roles with the Admin Option in a separate operation from the roles you want to add without the Admin Option.

Note: When you grant the DBA and RESOURCE roles to a user or role with Oracle7 release 7.2.2 or later, the user or role is also granted the UNLIMITED TABLESPACE system privilege. When you revoke either role from a user or role, the UNLIMITED TABLESPACE system privilege is also revoked. The UNLIMITED TABLESPACE can also be revoked independent of the DBA and RESOURCE roles.

From the SQL Worksheet, use the GRANT command to grant privileges on a column in a table or view. For information about the GRANT command, see the *Oracle Server SQL Reference*.

System Privileges: A scrolling list of the system privileges that you are able to grant to users. These are the system privileges you have been granted with the Admin Option. If you have the GRANT ANY PRIVILEGE system privilege, all privileges are listed.

Select the privileges that you want to add to the selected users. Click the With Admin Option box to allow the user to grant the system privileges to other users or roles.

Attention: You must add the privileges with the Admin Option in a separate operation from the privileges you want to add without the Admin Option.

Object Privileges

A tree listing of schemas in the database and objects in the schemas displays in the Object window. Click on the '+' to the left of a folder icon next to display the object types contained in the schema and then click on the '+' to the left of the object type folder to display available objects. Select the objects that you want to grant privileges for.

After the object is selected, the available privileges for the object are displayed to the right in the Privileges scrolling list.

You can grant an object privilege that you have been granted with the Grant Option. If you are the owner of the object, you can grant all privileges on the object. Select the privileges you want to grant for the selected objects.

Attention: You must add the privileges with the Grant Option in a separate operation from the privileges you want to add without the Grant Option.

Removing Privileges or Roles from a User

To remove the roles or privileges that are currently assigned to a user:

1. From the Users folder in the tree list, click on the '+' to the left of the user to expand the Roles Granted, System Privileges Granted, and Object Click on the '+' to the left of the folder icon to display the privileges or roles that have been assigned to the user.
2. Select the privilege or role that you want to remove from a user.
3. Choose Revoke Privilege from the User menu or Revoke from the the context-sensitive menu to remove the selected privilege or role from the user.

Privileges Multi-Column Lists

The Privileges multi-column scrolling list displays when a Roles Granted, System Privileges Granted or Object Privileges Granted object type folder is selected in the tree list. The list contains information about privileges assigned to the user or role.

Roles Granted

The columns in the Roles Granted list include:

Role

Name of the role.

Admin option

Whether the role was granted with the Admin option to the user or role.

Default

Indicates the role as a default for the user upon system logon

System Privileges Granted

The columns in the System Privileges Granted list include:

System Privilege

Name of the system privilege.

Admin Option

Whether the privilege was granted with the Admin option to the user or role.

Object Privileges Granted

The columns in the Object Privileges Granted list include:

Object Privilege

Name of the object privilege.

Grant option

Whether the privilege was granted with the Grant option to the user.

For more information on these columns, see the description of the Create User property sheet in the section, *Creating a User* on page 9-8.

Roles Folder

The Roles object type folder contains information about the roles defined in your database arranged alphabetically in a tree structure. An individual role can be expanded to show the system privileges, object privileges, and roles granted to the role.

When you select:

- the Roles folder, a multi-column list displays available roles and authentication, if necessary.
- an individual role, the Roles property sheet displays the detailed settings for the role.
- the Roles Granted, System Privileges Granted, or Object Privileges Granted folder under an individual role, a multi-column list of the privileges or roles granted to the role displays. See *Privileges Multi-Column Lists* on page 9-17.

Roles are named groups of privileges granted to users or other roles. For information about managing roles, see the *Oracle Server Concepts*, the *Oracle Server Administrator's Guide*, and the *Oracle Server SQL Reference*.

Roles Multi-Column List

A Roles multi-column list displays when a Roles or Roles Granted folder is selected in the tree list. The multi-column scrolling list contains a row of summary information for each of the roles in the roles folder.

If the folder is named Roles Granted and is contained in a user or role, the list only contains information about roles assigned to the user or role. The columns of the Roles multi-column list include:

Role

Name of the role.

Authentication

Type of authentication used: none, external, password, or global.

For more information on these columns, see the description of the Create User property sheet in the section, *Creating a User* on page 9-8 .

Suggestion: If a multi-column list that is wider than the window display area, you can increase the viewing area by resizing the application window or dragging the splitter between left and right section of the window.

Creating a Role

To create a new role:

1. Choose Create from the Role menu or one of the context-sensitive menus. The General page of the Create Role property sheet appears.
2. Fill in the property sheet with the desired parameters.
3. Click the Create button to apply the changes you have made to the property sheet and create the new role. For details on the property sheet command buttons, see *Property Sheets* on page 7-17 .

The Create Role property sheet contains the following pages.

- General (Basic UI mode)
- Roles/Privileges (See *Create User Property Sheet: Roles/Privileges Page* on page 9-10)
- Object Privileges (See *Create User Property Sheet: Object Privileges Page* on page 9-12)

Create Role Property Sheet: General Page

The General page allows you to enter the following information:

Role

Name of the role to be created. Enter the name of the new role.

Authentication

Method used to enable the role.

None specifies that a user granted the role may enable it without specifying a password.

Global: Specifies that a user granted the role may enable it globally amongst multiple databases. The global authorization option is only available with Oracle 8 databases.

External: Specifies that the operating system or an external security utility to verify the role.

Password: Specifies that a password is required to enable the role. Enter the password in the `Enter Password` entry field. Enter the password again in the `Confirm Password` entry field to verify the new password.

Create Like Role

If you want to create a new role with the similar parameters as an existing role:

1. Select the desired role from the tree list.
2. Choose Create Like menu option from the Role menu.
3. Modify the property sheet as desired.
4. Click the Create button.

You can also perform this operation by selecting a role in the tree list and choosing the Create Like menu option from the context-sensitive menu. You must enter the name of the new role and enter a new password if the Password button is selected.

This property sheet is identical to the Create Role property sheet. Refer to *Creating a User* on page 9-8 for information about Create Role property sheet.

Modifying a Role

To alter the property sheet information for an existing role:

1. Select the role to be altered from the tree list. The property sheet for that role appears.
2. Modify the property sheet as desired.
3. Click the Apply Button.

You can also modify a role by selecting a role from the Role multi-column list with the right mouse button to call up the context-sensitive menu. Select Quick Edit to bring up the Quick Edit property sheet appears.

The Role property sheet is identical to the Create Role property sheet except that the name is read-only. Refer to *Creating a User* on page 9-8 for information about the property sheet.

If you want to add privileges or roles to multiple roles, use the Add Privileges to Roles menu item. See *Adding Privileges or Roles to Roles* on page 9-22.

If you alter an object, such as a user named DAVE or a role named CLERK, in any location of the tree list, all instances of the object in the tree list are changed.

Removing a Role

To remove a role that is no longer needed:

1. Select a role in the tree list.
2. Choose Remove from the Role menu. The Remove Role alert box appears.
3. Click Yes to remove the role.

You can also remove a role by selecting the role to be dropped from the Roles object type folder in the tree list and choosing Remove from the context-sensitive menu.

Adding Privileges or Roles to Roles

To assign subroles and grant individual privileges to multiple roles, or add roles and grant privileges to multiple roles:

1. Choose Add Privileges to Roles from the Role menu to display the Add Privileges to Roles dialog box. A scrolling list of roles is displayed in the top half of the dialog box.
2. Select the roles in the list that you want to add privileges or roles to.
3. Select System Privileges, Object Privileges, or Roles from the drop-down list. The display in the bottom half of the dialog box varies according to your selection.

Roles

If you selected Roles, the roles that you can grant to a role display in a scrolling list. These are roles you have created and roles you have been granted with the Admin Option. If you have the GRANT ANY ROLE system privilege, all roles are listed.

1. Select the roles that you want to add to the role.
2. Click the With Admin Option box to allow the role to grant the role to other users or roles. If you grant a role with the Admin Option, the role can also alter or drop the role.
3. Click the Apply button to add the selected roles to the role.

Attention: You must add the roles with the Admin Option in a separate operation from the roles you want to add without the Admin Option.

System Privileges

A scrolling list of the system privileges that you are able to grant to a role. These are the system privileges you have been granted with the Admin Option. If you have the GRANT ANY PRIVILEGE system privilege, all privileges are listed.

1. Select the privileges that you want to add to the role.
2. Click the With Admin Option box to allow the role to grant the system privileges to other users or roles.
3. Click the Add button to add the selected system privileges to the role.

Attention: You must add the privileges with the Admin Option in a separate operation from the privileges you want to add without the Admin Option.

Object Privileges

A tree listing of schemas in the database and objects within the schemas displays in the Object window. Click on the '+' to the left of a folder icon to display the objects contained in the schema, then select the objects that you want to grant privileges for.

After the object is selected, the available privileges for the object is displayed to the right in the Privileges scrolling list.

1. Select the privileges you want to grant for the selected object. The scrolling list includes the object privileges you can grant on an object.
2. Click the Apply button to add the selected object privileges to the role.
3. Click on the Apply button to save any changes you have made in the dialog box. For details on the dialog box command buttons, see *Dialog Boxes* on page 7-15.

Removing Privileges or Roles from a Role

To remove privileges or subroles that are assigned to a role in the Roles folder:

1. From the tree list, select the privilege or subrole that you want to remove from a role. If necessary, click on the '+' to the left of the folder icon to display the privileges or subroles that have been assigned to the role.
2. Choose the Remove from Role menu.

To conveniently remove multiple privileges or subroles from a single role, use the appropriate property sheet.

Profiles Folder

The Profiles folder contains information about the profiles defined for the database arranged alphabetically in a tree structure. An individual profile can be expanded to show the users that have been assigned the profile.

When you select:

- the Profiles object type folder, a multi-column list displays a row of summary information for each of the profiles in the database.
- an individual profile, the edit property sheet displays the detailed settings for the profile.

A profile is a set of limits on database resources. When you assign a profile to a user, that user cannot exceed the limits set in the profile.

Oracle automatically creates a default profile named DEFAULT. The DEFAULT profile initially defines unlimited resources. You can alter the DEFAULT profile to change any of its resource limits.

Any user who is not explicitly assigned a profile is subject to the limits defined in the DEFAULT profile. Also, if the profile that is explicitly assigned to a user omits a limit for a resource or specifies the value DEFAULT for a limit, then the user is subject to the limit on that resource as defined in the DEFAULT profile.

Attention: The initialization parameter RESOURCE_LIMIT must be set to TRUE to enforce the limits set in database profiles. For more information, see the *Oracle Server Reference*.

For more information about profiles, see the *Oracle Server Concepts*, the *Oracle Server Administrator's Guide*, and the *Oracle Server SQL Reference*.

Profile Multi-Column Lists

The profiles multi-column list displays when the Profiles folder is selected in the tree list. The scrolling list contains a row of summary information for each of the profiles in the Profiles folder.

The columns include all the fields on the pages of Create Profile property sheet. For more information on these columns, see *Creating a Profile* on page 9-25.

Creating a Profile

To create a profile:

Choose Create from the Profile menu. The Create Profile property sheet appears. This property sheet contains the General and Password pages.

Create Profile Property Sheet: General Page

Name

This field allows you to enter the name of a new profile.

Details

These fields determine the amount of time allocated to the CPU per Session, CPU per Call, Connect Time, and Idle Time for this profile. The fields are:

CPU/Session: Total amount of CPU time allowed in a session. The limit is expressed in seconds.

CPU/Call: Maximum amount of CPU time allowed for a call (a parse, execute, or fetch). The limit is expressed in seconds.

Connect Time: Maximum elapsed time allowed for a session. The limit is expressed in minutes.

Idle Time: Maximum idle time allowed in a session. Idle time is a continuous period of inactive time during a session. Long-running queries and other operations are not subject to this limit. The limit is expressed in minutes.

You can enter a value in a field or choose from the drop-down list adjacent to the field. Click on the down-arrow to display the list. The drop-down list provides the following choices:

Database Services

These fields determine the database services allocated to this profile. The fields are:

Concurrent Sessions: Maximum number of concurrent sessions allowed for a user.

Reads/Session: Total number of data block reads allowed in a session. The limit includes blocks read from memory and disk.

Reads/Call: Maximum number of data block reads allowed for a call (a parse, execute, or fetch) to process a SQL statement.

Private SGA: Maximum amount of private space a session can allocate in the shared pool of the System Global Area (SGA). The Private SGA limit applies only if you are

using the multi-threaded server architecture. The limit is expressed in kilo bytes (KBytes).

Composite Limit: Total resource cost for a session. The resource cost for a session is the weighted sum of the CPU time used in the session, the connect time, the number of reads made in the session, and the amount of private SGA space allocated.

You can enter a value in a field or choose from the drop-down list adjacent to the field. Click on the down-arrow to display the list. The drop-down list provides the following choices:

Attention: In the SQL Worksheet, you can use the SQL command ALTER RESOURCE COST to specify the weights for the resources in the Composite Limit. For information about the ALTER RESOURCE COST command, see the *Oracle Server SQL Reference*.

Drop-down List Selections

Default: Use the limit specified for this resource in the DEFAULT profile.

Unlimited: The user's access to this resource is unlimited.

Values: Select one of the existing values. The default values vary by field and are common values for the field. If you have entered a value in the field, that value appears in the drop-down list.

Create Profile Property Sheet: Password Page (Oracle 8)

The Password page allows you to set account password parameters. This page consists of the following:

Expire Password

Expire in: Limits the number of days after which a password expires. Select a value from the drop-down list, or enter a specific value.

Lock: Limits the number of days during which a password can be changed following the first successful login after password expiration.

Keep Password History

Keep: Specifies the number of times a password must be changed before it can be reused. Keep for is disabled if a value is specified in this field.

Keep for: Limits the number of days before a password can be reused after it expires. Keep is disabled if a value is specified in this field.

Enforce Password Complexity

Allows a PL/SQL routine to be used for password verification when users who are assigned this profile log into a database. This PL/SQL routine must be locally available for execution on the database to which this profile applies.

Oracle provides a default script (utlpwdmg.sql), however, you can also create your own routine, or use third-party software as an alternative. The password verification routine must be owned by SYS.

NULL (no password verification) is set by default.

Lock account on failed logon

Lock after: Limits the number of failed logon attempts allowed before a user is locked out from the account.

Lock for: Specifies the number of days the account is locked after failing the specified number of logon attempts. If UNLIMITED is specified, only the database administrator can unlock the account.

Note: If Default is selected for any password options, values defined in the Default profile are used.

Create Like Profile

To create a new profile that has identical parameter settings to an existing profile:

1. Select the profile to be copied from the tree list.
2. Choose Create Like from the Profile menu. The Create Profile property sheet appears.
3. Specify a profile name and modify the property sheet parameters if necessary.
4. Click the Create button.

You can also perform this operation by selecting a profile from the tree or multi-column list and choosing the Create Like menu option. You must enter the name of the new profile.

The Create Like property sheet is identical to the Create Profile property sheet. See *Creating a Profile* on page 9-25 for information about the property sheet.

Altering a Profile

To alter the resource limits for an existing profile:

1. Select the profile to be altered from the tree list. The Profile property sheet appears.
2. Modify the resource limits as necessary.
3. Click the Apply button.

You can also use the Quick Edit property sheet to modify a profile by selecting a profile from the multi-column list using the right mouse button and choosing Quick Edit from the context-sensitive menu.

The Quick Edit property sheet is identical to the Create Profile property sheet except that the name field is read-only. See *Creating a Profile* on page 9-25 for information about the property sheet.

In the SQL Worksheet, you can use the SQL command ALTER RESOURCE COST to specify the weights for the resources in the Composite Limit. For information about the ALTER RESOURCE COST command, see the *Oracle Server SQL Reference*.

Showing Profile Dependents and Dependencies

To display dependents for a particular profile:

1. Select a profile from the tree list.
2. Choose Show Dependencies from the Profiles menu.

Dependencies are organized by user. Expanding individual users in the Dependencies Viewer displays the schema objects, roles, and privileges associated with that user.

Removing a Profile

To remove a profile that is no longer needed:

1. Select the profile to be deleted from the tree list.
2. Choose Remove from the Profile menu. The Remove Profile alert box appears.

The Remove Profile alert box indicates if the profile you wish to drop is assigned to any users. If you drop a profile that is assigned to users, Security Manager assigns the DEFAULT profile to them.

3. Click OK.

Note: You cannot drop the DEFAULT profile.

Assigning a Profile to Users

To assign a profile to multiple users in the database:

1. Choose the Assign Profile to Users menu option from the Profile menu. The Assign Profile dialog box appears.
2. Select the profile that you want to assign from the drop-down list.
3. In the scrolling list, select the users that you want to assign the profile to.
4. Click the Apply button to assign the selected profile to the user(s). You can click OK to assign the profile and close the dialog box.

Managing Instances and Sessions

This chapter describes how to use Instance Manager to manage instances and sessions. With Instance Manager, you can start up and shut down a database, view and edit the values of initialization parameters, resolve in-doubt transactions, and manage users' sessions. This chapter assumes that you have read [Chapter 7, "Overview of the Database Tools"](#) and are familiar with the interface elements of the database tools.

This chapter contains information on the following topics:

- [Starting Instance Manager](#)
- [Instance Manager Menus](#)
- [Database Object](#)
- [Initialization Parameters Folder](#)
- [Stored Configurations](#)
- [Sessions Folder](#)
- [In-Doubt Transactions Folder](#)

Note: If you have SYSOPER privileges, you can only startup and shutdown the database from the Console using the context-sensitive menu.

Starting Instance Manager

To start Instance Manager, click on the Instance icon in the Launch Palette or choose Instance Manager from the Console Tools menu.

- If you have already connected to a database, the Instance window displays.
- If a valid database connection has not been made, the Login Information dialog box displays. See *Connecting to an Instance* on page 1-23.

After Instance Manager has successfully connected to a database, the Initialization parameters, Stored Configurations, In-Doubt Transactions, and Sessions folders display in a tree list on the left side of the Instance window. These folders are contained in the database folder which displays the name of the database the application is currently connected to.

The display on the right side of the window is determined by the objects selected on the left side of the screen. The right side may contain a multi-column scrolling list, property sheet, or dialog box.

An example of an Instance Manager window is shown in Figure 10-1.

Figure 10–1 *Instance Manager*



For general information, see:

- *Application Objects* on page 7-3
- *Multi-Column Lists* on page 7-14
- *Dialog Boxes* on page 7-15
- *Property Sheets* on page 7-17

Instance Manager Menus

Instance Manager includes the three standard menus, File, View, and Help, and the Database, Sessions, Transactions, and Configuration menus. The options for each of these menus are described below. For information on the standard menus, see *Application Menus* on page 7-10.

Database Menu

The Database menu provides access to all options pertaining to Oracle instance startup and shutdown.

Startup

Displays the Startup dialog box that provides startup options and parameter file/configuration selection for instance startup.

Shutdown

Displays the Shutdown Options dialog box that provides standard options for instance shutdown.

Mount

Mounts an Oracle database to an instance that has been started.

Open

Opens an Oracle database that has been mounted but currently is not open.

Archive Log

Toggles ARCHIVELOG mode on or off. If the database is mounted and open, Instance Manager displays the Shutdown Options dialog box allowing you to shut down the database before changing the ARCHIVELOG mode. If the database is not mounted, Instance Manager asks if you want to open the database in a mounted state.

Automatic Archival

Toggles automatic archiving of online redo log files on or off.

Note: Changes to archive log mode only affects the current instance. For permanent changes, you must edit the initialization parameter file.

Manually Archive

Allows you to manually archive online redo log files.

Current: Archive the current online redo log file group and force a log switch.

All: Archive all redo log file group that is full but has not been archived.

Next: Archive the next archive redo log file group that is full but has not been archived.

Sessions Menu

The Sessions Menu provides access to all options that apply to database session management.

Disconnect

Disconnects any session that is selected from the Sessions container in the tree list.

Immediate: Kills the selected session immediately.

Post Transaction: Wait until all database transactions are complete before killing the selected session.

Restrict

Restricts database access only to those users with the RESTRICTED SESSION system privilege.

Allow All

Permits database access to all users.

Transactions Menu

The Transactions menu provides access to all options that apply to in-doubt transactions.

Force Commit

Commits any transaction selected from the In-Doubt Transactions folder in the tree list.

Force Rollback

Rolls back any transaction selected from the In-Doubt Transactions folder in the tree list.

Configuration Menu

The Configuration menu allows you to manage any configurations selected from the Stored Configurations folder in the tree list.

Remove

Deletes any configuration selected from the tree list.

Export to File

Exports any configuration selected from the tree list to an ASCII file.

Attention: If you alter an object contained in another object, all instances of the object in the database are changed.

Database Object

Database Object The database object displays the name of the target database. When selected, the database property sheet displays providing you with information regarding database status and archive mode. This property sheet contains the following information.

Database Property Sheet: Status Page

The Status page contains information about the status of the current database, database version, and any installed options. This page also allows you to change the state of the database by selecting one of the Database State options and clicking Apply.

Database State

Shutdown: Database down.

Instance Started: Instance started but no database mounted.

Database Mounted: Database mounted but not open.

Database Open: Database mounted and open.

Database Version

Database version and any installed options.

Attention: If the database is not running, the message "ORACLE not available" is displayed.

Database Property Sheet: Information Page

The Information page contains displays the current state of redo log archival and the system global area (SGA).

Archive Information

Log Mode: Specifies whether the connected database is in ARCHIVELOG mode or NOARCHIVELOG mode.

Last SCN Archived: The last system change number that was archived. The SCN uniquely identifies the last committed database transaction.

Archive Destination: Specifies the destination where the archive log files are to be created. If you are archiving to disk, it is recommended that a dedicated disk with sufficient storage be used.

Archive Format: Specified the naming convention for the archived log files. ARC, appended with the log sequence number (%S) and the thread number (%T), is used to form a unique filename for the archive redo log.

Automatic Archival: Allows you to specify that redo log files be automatically archived. To enable Automatic archiving upon instance startup, set the LOG_ARCHIVE_START initialization parameter to TRUE. This option is enabled when the connected database is running in ARCHIVELOG mode.

SGA Information

Database Buffers: Size of the database buffer cache (in bytes).

Fixed Size: Memory allocated to the area of the SGA that contains general information about the state of the database and the instance. No user data stored here (in bytes).

Redo Buffers: Size of the redo log buffer (in bytes).

Variable Size: Memory allocated to variable size data structures (in bytes).

Database Property Sheet: Resource Limit Page

The page allows you to instantly view resource settings in the Initialization Parameter file versus current usage.

Resource Name

Name of the initialization parameter.

Current Utilization

Current level usage of the initialization parameter.

Max Utilization

Highest recorded usage of the initialization parameter.

If you want to change any of the settings, select the Initialization Parameter container in the Navigator and edit the parameters from the Initialization Parameters property sheet. For a detailed description of the parameters, see your *Oracle Server Reference* manual.

Starting Up a Database

To start up an instance:

1. Select the Database folder. The Database property sheet appears.
2. Go to the Status page and select the desired startup option.
3. Click Apply.

Alternatively, you can access the Startup page of the Database property sheet by choosing Startup from the Database menu. For Oracle Enterprise Manager Release 1.3.6, startup of an Oracle8 database using Instance Manager is not permitted.

Attention: Before starting up a release 7.1 or later database, you must connect as SYSDBA or SYSOPER. For information about connecting as SYSDBA or SYSOPER, see your *Oracle Server Documentation*. For release 7.0, you must be connected as INTERNAL before starting up the database. For information about starting up a database, see your *Oracle Server Administrator's Guide*.

Shutting Down a Database

To shut down a database:

1. Select the Database folder. The Database property sheet appears.
2. Select the Shutdown option.
3. Click Apply.

Alternatively, you can access the Shutdown page of the Database property sheet by choosing Shutdown from the Database menu. For Oracle Enterprise Manager Release 1.3.6, shutdown of an Oracle8 database using Instance Manager is not permitted.

If you are shutting down a remote database, make sure you have a local copy of the INIT<SID>.ORA file or a stored configuration before attempting to restart the database.

Other applications create separate connections when you start them. When performing a shutdown in normal mode, remember to close these database connections, or the shutdown will not complete.

Attention: Before shutting down a release 7.1 or later database, you must connect as SYSDBA or SYSOPER. For information about connecting as SYSDBA or SYSOPER, see the *Oracle7 Server Documentation Addendum*. For release 7.0, you must be connected as INTERNAL before shutting down the database. For information about shutting down a database, see the *Oracle Server Administrator's Guide* for your database release.

Mounting or Opening a Database

If you have previously started an instance without mounting the database, you can mount the database by choosing Mount from the Database menu or selecting the desired option from the Status page of the Database property sheet. The Mount menu command mounts the database in exclusive mode, allowing the database to be mounted by only one instance at a time.

If you have previously mounted a database, you can open the database by choosing Open from the Database menu or selecting the desired option from the Database property sheet: Status page. The database is opened and is accessible to all users.

Initialization Parameters Folder

Initialization parameters specify the operational characteristics of a database. Instance Manager allows you view and edit these parameters. When you select the Initialization Parameters icon in the tree list, the Initialization Parameter property sheets appear.

Initialization Parameters Property Sheet

The Initialization Parameters property sheets consist of two basic and two advanced pages each with a multi-column scrolling list containing information about the parameters defined in the initialization parameter file used to start up your instance.

You can sort the Initialization list on each of the columns by clicking on the column heading. You can edit the values of parameters that can be updated.

Attention: In order to start up or shut down a remote database, the INIT<SID>.ORA file must reside on both the local (Console) and remote (remote database) systems. The INIT<SID>.ORA file on both machines must be identical.

Initialization Parameters Property Sheet: Basic Tuning Page

The Basic Tuning property sheet contains all initialization parameters that are considered essential for basic database operation.

Property sheet columns are defined as follows:

Parameter Name

Name of the initialization parameter.

Value

An editable field displaying the current value of the initialization parameter while the database is running. To modify a parameter, select the desired Value field, change the parameter, then click apply. Instance Manager prompts you if database shutdown is required and if you wish to save the changes as a Stored Configuration.

See the Oracle Server Reference manual for specific information about initialization parameters.

Initialization Parameters Property Sheet: Instance Specific Page

The Instance Specific property sheet contains initialization parameters that pertain to a specific Oracle instance. Normally, these parameters will differ from one database to another.

See *Initialization Parameters Property Sheet: Basic Tuning Page* on page 10-11 for an explanation of page columns.

Initialization Parameters Property Sheet: Advanced Tuning Page (Advanced UI mode)

The Advance Tuning property sheet lists initialization parameters that normally fall beyond the realm of day-to-day database administration. Normally, these parameters are used for performance monitoring or system tuning.

See *Initialization Parameters Property Sheet: Basic Tuning Page* on page 10-11 for an explanation of page columns.

See your Oracle Server Tuning Guide for database tuning information.

Initialization Parameters Property Sheet: Derived Page

The Derived property sheet, which appears when Instance Manager is run in Advanced UI mode, lists initialization parameters that normally fall beyond the realm of day-to-day database administration. Normally, these parameters are used for performance monitoring, or system tuning.

See *Initialization Parameters Property Sheet: Basic Tuning Page* on page 10-11 for an explanation of page columns.

See your Oracle Server documentation for database tuning information.

Editing Initialization Parameters

To edit any initialization parameter:

1. Select a specific parameter from any initialization parameter list.
2. Enter a new value. Some fields in the list consist of drop-down menus.
3. Click Apply.
4. You can also click Save to save your changes as a Stored Configuration.

Applying non-dynamic initialization parameter changes requires that the database first be shut down. Regardless of whether you choose to shut down the database, Instance Manager asks if you wish to save the current set of parameters as a configuration file. See *Stored Configurations* on page 10-14 for more information.

Note: While modifying initialization parameters, nothing is written to the initialization parameter file INIT.ORA. Export the changes to a file by selecting Export to File from the Configuration menu.

Resetting Edits

You can cancel any changes you have made to the parameter values with the Reset command button. Reset only cancels changes you made since the last Apply.

Applying Edits

You can save edits you make to parameter values by using the Apply command button. Any changes you make appear in the New Value column of the initialization parameter list except where parameters are dynamic. Whenever you apply an edit, the following actions occur:

- If you have made changes to parameters that can be dynamically updated, the changes take effect when you click the Apply button.
- If you have made changes to parameters that cannot be dynamically updated, the changes take effect after you have saved the changes and restarted the database.

Stored Configurations

Stored configurations allow you to create multiple database startup configurations without the need to track files initialization parameter files (INIT<SID>.ORA). Stored configurations exist in the registry and not as external files.

Stored Configuration Folder

Clicking on the Stored Configurations folder in the tree list displays the stored configuration multi-column list. This list consists of the following columns:

Configuration Name

User-specified name for the stored configuration.

Description

User-specified description of the new stored configuration.

Creating a Stored Configuration

To create a stored configuration:

1. Click Save from one of the initialization parameter property sheets (Basic Tuning, Instance Specific, Advanced Tuning, or Derived). The Save Configuration property sheet appears.
2. Enter a configuration name.
3. Enter any comments you wish to appear in the Comments column of the Stored Parameters multi-column list.

Stored Configuration Property Sheet

Once a stored configuration is created, it appears in the tree list within the Stored Configurations folder. Selecting an individual stored configuration displays an initialization parameter property sheet containing a single page listing all parameters of the stored configuration. See *Initialization Parameters Property Sheet* on page 10-11. for explanatory information.

Editing a Stored Configuration

To edit a stored configuration:

1. Select an existing stored configuration from the Stored Configuration folder in the tree list. A single property sheet containing a list of initialization parameters appears. The parameters for stored configurations, are not categorized into Basic Tuning, Instance Specific, Advance Tuning, and Derived.
2. Select the desired initialization parameter.
3. Enter a new value.
4. After making the desired change(s):
 - To apply the edit, click Apply. Instance Manager prompts you if database shutdown is required.
 - To save the edit to a configuration file, click Save. You can either overwrite the selected stored parameter or specify that the changes be saved to a new stored configuration.

Deleting Parameters from a Stored Configuration

To delete an initialization parameter from a stored configuration:

1. Select the desire parameter from the initialization parameter property sheet.
2. Choose Delete from the context-sensitive menu.

Adding a Parameter to a Stored Configuration

To add an initialization parameter to a stored configuration:

1. Select the desired stored configuration from the tree list.
2. Enter the name of the new initialization parameter in one of the blank entry fields.
3. Enter the startup value for the initialization parameter.
4. Click OK.

Note: For a full list of initialization parameters and permissible values, see your *Oracle Server Reference Manual*.

Creating a File from a Stored Configuration

Stored configurations are stored in the Windows NT registry. Hence, configurations created on a specific Windows server are only available on that machine. If you wish to make a configuration available to other machines, Instance Manager allows you to export stored configurations to an ASCII file.

To create a stored configuration file:

1. Select a stored configuration from the tree list.
2. Choose Export to File from the Configuration menu. You can also call up this option from the context-sensitive menu from the stored configuration in the tree list. A standard file Save As dialog box appears.
3. Specify the desired file information and click OK.

Deleting a Stored Configuration

To delete a stored configuration:

1. Select the desired configuration from the tree list.
2. Choose Remove from the Configuration menu. Instance Manager asks if you want to remove the stored configuration.
3. Click Yes.

Sessions Folder

The Session list contains information about the users connected to the database. You can sort the Sessions list on each of the columns by clicking on the column heading.

Sessions Multi-column List

When you click on the Sessions folder in the tree list, the Sessions multi-column list displays. The list consists of the following columns:

Session ID

Session identifier

Status

Whether a session is ACTIVE or INACTIVE

Username

Oracle USERID associated with the session.

Schema Name

Schema name associated with the user.

OS User

Operating system user name.

Terminal

Operating system terminal name.

Machine Name

Operating system machine through which the user is connected.

Program

Executable running through the session.

Sessions Property Sheet

The Sessions property sheet consists of a single General page. Columns of the Session list are described below:

Session ID

Session identifier.

Serial #

Session serial number, used to uniquely identify a session. In combination with the SID, guarantees that session-level commands are applied to the correct session in the event that the session ends and another session begins with the same session identifier.

Status

Whether a user session is active or inactive.

Username

Oracle username associated with the session.

Schema Name

Schema name associated with the user.

Program

Program you are currently running.

OS Information

User: Operating system username.

Terminal: Operating system terminal name.

Machine: Operating system machine through which the user is connected.

Disconnecting a User's Session

To disconnect a user's session:

1. Select the session to disconnect from the Sessions multi-column list.
2. Choose Disconnect Immediate or Post-Transaction from the Session menu.

Attention: When you disconnect a session, the session is not actually terminated until the user tries to execute a database operation.

Restricting or Allowing Sessions

To make the database accessible only to users with the `RESTRICTED SESSION` system privilege, choose `Restrict` from the `Session` menu. Only users with the `RESTRICTED SESSION` system privilege are allowed to connect. Users already connected are not affected.

To make the database accessible to all users, choose `Allow All` from the `Database` menu. All users with the `CREATE SESSION` system privilege are allowed to connect.

See [Chapter 9, "Managing Database Security"](#) for more information regarding roles and privileges.

In-Doubt Transactions Folder

The In-Doubt Transactions folder contains information about distributed transactions that failed in the `PREPARED` state. You can sort the Transactions list on each of the columns by clicking on the column heading.

Additional Information: For information about distributed transactions, see the *Oracle Server Concepts*.

In-Doubt Transactions Multi-column List

The columns of the In-Doubt Transactions list are described below:

Global ID

Global identifier for the transaction.

Local ID

Identifier on the local database for the transaction.

State

The state of the transaction: `collecting`, `prepared`, `committed`, `heuristic commit`, or `heuristic abort`.

Advice

Suggested action: `C` (`Commit`), `R` (`Rollback`), or `null` (`no advice`).

Commit Comment

Comment given with the `COMMENT` clause of the `COMMIT WORK` command.

In-Doubt Transactions Property Sheet

The In-Doubt Transactions property sheet displays information about distributed transactions in which a commit was interrupted by a system, network, or any failure resulting from external factors.

Local ID

Name of the node that references data on other nodes to complete its part in the distributed transaction.

Global ID

Name of the node where the distributed transaction originates. The database application issuing the transaction is directly connected to this node.

State

State of the in-doubt transaction: collecting, prepared, committed, forced commit, forced rollback.

Advice

Suggested resolution: C for commit, R for rollback, null for states not requiring immediate action.

Date/Time Information

Failure Time: Date and time of transaction failure.

Last Automatic Retry: Last attempt by the RECO (recover) background process to resolve the transaction discrepancy.

Last Manual Force: Date and time of the last forced rollback or commit.

Commit Comment: Optional text entry field. Contents of this field are displayed

Forcing a Commit or Rollback

To commit an in-doubt transaction:

1. Select the In-Doubt Transactions folder in the tree list. This displays the In-Doubt Transactions multi-column list.
2. Select the transaction to be committed from the In-Doubt Transactions multi-column list.
3. Choose Force Commit from the Transaction menu.

To roll back an in-doubt transaction:

1. Select the In-Doubt Transactions folder in the tree list to display the In-Doubt Transactions multi-column list.
2. Select the transaction to be rolled back from the multi-column list
3. Choose Force Rollback from the Transaction menu.

Attention: You cannot roll back an in-doubt transaction to a savepoint.

Managing Schema Objects

This chapter describes how to use Schema Manager to administer schema objects in your database. This chapter assumes that you have read [Chapter 7, "Overview of the Database Tools"](#) and are familiar with the interface elements of the database tools.

With Schema Manager, you can create, alter, or drop schema objects such as clusters, database links, execution plans, PL/SQL, sequences, snapshots, snapshot logs, synonyms, tables, and views.

The topics in this chapter are:

- [Starting Schema Manager](#)
- [Schema Menus](#)
- [Schema Folders](#)
- [Schema Property Sheets](#)
- [Array Types Folder](#)
- [Clusters Folder](#)
- [Database Links Folder](#)
- [Functions Folder](#)
- [Indexes Folder](#)
- [Object Types Folder](#)
- [Package Bodies Folder](#)
- [Packages Folder](#)
- [Procedures Folder](#)

-
- Queue Tables
 - Refresh Groups Folder
 - Sequences Folder
 - Snapshot Logs Folder
 - Snapshots Folder
 - Synonyms Folder
 - Table Types Folder
 - Tables Folder
 - Triggers Folder
 - Views Folder

Starting Schema Manager

To start Schema Manager, click on the Schema Manager icon in the Administrator Console or choose Schema Manager from the Tools menu.

- If you have already connected to a database, the Schema window appears.
- If a valid database connection has not been made, the Connect dialog box appears. See *Connecting to an Instance* on page 1-23.

Note: You can change the database connection with the Change Database Connection option in the File menu. For information, see *Application Menus* on page 7-10.

After Schema Manager successfully connects to a database, the schema navigator displays on the left side of Schema Manager window. The navigator consists of a series of schema object containers, one for each type schema object contained in the main Schema Objects folder. The name of the connected database is displayed next to the Schema Objects folder.

The display on the right side of the window is determined by the objects selected on the left side of the screen. The right side may contain a multi-column scrolling list, property sheet.

For general information, see:

- *Application Objects* on page 7-3
- *Multi-Column Lists* on page 7-14
- *Dialog Boxes* on page 7-15
- *Property Sheets* on page 7-17

An example of a Schema Manager window is shown in [Table 11-1, "Schema Manager Window"](#).

Figure 11–1 *Schema Manager Window*



Schema Menus

Schema Manager includes four standard menus, File, View, Log, and Help, plus the Object menu. The Object menu contains the following menu options:

Create

Creates a new schema object.

Create Like

Creates a new schema object based on the currently selected object in the tree list.

Remove

Removes the selected schema object in the tree list.

Show Dependencies

Displays database objects that rely on a selected schema object and any objects that the selected schema object relies on.

Grant Privilege on

Grant object privileges to a user or role.

Create Synonym for

Creates a new synonym for a schema object.

Create Index on

Allows you to quickly create indexes on tables and clusters selected from the tree list.

Truncate

Removes a partition from a table.

Drop Storage: Deallocates space previously used by the partition and make that space available for use by other schema objects.

Reuse Storage: Specifies that space previously used by the partition remains allocated to the partition. This space is only available for inserts and updates to the same partition.

Local Indexes

Make Unusable: Marks the local index partitions corresponding to a table partition unusable.

Rebuild Unusable: Rebuilds all unusable local index partitions corresponding to a table partition. This operation is not parallelized.

Exchange

Converts a partition into a non-partitioned table by exchanging the data segments.

Split

Split an existing partition into multiple partitions (Oracle8)

Enqueue (context-sensitive: queue selected in navigator)

Start the selected queue for enqueue and dequeue.

Dequeue (context-sensitive: queue selected in navigator)

Stop the selected queue for enqueue and dequeue.

Schema Folders

The Schema folder contains all database objects pertinent to database schema. Specific schema object types are listed alphabetically under the Schema Object folder.

Alternatively, you can view the navigator by specific schemas defined within the connected database rather than by schema type. To view the navigator by defined schemas, select **By Schema** from the View menu. All schemas defined in the database are listed alphabetically in the navigator as containers. You expand the specific schema folder in the navigator to view the complete list of schema object types defined for that schema.

Schema Property Sheets

There are two property pages that are used in several schema object property sheets. They are available when **Advanced UI Mode** is selected from the View menu. The common property sheets are as follows:

- Storage
- Options

Schema Object Property Sheet: Storage Page

The Storage page allows you to specify database storage parameters for a particular schema object. Schema Manager provides two methods for defining storage parameters:

- **Explicit (default setting):** User specifies all storage parameters for the index.
- **Auto Calculation:** Schema Manager calculates optimum storage parameters using methods recommended in the Oracle Administrator's Guide. In order to use Auto Calculation, you must have select privilege on v\$type_size.

The Storage page layout and options depend on the method chosen.

Explicit

Extents

Initial Size: Size of the object's first extent. You can enter a value, the minimum being the size of one data block. The default value is the size of five data blocks.

Next Size: Size of the next extent to be allocated to the object. The default value is the size of five datablocks. The smallest permissible value is the size of one datablock.

Increase Size by: Percent by which each extent grows (after the second extent) compared to the previous extent.

Minimum Number: Total number of extents allocated when a segment is created. The default value is one. You can enter a value of one or greater.

Maximum Number: Total number of extents, including the first, that an Oracle database can allocate for the object. There are two extent options available, Unlimited and Value.

- **Unlimited:** When selected, allows you to create a number of extents that is only limited by the amount of contiguous free space in a tablespace.
- **Value:** When selected, allows you to specify the maximum number of extents. If you do not specify a value, the default value applies. The default and maximum values depend on the data block size. You can enter a value of one or greater. The default value varies.

Space Usage

% Free: Percentage of space in each of the object's data blocks that is reserved for future updates to the object. You can enter values from 0 to 99.

% Used: Minimum percentage of used space that an Oracle database maintains for each data block of the object. A block becomes a candidate for row insertions when its used space falls below the % Used value. You can enter a value from 1 to 99. The default value is 40.

Number of Transactions

Initial: Initial number of transaction entries allocated within each data block allocated to the object. You can enter a value from 1 or 2 (for clusters and indexes) to 255.

Maximum: Maximum number of concurrent transactions that can update a data block allocated to the object. You can enter a value from 1 to 255.

Freelists

Lists: Number of free lists for each of the free list groups for the table, cluster, or index. You can enter a value of one or greater. The default value is one.

Groups: Number of groups of free lists for a table, cluster, or index. You can enter a value of one or greater. The default value is one.

Auto Calculation

Note: In order to use Auto Calculation, you must have select privilege on v\$type_size.

Size characteristics of table

Initial # rows: number of rows in the indexed table. Schema Manager uses the number of rows to determine the size of the first extent.

Growth rate: expected table growth rate in rows/unit of time. Enter the number of rows, then select day, week, month, or year from the drop-down list. Schema manager uses the growth rate to determine the size (in bytes) of the next extent.

Insert Activity

Determines the percentage of space to leave free for updates and insertions within each of the index's data blocks.

High: sets the percentage of reserved space for index updates to 40 percent.

In Ascending Order: available when Insert Activity is High, specifies that the index be created in ascending order. When selected, the percentage of space reserved for updates and insertions is set to 5 percent.

Low or None: sets the percentage of reserved space for index updates to 40 percent.

Schema Object Property Sheet: Options Page**Parallel**

When checked, specifies parallel execution of an operation.

Degree: The degree of parallelism for an operation on a single instance.

- **Default:** Default number of query servers used. Calculated from the number of CPUs and the number of DEVICES storing tables to be scanned in parallel.
- **Value:** User-specified number of query servers.

Instances: Maximum number of instances allowed to participate in parallel query processing.

- **Default:** Use all available instances. This option only applies to those instances using the Oracle Parallel Server.
- **Value:** User-specified number of instances.

Array Types Folder

An array is an ordered set of data elements, of all the same datatype, with size corresponding to the number of elements in the array. Arrays can be of variable size (VARRAY). When defining an array type, you specify the maximum size. Creating an array type does not allocate space; it defines a datatype which you can use as

- The datatype of a column of a relational table
- An object type attribute
- A PL/SQL variable, parameter, or function return type.

For detailed information on using VARRAYs, see *Oracle8 Server Application Developer's Guide*.

Array Types Multi-column List

The Array Types multi-column list displays when the Array Types folder is selected from the navigator. The multi-column list contains the following columns:

Schema (not displayed if specific schemas are selected in the Navigator)
Schema to which the array type belongs.

Array Type Name
Name of the array type.

Limit
Maximum number of elements allowed in a variable array of the type currently being defined/edited.

Element Type Schema
Schema to which the datatype (for individual elements within the array) belongs.

Element Type Name
Name of the datatype (for individual elements within the array).

Creating an Array Type

To create an array type:

1. Choose Create from the Object menu. The Create Object dialog box appears.
2. Select Array Type from the Create Object list.
3. Click OK. The Array Type property sheet appears.
4. Enter the requisite information.
5. Click Create.

Array Types Property Sheet: General Page

This single-page property sheet allows you to define the requisite attributes for an array type.

Name

Name of the variable array type.

Schema

Schema to which the variable array type belongs.

Size

Maximum number of elements allowed in a variable array of the type currently being defined/edited.

Datatype: Specifies the datatype for elements in the variable array.

Schema: Schema to which the datatype belongs. When a datatype is defined as one of the built-in types (for example, VARCHAR2), schema does not apply. <NONE> is displayed in this case.

Type: Datatype of the element(s) in the array.

Length: Number of bytes allowed for values defined in the column (for CHAR, VARCHAR2, NUMBER, RAW datatypes).

Precision (for NUMBER datatype): Number of digits to the right of the decimal point. If CHAR is selected in the Datatype drop-down list, you may specify length. If VARCHAR2 or RAW is selected, you may specify a length. If NUMBER is selected, you may specify a length, a length and a precision, or neither.

User defined object will be a reference: Specifies that the elements contained in the variable array be a REF. This option is only available if a user-defined datatype is selected from the Type drop-down list.

Create Like Array Type

To create an array type with parameters that are identical to an existing type:

1. Select an existing array type from the Navigator.
2. Choose Create Like from the Object menu. The Array Type property sheet appears with the appropriate parameters filled in.
3. Specify a new array type name and modify, if necessary, any property sheet parameters.
4. Click Create.

Editing an Array Type

To edit an existing array type:

1. Select an array type from the Navigator.
2. Modify the size, element type, or element schema as necessary.
3. Click Apply.

Clusters Folder

The Clusters folder consolidates information about clustered tables.

A cluster is a group of tables that share the same data blocks because they share common columns. Clustered tables are often joined in queries.

Clusters Multi-column List

The Clusters multi-column list displays when the Clusters folder is selected from the navigator. The Clusters multi-column list contains the following columns:

Schema

Name of the schema to which the cluster belongs. This column is not shown when viewing a particular schema.

Cluster

Name of the cluster.

Type

Type of cluster (index or hash).

Tablespace

Tablespace in which the cluster is stored.

Creating a Cluster

To create a new cluster:

1. Choose Create from the Object menu. The Create Object dialog box appears.
2. Select Cluster from the Create Object list.
3. Click OK. The Create Cluster property sheet appears.

The Cluster property sheet contains the following pages:

- General
- Storage (Advanced Mode only, see *Schema Object Property Sheet: Storage Page* on page 11-7)
- Options (Advanced Mode only, see *Schema Object Property Sheet: Options Page* on page 11-9)

Cluster Property Sheet: General Page

This page allows you to define the cluster name, type, and cluster key column attributes. The General page contains the following:

Name

Name of the cluster being defined.

Schema

Name of the schema containing the cluster. The drop-down list contains all schemas available in the current database.

Tablespace

Name of the tablespace to which the cluster will belong.

Size

An estimate of the total amount of space required to store the average number of rows associated with each cluster key value or hash value (kilobytes or megabytes)

Type: Index Cluster

When selected, indicates that an index is specifically defined for the cluster.

Type: Hash Cluster

When selected, indicates that related rows are stored in the same data block.

Hashkeys: Maximum number of unique hash values that can be generated by the hash function.

Hash Function: Specify that Oracle's internal hash function use the current function (Default) or bypass the Oracle internal hash function and specify a SQL expression.

Cluster Key Columns

Column: Name of the column

Datatype: CHAR, VARCHAR2, NUMBER, LONG, RAW, LONG RAW, DATE, ROWID, MLSLABEL

The following buttons appear when creating a new cluster.

New: Displays the Column Details dialog box.

Edit: Displays the Column Details dialog box for the cluster key column selected in the list.

Remove: Deletes a cluster key column that is selected in the list.

Column Details Dialog Box

Note: This dialog box displays when you click on the New or Edit button.

Name

Name of the table column being defined or edited. You can specify a valid Oracle identifier as the name of the column. This field is disabled for columns that have already been defined in the database.

Datatype

Column's Oracle datatype. See *Oracle Server SQL Reference* for specific information on Oracle datatypes.

You can select one of the following items: CHAR, VARCHAR2, NUMBER, RAW, DATE, ROWID, MLSLABEL

Length: Number of bytes allowed for values defined in the column (for CHAR, VARCHAR2, NUMBER, RAW).

Precision (for NUMBER datatype): Number of digits to the right of the decimal point. If CHAR is selected in the Datatype drop-down list, you may specify length. If VARCHAR2 or RAW is selected, you may specify a length. If NUMBER is selected, you may specify a length, a length and a precision, or neither.

Length

Maximum size for values of a given datatype. For specific information on datatypes, see *Oracle Server SQL Reference*.

Precision (Visible when Datatype is Number)

For Number datatypes, value specifies the precision (total number of digits). Precision values can range from 1 to 38.

Create Like Cluster

To create a cluster with parameters that are identical to an existing cluster:

1. Select a cluster from the clusters multi-column list.
2. Choose Create Like from the Object menu. The Create Cluster property sheet appears with the appropriate parameters filled in.
3. Specify a new cluster name and modify (if necessary) any property sheet parameters.
4. Click Create.

Editing a Cluster

To edit an existing cluster:

1. Select a specific cluster in the navigator. The Clusters property sheet appears with the current parameters.
2. Modify the parameter entries as necessary.
3. Click Apply.

Database Links Folder

The Database Links folder contains database links defined within the connected database.

A database link allows you to access a remote database. For information about database links, see *Oracle Server SQL Reference*.

Database Links Multi-column List

The Database Links multi-column list displays when the Database Links folder is selected in the navigator. The list contains information about the status of links defined within the connected database. For information about database links, see *Oracle Server Concepts*.

The columns of the Database Links list are described below:

Name

The name of the schema containing the link. This column is not shown when viewing a particular schema.

Database Link

The name of the database link.

Username

The name of the user the database connects as when connecting to the remote database.

Host

SQLNet connect string to the host.

Created

Date on which the link was created.

If you view the navigator by schema, a Database Links folder appears in an alphabetical list of all schema objects under each schema in the navigator. When the Database Links folder is selected, the multi-column list is displayed.

Creating a Database Link

To create a new database link:

1. Choose Create from the Object menu.
2. Select Database link from the list in the Create Object dialog box. This object will be selected by default if the Database Links folder is selected in the navigator.
3. Click OK. The Create Database Link property sheet appears.
4. Define the parameters of the new database link.
5. Click Create.

The Database Link property sheet consists of a single General page:

Database Link Property Sheet: General Page

The General page contains the following information:

Name

The name of the database link being defined. You can specify any valid Oracle identifier.

Public

Specifies that the database link be available to all users (public) or only the owner (private).

Connection Details

Current User (Oracle 8 only): Need to do determine function.

Connected User: Specifies that the database link use the username and password of the user accessing the database link.

Fixed User

Username: Username used to connect to the remote database. If the Username field is left blank during the creation process, the newly defined link acquires the username of the person accessing the remote database.

Password: Password used to connect to the remote database. If the Username and Password fields are left blank, the database link uses the username password of the person accessing the database link.

Service Name: Database specification of a remote database.

Test

Checks the validity of the Username, Password, and Service Name when a database link is created. This button is only available when altering an existing database link.

Creation Date

Displays the date of database link creation. By default, the current date is displayed.

Create Like Database Link

To create a database link with parameters that are identical to an existing database link:

1. Select a specific link from the database link multi-column list.
2. Select Create Like from the Object menu. The Create Database Link property sheet appears with same connection details and service name as the original link.
3. Enter a new name and modify any parameters (if necessary).
4. Click Create.

Viewing a Database Link

To look at the parameters for an existing link:

Select a specific database link from the navigator. A property sheet with the link's current parameters is displayed. The property sheet is nearly identical to the Create Database Link property sheet with the following exceptions:

- All fields are disabled.
- The Creation date displays the date the database link was created.

Functions Folder

The Functions folder contains an alphabetical list of functions defined within the database.

A function is a PL/SQL subprogram that executes an operation and returns a value at the completion of the operation. A function can be either built-in or user-named. For information about functions, see *Oracle Server SQL Reference*.

Functions Multi-column List

The Functions multi-column list displays when the Functions folder is selected in the navigator. The list contains information about the status of links defined within the connected database.

The columns of the Functions list are described below:

Schema

Name of the schema containing the function. This column is not shown when viewing a particular schema.

Function

Name of the function.

Created

Date on which the function was created.

Last Modified

Date the function was last modified.

Status

Current status of the function (Valid or Invalid).

Creating a Function

To create a new function:

1. Choose Create from the Schema Manager Object menu.
2. Select Function from the list in the Create Object dialog box. The Create Function property sheet appears.
3. Fill in the requisite parameters and enter the PL/SQL source for the function.
4. Click Create.

The property sheet contains a single General page.

Function Property Sheet: General Page

The General page consists of the following information:

Name

Name of the function being defined.

Schema

Name of the schema containing the function. The drop-down list contains all schemas available in the current database.

Recompile

Recompile the PL/SQL source for the function so that it can be stored and executed from within the database.

Dates

Created: Date the function was originally created.

Last Modified: Date the function was last modified.

Status

Current status of the Function (valid or invalid).

Show Errors: Displays the Error Message dialog box listing the errors encountered while trying to compile the function source. This button is only enabled when the function status is invalid. For detailed error message information pertaining to PL/SQL, see the *Oracle Server Messages Manual*.

Source

Function SQL statements or PL/SQL constructs.

Create Like Function

To create a function with parameters and code that is identical to an existing function:

1. Select a function from the function's multi-column list.
2. Select Create Like from the Object menu. The Create Function property sheet appears with parameters PL/SQL source identical to the original function.
3. Enter a new Name.
4. Click Create.

Altering a Function

To alter an existing function:

1. Select a specific function from the navigator. A property sheet with the function's current parameters and source code displays.
2. Modify the parameters and/or PL/SQL source as necessary.
3. Click Apply after changing the desired parameters to recompile the function and save the changes.

Note: To change the parameters and any code modifications back to their original setting, click Revert.

Indexes Folder

The Indexes folder contains table indexes defined within the database. An index is an optional structure associated with a table, used by Oracle to locate rows of that table quickly and can be used to guarantee that every row is unique. For information about indexes, see the *Oracle Server SQL Reference*, or the *PL/SQL User's Guide and Reference*.

Indexes Multi-column List

The Indexes multi-column list appears when you select the Indexes folder in the navigator. This multi-column list contains the following information:

Schema

Name of the schema. This column is not shown when viewing a particular schema.

Index

Name of the index.

Table Owner

Owner of the table for which the index is created.

Table

Table for which the index is created.

Table Type

Type of table (table or cluster).

Status

Status of the table (Valid or Invalid).

When you view the navigator by schema, an Indexes folder appears listing all available schema in alphabetical order. Selecting the Indexes folder displays a multi-column list of indexes contained in the parent schema.

Creating an Index

To create a new index:

1. Choose Create from the Object menu
2. Select Indexes from the list in the Create Object dialog box. The Create Index property sheets appear.
3. Enter the requisite parameter values.
4. Click Create.

The Create Index property sheet consists of up to three pages:

- General
- Storage (Advanced Mode only, see *Schema Object Property Sheet: Storage Page* on page 11-7)
- Options (Advanced Mode only, see *Schema Object Property Sheet: Options Page* on page 11-9)

Note: You can also call up the Index property sheet by selecting a table from the navigator and choosing Create Index On from the Object menu.

Index Property Sheet: General Page

The General page of the index property sheet allows you to define an index on one or more columns of a table or a cluster. The General page consists of the following information:

Name

Name of the index. When creating an index you can specify a valid Oracle identifier

Schema

Schema in which the index will be defined. The default schema for a new index is the user's default schema. When creating an index link, you can also select any schema defined in the connected database from the Schema drop-down list.

Tablespace

Name of the tablespace to which the index belongs.

Index On

Table/Cluster: Specifies whether the index is to be placed on a table or cluster.

Schema: Drop-down list displaying available schema to which the index can belong.

Table: Drop-down list displaying all available tables belonging to the selected schema.

Index Type (Oracle 8 only): Specifies whether a partitioned index is local (equi-partitioned with the underlying table) or global (partitioning defined explicitly when you create the index). Partitioning of global indexes must be maintained manually.

Table Spreadsheet

- *Table Columns*: Column names of the table selected in the Table drop-down list.
- *Order*: Order of the columns selected for the index. Columns are ordered in the sequence in which it is added to the index.

To add a column to the index, click on the desired column in the spreadsheet. The order of the column appears in the Order column.

To remove a column from the index click on the desired column entry in the spreadsheet. The column are automatically reordered.

Options

Unique: Specifies that the value of the column (or combination of columns) in the table to be indexed must be unique.

Bitmap: Specifies that the index is to be created as a bitmap rather than as a B-tree. This option does not apply when creating a global partitioned index.

Sorted (Default): If deselected, indicates to Oracle that rows are stored in the database are in ascending order and therefore do not have to be sorted when creating the index.

Recoverable: Specifies that the creation of the index will be logged in the redo log file.

Create Like Index

To create an index with parameters that are identical to an existing index:

1. Select a index from the multi-column list.
2. Select Create Like from the Object menu. The Create Index property sheet appears with the appropriate parameters filled in.
3. Enter a new name and modify the parameter settings (if necessary).
4. Click Create to close the property sheet and create the new index.

Altering an Index

To alter an existing index:

1. Select a specific index from the navigator. A property sheet with the index's current parameters displays.
2. Modify the property sheet parameters as necessary.
3. Click Apply.

Object Types Folder

The Object Types folder consolidates information about user-defined datatypes (object types) within the connected database. Object types are schema objects that can represent real-world entities, such as a purchase order used by an inventory control application.

Object Types Multi-column list

The Object Type multi-column list displays when the Object Types folder is selected in the Navigator. The list contains the following columns:

Schema

Schema to which the object type belongs. Column is not displayed if an individual schema is selected in the Navigator.

Object Type Name

Name of the object type.

Incomplete

Whether or not the object type is incomplete (NO-object type has at least one attribute or method) or complete (YES-object type has no attributes or methods).

Number of Attributes

Total number of attributes that are defined for the object type.

Number of Methods

Total number of methods defined for the object type.

Creating an Object Type

To create a new object type:

1. Choose Create from the Object menu. The Create Object dialog box appears.
2. Select Object Type from the Create Object list.
3. Click OK. The Object Types property sheet appears.
4. Click Create after defining the requisite parameters in the property sheet.

Object Types Property Sheet: General Page

The General page allows you to specify the name of the object type and the schema within which to create it. Object types are abstractions of real-world entities, such as a purchase order. You have the option of defining either an incomplete (no attributes or methods specified) or a complete object type.

Name

Name of the user-defined object type.

Schema

Schema in which the object is to reside.

Show Errors/Hide Errors: Opens or closes the Error Messages dialog box. This button is available when the object status is invalid.

Attributes

Multi-column list displaying all currently defined attributes.

Add: Displays the Object Type Attribute property sheet allowing you to define a new object attribute. Newly created attributes appear in the multi-column list.

Edit: Displays the Object Type Attribute property sheet allowing you to modify the attribute selected in the multi-column list.

Remove: Deletes the attribute selected in the Attributes multi-column list.

Methods

Multi-column list displaying all available object methods.

Add: Displays the Object Type Method property sheet allowing you to define a new method for the object. Newly created methods appear in the Methods multi-column list.;

Edit: Displays the Object Type Method property sheet allowing you to modify a method selected in the multi-column list.

Remove: Deletes the method selected in the Attributes multi-column list.

Adding Attributes to Object Types

To add attributes to an object type:

1. Select an existing object type from the Navigator, or create a new object type.
2. From the Object Type property sheet, click Add in the Attributes area. the Object Attribute property sheet displays.
3. Enter the requisite information and click OK. The new attribute appears in the Attribute multi-column list.
4. Click Apply or Create.

Object Attribute Property Sheet: General Page

This page allows you to define the attributes of a user-defined object. Attributes are built-in types or other user-defined types which model the structure of real world entities. Attributes are the variables that form the data structure of a user-defined object.

Attribute Name

Name of the object attribute.

Object Schema

Schema containing the object to which the attribute belongs.

Object Name

Name of the object to which the attribute belongs.

Datatype

Schema: Attributes's datatype schema. When an attribute is defined as one of the built-in types (for example, VARCHAR2), schema does not apply. Hence, <NONE> is displayed.

Type: Attribute datatype. The drop-down list displays all standard Oracle built-in types. If a schema containing user-defined types is selected, these types also appear in the list.

Length: Number of bytes allowed for values defined in the column (for CHAR, VARCHAR2, NUMBER, RAW).

Precision (for NUMBER datatype): Number of digits to the right of the decimal point. If CHAR is selected in the Datatype drop-down list, you may specify length. If VARCHAR2 or RAW is selected, you may specify a length. If NUMBER is selected, you may specify a length, a length and a precision, or neither.

User defined object will be a reference (REF): Specifies that the attribute be a REF datatype. A REF datatype allows the attribute to "refer" or point to an object stored elsewhere. This option is only available if a user-defined datatype is selected from the Type drop-down list.

Editing or Removing Object Attributes

You can edit or remove an attribute from an object type by selecting the desired attribute in the method multi-column list and clicking on Edit (to bring up the Object Method property sheet) or Remove to delete the selected method.

Adding Methods to Object Types

To add methods to an object type:

1. Select an existing object type from the Navigator, or create a new object type.
2. From the Object Type property sheet, click Add in the Methods area. the Object Method property sheet displays. Enter the requisite parameters and click OK. The new method appears in the Method multi-column list.
3. Click Apply (edit mode) or Create (create mode).

The Object Method property sheet consists of the three pages:

- General
- Body
- Pragma

Object Method Property Sheet: General Page

The object method property sheet allows you define functions or procedures written in PL/SQL and stored in the database.

Method Name

Name of the method being created.

Object Schema

Schema of the object to which the method belongs.

Object Name

Name of the object to which the method belongs.

Method Type

Drop-down list allows you to select one of four possible types: Procedure, Function, Map Function, or Order Function. Method type selection determines which fields are active in the Return Type area.

Parameters: An editable spreadsheet allowing you define method input parameters (one parameter per row).

Name: Name of the parameter.

Schema: Schema to which the parameter belongs. If a built-in datatype is selected, <NONE> is displayed. Built-in types are common to all schemas.

Type: Datatype of the parameter. Drop-down list displays all built-in datatypes in addition to any user-defined types belonging to the selected schema.

Ref: Indicates whether the parameter is a reference. The field is enabled if the selected datatype can be referenced.

Direction: Specifies the behavior of method. There are three modes: IN (default-allows you to pass values to the subprogram being called), OUT (allows you return values to the caller of a subprogram), and IN OUT (allows you to pass initial values to the subprogram being called and return updated values to the caller).

Note: New additions (defined but not yet created) are graphically identified by a hand with a plus sign. After selecting an entry in the spreadsheet, you can use the context-sensitive menu (right mouse button) to drop the selected parameter.

Return Type: Allows you to specify the datatype of the result being returned from the defined method.

Schema: Schema to which the return value type belongs.

Type: Datatype of the return value.

User defined object will be a reference: Specifies that the method return value be a REF datatype. This option is only available if a user-defined datatype is selected from the Type drop-down list.

Object Method Property Sheet: Body Page

This page allows you to define/edit the PL/SQL code which makes up the various methods of the object type body. An empty object type body (which includes its heading and the keyword END;) is automatically provided. Within the edit area, you specify the individual headings and bodies of the various methods that belong in the object type body.

Object Method Property Sheet: Pragma Page

This page allows you to set method restrictions (pragma) denying member functions read/write access to database tables, packaged variables, or both. Changes made to pragma on this page override the settings specified at the object type level. Pragma are processed when the method is compiled. See the *Oracle8 PL/SQL User's Guide and Reference* for more information on user-defined exceptions.

Override Default Method restrictions (Pragma): When selected, overrides object type settings and allows you to choose any or all of the following restrictions for the method currently being defined or edited:

Writes No Database State (WNDS): Does not modify database tables.

Writes No Package State (WNPS): Does not modify packaged variables.

Reads No Database State (RNDS): Does not query database tables.

Reads No Package State (RNPS): Does not reference packaged variables.

Editing or Removing Object Methods

You can edit or remove methods from an object type by selecting the desired method in the method multi-column list and clicking on Edit (to bring up the Object Method property sheet) or Remove to delete the selected method.

Package Bodies Folder

The Package Bodies folder consists of a list of package bodies defined within the connected database. Package bodies are organized by schema within the Package Bodies folder.

A package consists of two parts: the specification and the body. The package body defines all constructs (public and private) of the package while the specification declares all public constructs. For information about package bodies, see *Oracle Server SQL Reference*, or the *PL/SQL User's Guide and Reference*.

Package Bodies Multi-column List

The Package Bodies multi-column list displays when the Package Bodies folder is selected in the navigator. The list contains all package bodies defined within the connected database.

The columns of the package bodies list are described below:

Schema

Name of the schema containing the package body.

Package Body

Name of the package body.

Created

Date on which the package body was created.

Last Modified

Date on which the package body was last modified.

Status

Current status of the package body (Valid or Invalid).

Creating a Package Body

To create a new package body:

1. Choose Create from the Schema Manager Object menu.
2. Select Package Body from the list in the Create Object dialog box. The Create Package Body property sheet appears.
3. Enter the Name, Schema and PL/SQL source for the package body.
4. Click Create to compile the package body and close the property sheet.

The Package Body property sheet contains a single General page.

Package Body Property Sheet: General Page

The General page consists of the following information:

Name

Name of the package body being defined.

Schema

Name of the schema containing this package body. The drop-down list contains all schemas available in the current database.

Recompile

Recompile the PL/SQL source for the package body so that it can be stored and executed from within the database.

Source

PL/SQL source code for the package body.

Dates

Created: Date the package body was originally created.

Last Modified: Date the package body was last modified.

Status

Current status of the package body (valid or invalid).

Show Errors/Hide Errors: Opens or closes the Error Messages dialog box. This button is available when the package body status is invalid.

Create Like Package Body

To create an package body with parameters that are identical to an existing one:

1. Select a package body from the navigator.
2. Select Create Like from the Object menu. The Create Package Body property sheet appears.
3. Enter a Name and modify property sheet parameters and/or PL/SQL source (if necessary).
4. Click Create to compile the package body and close the property sheet.

Editing a Package Body

To edit an existing package body:

1. Select a package body in the navigator. The Package Body property sheet appears with the Name and Schema fields greyed out.
2. Modify the PL/SQL source.
3. Click Apply to recompile the source code and commit the changes.

Packages Folder

The Package folder consists of a list of all packages defined within the connected database. Packages within this folder are organized by schema.

A package is a group of related procedures and functions, together with cursors and variables they use, stored together in the database for continued use as a unit. For information about packages, see *Oracle Server SQL Reference*, or the *PL/SQL User's Guide and Reference*.

Packages Multi-column List

The Packages multi-column list appears when the Packages folder is selected in the navigator. The list contains information about:

Schema

Name of the schema containing the package. This column is not displayed when viewing a particular schema.

Package

Name of the package.

Created

Date on which the package was created.

Last Modified

Date on which the package was last modified.

Status

Current status of the package (valid or invalid).

Creating a Package

To create a new package:

1. Choose Create from the Schema Manager Object menu.
2. Select Package from the list in the Create Object dialog box. The Create Package property sheet appears.
3. Enter the Name, Schema and PL/SQL source for the package. For more information on creating packages, see *Oracle PL/SQL User's Guide and Reference*.
4. Click Create to compile the package source and create the new package.

The Package property sheet consists of a single General page.

Packages Property Sheet: General Page

The General page consists of the following information:

Name

Name of the package being defined.

Schema

Name of the schema containing this package. The drop-down list contains all schemas available in the current database.

Recompile

Recompile the PL/SQL source for the package so that it can be stored and executed from within the database.

Dates

Created: Date the package was originally created.

Last Modified: Date the package was last modified.

Status

Current status of the package (valid or invalid).

Show Errors/Hide Errors: Opens or closes the Error Messages dialog box. This button is available when the package status is invalid.

Source

PL/SQL source code for the package.

Create Like Package

To create an package with parameters that are identical to an existing package:

1. Select a package from the navigator
2. Select Create Like from the Object menu. The Create Package property sheet appears with the appropriate information filled in.
3. Enter a name and modify the PL/SQL source (if necessary).
4. Click Create to compile the source and create the package.

Editing a Package

To edit an existing package's PL/SQL source:

1. Select a specific package in the navigator. A property sheet with the package's current parameters is displayed.
2. Modify the source code.
3. Click Apply to compile the source code and commit the changes.

Procedures Folder

The Procedures folder consists of a list of all procedures defined within the connected database. Procedures in the navigator are organized by schema.

A procedure groups a set of SQL and PL/SQL statements together to perform a specific task. For more information on procedures, see *Oracle Server SQL Reference, Vol. II*, or the *PL/SQL User's Guide and Reference*.

Procedures Multi-column List

The Procedures multi-column list appears when the Procedures folder is selected in the navigator. The list contains the following information:

Schema

Name of the schema containing the procedure. This column is not shown when viewing a particular schema.

Procedure

Name of the procedure.

Created

Date the procedure was originally created.

Last Modified

Date the procedure was last modified.

Status

Current status of the procedure (Valid or Invalid).

Creating a Procedure

To create a new procedure:

1. Choose Create from the Schema Manager Object menu.
2. Select Procedure from the list in the Create Object dialog box. The Create Package property sheet appears.
3. Enter the Name, Schema, and SQL or PL/SQL source code.
4. Click Create to compile the source code and create the procedure.

The property sheet consists of a single General page.

Procedures Property Sheet: General Page

The General page consists of the following information:

Name

Name of the procedure being defined.

Schema

Name of the schema containing this procedure. The drop-down list displays the names of all available schemas.

Recompile

Recompile the PL/SQL source for the procedure so that it can be stored and executed from within the database.

Source

PL/SQL source code for the procedure.

Status

Current status of the procedure (valid or invalid).

Show Errors: Displays the Error Message dialog box listing the errors encountered while trying to compile the function source. This button is only enabled when the function status is invalid.

Dates

Created: Date the procedure was originally created.

Last Modified: Date the procedure was last modified.

Create Like Procedure

To create an procedure with parameters that are identical to an existing one:

1. Select a procedure from the multi-column list.
2. Select Create Like from the Object menu. The Create Procedure property sheet appears with the appropriate parameters filled in.
3. Enter a new name and modify the procedure source code (if necessary).
4. Click Apply to compile the code and create the new procedure

Altering a Procedure

To alter an existing procedure:

1. Select a procedure in the navigator. The Procedure property sheet appears with the current parameters and source code.
2. Modify the procedure source code.
3. Click Apply to recompile the source and commit the changes.

Queue Tables

Advanced Queuing allows you to defer execution of a database request. After a request for work is entered, advanced queuing defers processing of that request until the requestor completes the task or process or transaction that created the request. For more information about Advanced Queueing, see the Advanced Queueing chapter in the *Oracle Server Application Developer's Guide*.

Queues Multi-column List

Selecting Queue Tables from the tree list displays multi-column list containing all queues in the database organized by schema. The columns are defined as follows:

Schema

Name of the schema containing the queue. The column is not displayed if an individual schema is selected.

Queue Name

Valid queue identifier.

Queue Table

Table (data repository containing one or more set of queues) created using the DBMS_AQADM package.

Type

Normal (Reply): An application server can communicate through queues.

Exception: If an application server is unable to process a queue message, it is ENQUEUED to an exception queue to be processed by another application server.

Enqueue

Enabled or Disabled. Ability to add a queue message onto a queue specified by a user.

Dequeue

Enabled or Disabled. Ability to remove a queue message from a queue specified by a user.

User Comment

Textual description of queue function.

Creating a Queue Table

To create a queue table:

1. Choose Create from the Object menu.
2. Select Queue Table from the list in the Create Object dialog box. The Queue Table property sheet appears.
3. Enter the requisite queue table information.
4. Click Create.

The Queue Table property sheet consists of the three pages:

- General
- Storage (See *Schema Object Property Sheet: Storage Page* on page 11-7)
- LOB Storage

Queue Table Property Sheet: General Page

Queue tables are database tables in which queues (repository for messages) are stored. This property sheet allows you to create queue tables.

Name

Name of the queue table to be created.

Schema

Schema in which queue table is to reside.

User Comment

User-specified description of the queue table. This comment is added to the queue catalog.

Tablespace

Tablespace in where the queue table is created.

Payload Type

Specifies the type of user data to be stored in the queue table (queue datatype). Queue table datatypes consist of the standard object types (Oracle-defined or user-defined). In addition the RAW datatype has been added, permitting queue tables to handle large objects (LOBs). Although LOBs can hold up to 4 gigabytes of RAW binary data, the size of the payload is limited to 32 KB.

Raw: Specifies that the queue table store payloads of type RAW. When the RAW datatype is selected the LOB Storage property sheet appears in addition to the General and Storage property sheets.

Object: Specifies that the queue table store conventional object type payloads. Datatypes can be either Oracle-defined or user-defined. The following drop-down lists are active when Object is selected:

Schema: Schema containing the user-defined type.

Type: Payload datatype to be stored in the queue table.

Options

Allow Subscriber: Permits a single message to be consumed by multiple consumers. (one-to-many). When this option is selected, a Subscribers page appears within the Queue property sheet.

Allow Message Group: Permits one or more messages belonging to a single queue to be grouped into set that can only be consumed by one user at a time (many to one). Message grouping allows you to segment complex messages into simple messages.

Sort List

Specifies the columns to be used as the sort key in ascending order. The allowed column names are Priority (priority) and Enqueue Time (enq_time). You click on a column entry to select it. If both columns are selected, then the first column selected (Order 1) defines the most significant order.

Note: If no order is specified, all queues in the queue table are sorted according to the enqueue time in ascending order.

Queue Table Property Sheet: Lob Storage

This page allows you to set parameters for large object (LOB) storage within the queue table. This page displays when RAW is selected as the payload type for a queue table (Queue Table: General page) or if a selected queue table heirarchically contains a user-defined object type with LOB attributes.

All (queue table creation only)

Applies LOB storage characteristics to all data segments within the queue table

Specific (queue table creation only)

Allows you to explicitly indicate the tablespace and storage characteristics for each LOB. The drop-down list displays existing LOBs.

Name

LOB column name.

Data Segment

Tablespace: Drop-down list displaying all tablespaces within current database. Selection determines the tablespace in which the LOB data segment is to reside.

Data Segment Storage: Standard storage parameter page allowing you to define specific data segment storage parameters.

Enable LOB storage in row: Specifies that the LOB value is stored in the row (inline) if its length is less than approximately 4000 bytes minus system control information (default value).

Number of blocks of LOB data that will be accessed at one time (CHUNK): Integer representing the unit of LOB allocation and manipulation by Oracle. This unit of LOB storage is calculated as the specified integer multiplied by the data block size.

Percent of LOB storage space to be used as old versions of LOB pages (PCTVERSION): Integer representing the maximum percentage of overall LOB storage space used for creating new versions of the LOB. The default value is 10.

LOB cache (On/Off): Specifies whether LOB data should be cached. Caching is off by default. You should turn LOB caching On if LOB data will be accessed frequently.

Generate full redo of LOB data pages (LOGGING): When selected, specifies that the LOB storage characteristics will be logged in the redo log file. This option is only available if LOB cache is Off.

Create Like Queue Table

To create another queue table with parameters that are identical to an existing queue table:

1. Select an existing sequence from the Navigator.
2. Select Create Like from the Object menu. The Queue Table property sheet displays.
3. Enter a name and modify the property sheet parameters as necessary.
4. Click Create.

Altering a Queue Table

To alter an existing sequence:

1. Select a specific queue table in the navigator. A property sheet with the table's current parameters is displayed.
2. Modify the property sheet parameters as necessary.
3. Click Apply.

Creating a Queue

To create a queue:

1. Choose Create from the Object menu.
2. Select Queue Table from the list in the Create Object dialog box. The Queue Table property sheet appears.
3. Enter the requisite queue table information.
4. Click Create.

The Queue property sheet consists of two pages:

- General
- Subscribers (only appears if an associated queue table had been created with the Allow Subscriber option)

Queue Property Sheet: General Page

A queue is a repository for messages. This property sheet allows you to define one of two types of queues: A user queue (Normal), or an exception queue.

Name

Name of the queue to be created.

Schema

Schema where the queue table containing the queue is located.

Table

Queue table containing the queue.

User Comment

User-specified description of the queue. The comment will be added to the queue catalog.

Queue Type

Normal: Specifies that the queue be a user queue for normal message processing.

Exception: Specifies that the queue be an exception queue. Exception queues handle messages that for some reason cannot be retrieved and processed.

Enqueue

Enabled: Specifies that ENQUEUE should be enabled on this queue.

Disabled: Disables enqueueing for this queue.

Note: A queue cannot be stopped if there are outstanding transactions against the queue.

Dequeue

Enabled: Specifies that DEQUEUE should be enabled on this queue.

Disabled: Disables dequeuing for this queue.

Note: Only dequeue operations are allowed for exception queues.

Retries

Max Retries: Number of times a dequeue with the REMOVE mode can be attempted on a message. A message is moved to an exception queue when the Max Retries limit is reached. By default, the limit is set to zero, meaning no retry is allowed.

Delay: Specifies the delay time before a message is scheduled for processing again after an application rollback. You use the drop-down list to select the time units (seconds, minutes, hours, or days). By default, the delay time is set to zero, meaning the message can be retried as soon as possible.

Message Retention

Retention Time: The time a message will be retained in the queue table after being dequeued from the queue. By default, the time is set to zero, i.e., not retained.

Queue Property Sheet: Subscribers Page

This page allows you to add subscribers to the queue in addition to scheduling the propagation of messages from a queue to a destination identified by a specific database link. The Subscribers page only appears if an associated queue table had been created with the Allow Subscriber option.

List of Subscribers

An editable spreadsheet allowing you to define the queue subscribers. The spreadsheet consists of three columns:

Name: Consumer of a message.

Note: For Oracle release 8.0.3, you can only define a name for a subscriber.

Address: Protocol-specific address of the message recipient. By default, the address is of the form SCHEMA.QUEUE

DBLINK: Database link associated with the message consumer.

To add a subscriber name, click on the desired cell to activate the cursor and type in the desired name, address, and database link. Entries may consist of up to 30 characters.

To remove an entry from the list, click on the spreadsheet button located to the immediate left of the name. Right click on the spreadsheet button to display the context-sensitive menu. Select Delete.

Note: You can perform standard editing functions (cut, copy, past, delete) by selecting the text in the spreadsheet and right clicking to display the context-sensitive menu.

List of Schedules

An editable spreadsheet allowing you to schedule the propagation of messages to a specific destination.

DB Link: Drop-down list displaying all database links (destinations) defined in the current database. You can unschedule a queue from a destination by clicking on the spreadsheet button to the immediate left of the database link, and then using the right mouse button to display the context-sensitive menu. Select Unschedule.

Note: If no database link is selected, the destination is set to the local database. Messages will be propagated to other queues within the local database.

Schedule (for): Displays the current schedule for the selected database link.

Start Time: Specifies the initial start time for the propagation window for messages from the source queue to the destination.

Next Time: Date function to compute the start of the next propagation window from the end of the current window. If this field is left blank, propagation will be stopped at the end of the current window. For example, to start the window at the same time every day, Next Time should be specified as 'SYSDATE + 1 - duration/86400'.

Duration: Specifies the duration of the propagation window in seconds, minutes, or hours. Leaving the Duration field blank specifies that the propagation window remain open until the propagation is unscheduled.

Latency: Specifies the minimum wait (seconds, minutes, or hours) in the propagation window for a message to be propagated after it is enqueued.

Create Like Queue

To create another queue with parameters that are identical to an existing queue table:

1. Select an existing sequence from the Navigator.
2. Select Create Like from the Object menu. The Queue property sheet displays.
3. Enter a name and modify the property sheet parameters as necessary.
4. Click Create.

Altering a Queue

To alter an existing sequence:

1. Select a specific queue table in the navigator. A property sheet with the table's current parameters is displayed.
2. Modify the property sheet parameters as necessary.
3. Click Apply.

Refresh Groups Folder

The Refresh Group folder contains all refresh groups defined within the connected database. Within the folder, refresh groups are organized by schema.

A refresh group consolidates multiple snapshots and allows you to easily refresh two or more snapshots to a single point in time.

For information about refresh groups, see *Oracle Server Distributed Systems, Volume II: Replicated Data*.

Refresh Groups Multi-column List

The Refresh Groups multi-column list appears when the Refresh Groups folder is selected in the navigator. The list contains the following information:

Schema

Name of the schema containing the refresh group. This column is not shown when viewing a particular schema.

Name

Name of the refresh group. Refresh group names must follow the same naming conventions as tables.

Number

Identification number of the refresh group.

Next Date

Next date the refresh will occur.

Interval

Function used to calculate the next time to refresh the snapshots in the group.

Creating a Refresh Group

To create a new refresh group:

1. Choose Create from the Object menu.
2. Select Refresh Group from the list in the Create Object dialog box. The Create Refresh Group property sheet appears.
3. Enter the requisite information.

To **add** a snapshot to the refresh group, select an item from the Available list and click the Up arrow. The snapshot appears in the In Group list.

To **remove** a snapshot from the refresh group, select an item in the In Group list and click the Down arrow.

4. Click Create.

The Refresh Group property sheet consists of two pages:

- General
- Snapshots

Refresh Group Property Sheet: General Page

The General page consists of the following information:

Name

Name of the Refresh group.

Schema

Name of the schema containing the refresh group.

Refresh

Refreshes the current refresh group immediately.

Broken

When checked, marks a job as broken so Oracle will not attempt to execute the job.

Delete group when last member deleted

Upon deleting the last refresh group member, the refresh group is automatically deleted.

Refresh

Next Date: Next date that you want a refresh to occur.

Interval: Function used to calculate the next time to refresh the snapshots in the refresh group.

Continue refresh despite conflicts: Select this option if you want the refresh to proceed even if there are outstanding conflicts logged in the DefError table for the snapshot's master. This option is active for updatable snapshots only.

Push changes from snapshot to master before refresh: Selected by default, this option allows you to push changes from the snapshot to its associated master before refreshing the snapshot. Otherwise, these changes may appear to be temporarily lost. This option is active for updatable snapshots only.

Rollback Segment

Use default rollback segment: When selected, the default rollback segment is used. Segment name and Segment Schema are greyed out.

Segment Name: Name of the rollback segment to use while refreshing snapshots.

Segment Schema: Name of the schema containing the rollback segment.

Refresh Group Property Sheet: Snapshots

The Snapshot page consists of the following information:

In Group

Scrolling list containing all snapshots in the refresh group.

Available

Scrolling list containing all available snapshots.

Create Like Refresh Group

To create a refresh group with parameters that are identical to an existing one:

1. Select a refresh group in the navigator.
2. Select Create Like from the Object menu. The Create Refresh Group property sheet appears with the appropriate parameters filled in.
3. Enter a name and modify any parameters (if necessary).
4. Click Create.

Altering a Refresh Group

To alter an existing refresh group:

1. Select a specific refresh group in the navigator. A property sheet with the refresh group's current parameters displays.
2. Modify the property sheet entries as necessary.
3. Click Apply.

Sequences Folder

The Sequence folder contains all sequences stored within the connected database. Sequences are organized by schema.

A sequence is a database object used to generate a serial list of unique numbers for numeric columns of a database's tables. Sequences simplify data entry in a multi-user system by automatically generating unique numerical values for the rows of a single table or multiple tables. For information about sequences, see *Oracle Server Concepts* and *Oracle Server Administrator's Guide*.

Sequences Multi-column List

The Sequences multi-column list displays when the Sequences folder is selected in the navigator. The list contains sequences defined within the connected database. For information about sequences, see the *Oracle Server SQL Reference*.

The columns of the Sequences multi-column list are described as follows:

Schema

Name of the schema containing the sequence. This column is not shown when viewing a particular schema.

Sequence

Name of the sequence.

Minimum Value

Smallest permissible value for the sequence.

Maximum Value

Largest permissible value for the sequence.

Interval

Incremental value for the sequence (value may be positive or negative).

Recent Value

Most recent value of the sequence.

Cycle

Indicates whether or not the sequence should continue to generate values after reaching its minimum or maximum value (Yes or No).

Order

Indicates whether or not the sequence numbers should be generated in order of request (Yes or No).

Cache Size

Number of values to be preallocated and stored by the database.

Creating a Sequence

To create a new sequence:

1. Choose Create from the Object menu.
2. Select Sequence from the list in the Create Object dialog box. The Create Sequence property sheet appears.
3. Enter the requisite sequence information.
4. Click Create.

The property sheet consists of a single General page.

Sequences Property Sheet: General Page

The General page contains the following information:

Name

Name of the sequence being defined. During the Create operation, you can specify a valid Oracle identifier as the name of the sequence.

Schema

Schema in which the sequence is defined. The default schema for a new sequence is the user's default schema. During the Create operation, you can choose a schema from those defined in the connected database from the drop-down list.

Type: Ascending

When selected, sequence values increase from the initial value towards the maximum value. This is the default setting when creating a sequence.

Type: Descending

When selected, the sequence values decrease from the initial value towards the minimum value.

Values

Minimum: Minimum allowed value of the sequence. When creating a sequence, this field is initially blank. If this field is blank when you select the Create button, a default value of one for an ascending sequence or -10^{26} for a descending sequence will be used.

Maximum: Maximum value of a sequence. When creating a sequence this field is initially blank. If left blank when the Create button is selected, a default value of 10^{27} for an ascending sequence and -1 for a descending sequence will be used.

Increment: Interval by which the sequence increases or decreases. When creating a sequence, this field is initially blank. If left blank when the Create button is selected, a default value of 1 is used. Only positive integers are accepted in this field.

Decrement: Interval by which the sequence decreases. When creating a sequence, this field is initially blank. If left blank when the Create button is selected, a default value of 1 is used. Only positive integers are accepted in this field.

Initial: The starting value of the sequence. If the field is blank when the Create button is selected, a default value of the sequence's minimum will be used for an ascending sequence. For a descending sequence, a default value of the sequence's maximum value is used.

Recent: The most recent value of the sequence. This field is active when altering an existing sequence.

Options

Cycle Values: When checked, specifies that the sequence should continue to generate values after reaching its minimum or maximum value. For ascending sequences, the minimum value is generated after the maximum is reached. For descending sequences, the maximum value is generated after the minimum is reached. If Cycle Values is not checked, the sequence will stop generating values when it reaches its minimum or maximum. When creating a sequence, this checkbox is unchecked by default.

Order Values: When checked, specifies that the sequence numbers are to be generated in order of request. When creating a sequence, this checkbox is unchecked by default.

Cache Size: Number of values to be pre-allocated and stored by the database. There are three options:

- **Default:** The default value is 20. This option is selected by default.
- **No Cache:** Specifies that values of the sequence not be pre-allocated.
- **Size:** The smallest acceptable value is 2. For sequences that cycle, this value must be less than the number of values in the cycle. If the maximum number of values the sequence can generate is less than the cache size, the cache size will be automatically changed to the maximum number of possible values.

Create Like Sequence

To create another sequence with parameters that are identical to an existing sequence:

1. Select an existing sequence from the navigator.
2. Select Create Like from the Object menu. The Create Sequence property sheet displays.
3. Enter a name and modify the property sheet parameters as necessary.
4. Click Create.

Altering a Sequence

To alter an existing sequence:

1. Select a specific sequence in the navigator. A property sheet with the sequence's current parameters is displayed.
2. Modify the property sheet parameters as necessary.
3. Click Apply.

Snapshot Logs Folder

The Snapshot Logs folder contains a list of all snapshot logs created in the connected database.

A snapshot log is a table associated with the master table of a snapshot. The snapshot log contains information about changes made to its associated master table. The information in the snapshot log is used to refresh the snapshot.

Snapshot Logs Multi-column List

The Snapshot Logs multi-column list displays when the Snapshots folder is selected in the navigator. The multi-column list contains snapshot log information defined within the connected database. For information about snapshot logs, see *Oracle Server Concepts*.

The columns of the Snapshots multi-column list are described as follows:

Schema

Name of the schema containing the snapshot log. This column is not shown when viewing a particular schema.

Log Table

Name of the table storing the ROWID and timestamp of rows updated in the master table.

Log Trigger

Name of the AFTER ROW trigger on the master table.

Master Table

Name of the table associated with the snapshot log.

Creating a Snapshot Log

To create a new snapshot log:

1. Choose Create from the Object menu.
2. Select Snapshot Log from the list in the Create Object dialog box. The Create Snapshot Log property sheet appears.
3. Enter the necessary property sheet information.
4. Click Create.

The Snapshot Logs property sheet consists of two pages:

- General
- Storage (Advanced Mode only, see *Schema Object Property Sheet: Storage Page* on page 11-7)
- Options (Advanced Mode only, see *Schema Object Property Sheet: Options Page* on page 11-9)

Snapshot Logs Property Sheet: General Page

The General page contains the following:

On Master Table

Schema: Name of the schema to which the master table belongs. The drop-down list displays all schemas available to the user.

Table: Name of the master table. The drop-down list displays all tables belonging to the selected schema.

Tablespace

Tablespace to which the snapshot log belongs.

Log Table

Name of the table that stores the ROWID and timestamp of rows updated in the master table.

Log Trigger

Name of the AFTER ROW trigger on the master table that tells Oracle to insert the ROWIDs and timestamps of inserted, updated, and deleted rows into the master snapshot log.

Current Snapshots:

The number of snapshots associated with the snapshot log.

Create Like Snapshot Log

To create a snapshot log with parameters that are identical to an existing log:

1. Select an existing snapshot log from the navigator.
2. Choose Create Like from the Object menu. The Create Snapshot Log property sheet appears.
3. Enter a name and modify the property sheet entries (if necessary).
4. Click Create.

Altering a Snapshot Log

To alter an existing snapshot log:

1. Select a specific snapshot log in the navigator. A property sheet with the log's current parameters displays.
2. Modify the property sheet parameters.
3. Click Apply.

Snapshots Folder

The Snapshot folder contains of a list of all snapshots defined within the database.

A snapshot is a read-only copy of a master table located on a remote node. It is periodically refreshed to reflect changes made to the master table. A snapshot can be queried, but not updated; only the master table can be updated.

Snapshots Multi-column List

The Snapshots multi-column list displays when the Snapshots folder is selected in the navigator. The list contains all snapshots defined within the connected database.

The columns of the snapshots list are described below:

Schema

Name of the schema containing the snapshot. This column is not shown when viewing a particular schema.

Snapshot

Name of the snapshot.

Master Owner

Owner of the snapshot's master table.

Master Table

Name of the master table.

Master Link

Database link to the master site.

Last Refresh

Date of the last refresh (at the master site).

Type

Type of refresh for automatic refreshes: Complete, Fast, or Force.

Creating a Snapshot

To create a new snapshot:

1. Choose Create from the Schema Manager Object menu.
2. Select Snapshot from the list in the Create Object dialog box. The Create Snapshot property sheet appears.
3. Enter the requisite information.
4. Click Create.

The property sheet consists of four pages:

- General
- Cluster
- Storage (Advanced Mode only, see *Schema Object Property Sheet: Storage Page* on page 11-7)
- Options (Advanced Mode only, see *Schema Object Property Sheet: Options Page* on page 11-9)

Snapshot Property Sheet: General Page

The General page contains the following:

Name

Name of the snapshot being defined. During the Create operation, you can specify a valid Oracle identifier as the name of the sequence.

Schema

Name of the schema containing the snapshot. The drop-down list contains all available schemas within the connected database.

Tablespace

Name of the tablespace to which the snapshot belongs.

Refresh Type

Type of refresh for automatic refreshes: Complete, Fast, or Force.

Complete: Specifies a refresh that executes the subquery.

Fast: Specifies a refresh using only the updated data stored in the snapshot log associated with the master table.

Force: Specifies a fast refresh if one is possible or a complete refresh if fast refresh is not possible.

Start Date: Date snapshot refresh is to begin.

Next Date: Expression that calculates the interval at which the snapshot is updated.

Updatable

When selected, indicates that the snapshot is updatable. When used in conjunction with the Replication Option, updates are propagated to the master.

Snapshot Subquery

Snapshot query executed by Oracle when you create the snapshot. Results of the query are placed in the snapshot. (Create only)

Snapshot Property Sheet: Cluster Page

The Cluster page allows you to define the cluster to which a snapshot belongs. This property sheet contains the following:

Cluster

Cluster to which the snapshot belongs. The drop-down list displays all clusters contained within the selected schema.

Column Name: Name of the column to be used as part of the snapshot. Click Add to add the column name to the Snapshot Columns list.

Snapshot Columns: Scrolling list of all columns added to the snapshot. Use the Up/Down arrows to reorder the list of columns (column order must match the order in the cluster).

Remove: Deletes the selected snapshot column from the list.

Create Like Snapshot

To create a snapshot with parameters that are identical to an existing snapshot:

1. Select an existing snapshot from the navigator.
2. Choose Create Like from the Object menu. The Create Snapshot property sheet appears.
3. Enter a name and modify the property sheet parameters (if necessary).
4. Click Create.

Editing a Snapshot

To edit an existing snapshot:

1. Select a snapshot in the navigator. A property sheet with the snapshot's current parameters displays.

This property sheet is similar to the Create Snapshot property sheet except that the Subquery entry field is replaced by the following information:

Trigger

Associated AFTERROW trigger on the master table.

Log

Associated snapshot log.

Master Info

Owner: Name of the schema to which the snapshot belongs.

Table: Name of the master table.

Link: Database link to the database where the master table resides.

Can Log

Whether or not the snapshot can be logged.

Ref. Group

Associated refresh group.

Last Refreshed

Date the snapshot was last refreshed.

2. Modify the property sheet parameters as necessary.
3. Click Apply.

Synonyms Folder

The Synonyms folder contains a list of all synonyms defined within the connected database.

A synonym is an alias for a table, view, sequence, procedure, function, package or another synonym.

Synonyms Multi-column List

The Synonyms multi-column list displays when the synonyms folder is selected in the navigator. The list summarizes information for synonyms defined within the connected database.

The columns of the Synonyms multi-column list are described as follows:

Schema

Name of the schema containing the synonym. This column is not shown when viewing a particular schema.

Synonym

Name of the synonym.

Object Owner

Owner of the original database object.

Object

Name of the original database object.

Link

Database link to the the original schema object (e.g. table, view, snapshot).

Creating a Synonym

To create a new synonym:

1. Choose Create from the Object menu.
2. Select Synonym from the list in the Create Object dialog box. The Create Synonym property sheet appears.
3. Enter the requisite information.
4. Click Create.

The Create Synonym property sheet contains a single General page.

Synonym Property Sheet: General Page

The General page consists of the following information:

Name

Name of the synonym.

Schema

Name of the schema containing the synonym. The drop-down list contains all available schemas within the connected database.

As Alias For

Local Database/Remote Database: Allows you to create a synonym from an object on a local database or an accessible database on the network.

Object Type (Local)/DB Link (Remote): Drop-down list containing available database object types (tables, snapshots, etc.) for local databases. If Remote Database is selected the field label changes to DB Link. In this case, the drop-down list contains all database links defined for the current instance.

Schema: Drop-down list containing all available schemas for the current instance (local databases). If Remote Database is selected, the schema name must be typed in explicitly.

Object: Drop-down list containing all available database objects for the selected schema.

Create Like Synonym

To create another synonym with parameters that are identical to an existing synonym:

1. Select an existing synonym from the navigator.
2. Select Create Like from the Object menu. The Create Like property sheet is displayed.
3. Enter a name.
4. Click Create.

Viewing a Synonym

To view the parameters of an existing synonym, select a synonym in the navigator. A property sheet with the original synonym's current parameters displays.

Table Types Folder

A nested table is an unordered set of data elements, all of the same datatype. This property sheet allows you to define a nested table type. This table type can be used as

- the datatype of a column of a relational table (scalar type)
- an object type attribute
- a PL/SQL variable, parameter, or function return type

Table Types Multi-column List

The Table Types multi-column list displays when the Table Types folder is selected from the navigator. The multi-column list contains the following columns:

Schema (not displayed if specific schemas are selected in the Navigator)
Schema to which the table type belongs.

Array Type Name
Name of the table type.

Element Type Schema
Schema to which the datatype (for elements within a nested table) belongs.

Element Type Name
Name of the datatype (for elements within a nested table).

Creating a Table Type

To create a new table type:

1. Choose Create from the Object menu.
2. Select Table Type from the list in the Create Object dialog box. The Table Types property sheet appears.
3. Enter the requisite sequence information.
4. Click Create.

The Table Types property sheet consists of a single General page.

Table Types: General Page

The General page consists of the following elements:

Name

Table type name.

Schema

The schema to contain the table type.

Show Errors/Hide Errors

Opens or closes the Error Messages dialog box. This button is active when the table type status is invalid.

Datatype: Specifies the datatype for elements of the nested table.

Schema: Scheme to which the datatype belongs.

Type: Datatype of the element(s) in the nested table.

Length: Number of bytes allowed for values defined in the column (for CHAR, VARCHAR2, NUMBER, RAW datatypes).

Precision (for NUMBER datatype): Number of digits to the right of the decimal point. If CHAR is selected in the Datatype drop-down list, you may specify length. If VARCHAR2 or RAW is selected, you may specify a length. If NUMBER is selected, you may specify a length, a length and a precision, or neither.

User defined object will be a reference (for user-defined types): Specifies that an object of the table datatype be a REF. This option is only available if a user-defined datatype is selected from the Type drop-down list.

Create Like Table Type

To create a table type with parameters that are identical to an existing table type:

1. Select an existing table type from the navigator.
2. Choose Create Like from the Object menu. The Create Table Type property sheet appears.
3. Enter a name and modify the property sheet parameters (if necessary).
4. Click Create.

Editing a Table Type

To alter an existing table type:

1. Select a specific table type from the navigator. A property sheet with the table type's current parameters displays.
2. Modify the datatype as necessary.
3. Click Apply.

Tables Folder

The Tables folder contains a list of all tables defined within the connected database.

As the basic unit of storage in a relational database management system, a table represents entities and relationships, and consists of one or more units of information (rows), each of which contains the same kinds of values (columns). For information about tables, see *Oracle Server Concepts*.

Tables Multi-column List

The Tables multi-column list displays when the Tables folder is selected in the navigator. The list summarizes information for tables defined within the connected database.

The columns of the Sequences multi-column list are described as follows:

Schema

Name of the schema in which the table is defined. This column is not shown when viewing a particular schema.

Table

Name of the table.

Tablespace

Name of the tablespace to which the table belongs.

Rows

Number of rows in the table. (Shown only if the cost-based optimizer is disabled.)

Note: Entries in the Rows column only appear if ANALYZE has been run against the table. For more information on the ANALYZE command, see *Oracle Server SQL Reference*.

Creating a Table

To create a new table:

1. Choose Create from the Object menu.
2. Select Table from the list in the Create Object dialog box. The New Table dialog box appears. From here, you can choose to create a new table using either the Table Wizard or manually using the Create Table property sheet.

The Table Wizard is a rule-based series of dialogs that guide you the process of creating a table. See Oracle Enterprise Manager online help for more information about the Table Wizard.

3. Enter the requisite property sheet information.
4. Click Create.

The Create Table property sheet contains the following pages and dialog boxes:

- General
- Constraints
- Cluster Columns
- Storage (Advanced Mode only, see *Schema Object Property Sheet: Storage Page* on page 11-7)
- Options (Advanced Mode only, see *Schema Object Property Sheet: Options Page* on page 11-9)

Table Property Sheet: General Page

The General page allows you to define or view the primary attributes of a table such as column definitions, schema in which the table resides, or query used to populate a new table. The General page contains the following information:

Name

Name of the table being defined. During the Create operation, you can specify a valid Oracle identifier as the name of the table.

Schema

Schema in which the table is defined. The default schema for a new table is the user's default schema. During the Create operation, you can choose a schema from those defined in the connected database via the drop-down list.

Tablespace

Tablespace to which the table belongs.

Define Columns/Define Query/Object Table

This area of the property page changes according to the option selected.

Define Columns

An editable spreadsheet allowing you to edit or add columns to a database table. The columns of this spreadsheet consist of Column (name), Datatype, Default, Nulls? and Cluster Column. New additions (defined but not yet created) are graphically identified by a hand with a plus sign. After selecting an entry in the spreadsheet, you can use the context-sensitive menu (right mouse button) to carry out standard editing functions (cut, copy, paste, delete).

The spreadsheet consists of 10 columns:

Name: Name of the table column being defined or edited. You can specify a valid Oracle identifier as the name of the column. This field is disabled for columns that have already been defined in the database.

Schema: Schema containing the desired datatype. If a built-in datatype is selected (for example, VARCHAR2), schema does not apply and <NONE> is displayed.

Datatype: Drop-down list of user-defined object types contained within the selected schema.

Note: If a user-defined datatype is selected, the nested table dialog appears automatically. If you close the Nested Table dialog but need to access it again at a later time, you can call up the dialog by clicking on the spreadsheet button to

the immediate left of the Column Name and then use the right mouse button to call up the context-sensitive menu. Select Nested Tables...

Length: Number of bytes allowed for values defined in the column (for CHAR, VARCHAR2, NUMBER, RAW).

Precision (for NUMBER datatype): Number of digits to the right of the decimal point. If CHAR is selected in the Datatype drop-down list, you may specify length. If VARCHAR2 or RAW is selected, you may specify a length. If NUMBER is selected, you may specify a length, a length and a precision, or neither.

Ref?: When checked, indicates the column is a reference to a row in an object table (type Ref).

Nulls?: When checked, indicates that the column being defined is not permitted to contain null values.

Default Value: An expression that serves as the default value for this column in any rows for which the INSERT statement omits a value for the column.

Scope Schema: Schema containing the table to be referenced. Cell is active when Ref? is checked.

Scope Table: Table being referenced. Cell is active when Scope is checked.

Define Query

When selected, a scrolling editable text area appears allowing you to enter a SQL query to be used to populate the table. A query does not have to be specified (Create mode only).

Object Table

When selected, allows you to create an object table utilizing user-defined object types.

Schema: Schema containing the user-defined object type.

Object Type: Name of the user-defined object type. Upon selecting an object type, the Nested Table dialog is automatically displayed if there are any aggregated nested tables that for the selected object type. Note: VARRAY, Nested Table, or built-in types are not displayed.

Nested Tables: Displays the Nested Tables dialog box allowing you to create a storage table for each nested table column. The button is active when a user-defined object type is selected.

Object Type Spreadsheet: An editable spreadsheet displaying the attributes of the selected user-defined type. The spreadsheet consists of the following columns:

- **Name:** Attribute name. Cell is not editable.
- **Datatype:** Datatype of the attribute and schema containing that datatype. Cell is not editable.
- **Nulls:** When checked, indicates that the column being defined is not permitted to contain null values.
- **Default Value:** An expression that serves as the default value for this column in any rows for which the INSERT statement omits a value for the column.
- **Scope:** When checked, allows you to restrict the scope of the reference to a single table. Cell is active when Ref is checked.
- **Schema:** Schema containing the table to be referenced. Cell is active when Scope is checked.
- **Table:** Table being referenced. Cell is active when Scope is checked.

Table Property Sheet: Constraints Page

The Constraint page allows you to define or edit integrity constraints for a table. This page contains the following information.

Constraints on Table

An editable spreadsheet containing the following columns:

Name: Name of the table constraint being defined. To define a new constraint, you can enter a valid Oracle identifier as the name of the constraint. If no name is entered, a default name will be assigned by the database. This field is disabled if the column has already been defined in the database.

Type: The type of constraint. The drop-down list displays available constraint types: UNIQUE, PRIMARY, FOREIGN, CHECK.

Disable: Indicates that the constraint should be disabled (checked) or enabled ('x') when the constraint is created. Click on the spreadsheet cell to toggle this setting.

Referenced Schema: Schema referenced by the foreign key in the constraint being defined. The drop-down list displays all available schema for the database. This list is active only when FOREIGN is selected as the constraint type.

Referenced Table: Table referenced by the column being defined. This drop-down list displays all tables contained within the Referenced Schema.

Cascade on Delete: When checked, indicates that rows will automatically be removed from the child table if they reference values contained in the referenced key column of rows being removed from the parent table. You toggle this setting by clicking on the desired spreadsheet cell. This control is disabled if a query is defined at Create time or if the column has already been defined in the database.

Check Condition: Check constraint to be included in the table constraint currently being defined.

You can enter the text of the check constraint in this spreadsheet cell or a valid expression. For more information see the *Oracle Server SQL Reference* for your server release.

Constraint Definition

An editable spreadsheet displaying both table and referenced columns.

Table Columns: Column on which the constraint is to be placed. The drop-down list displays all columns in the table. This cell is active for UNIQUE and PRIMARY key constraints. If you are defining an object table, this will be a list of all attributes of the object type the object table is defined on.

Referenced Columns: Columns (unique or primary key) referenced by a foreign key in a referential integrity constraint. The drop-down list displays all columns available in the referenced table.

Table Property Sheet: Cluster Columns

The Cluster Columns page (active when creating a new table) allows you to associate database cluster key columns with columns in a newly created table.

Cluster

Drop-down list displaying all clusters available in the schema selected on the General page of the Table property sheet.

Table Column : Drop-down list displaying all columns in the new table.

Cluster Column: Cluster key columns contained within the selected cluster.

To clear the Table Column entries, reselect the current cluster from the drop-down list.

Dropping a Constraint on a Table

To drop a constraint from a table:

1. Using the left mouse button, click on the entry select button located to the immediate left of an individual row in the Constraints on the Table spreadsheet.
2. Click the right mouse button to display the pop-up menu.
3. Select Drop to delete the constraint, or Drop with Cascade to tell Oracle to drop the constraint and maintain referential integrity by also removing dependent foreign key values if a referenced primary or unique key value is removed.

Removing a Constraint Definition (Create Mode Only)

To remove a constraint definition:

1. Using the left mouse button, click on the entry select button located to the immediate left of an individual row in the Constraints on the Table spreadsheet. The table columns or referenced columns appear in the Constraint Definition spreadsheet.
2. In the Constraint Definition spreadsheet, select the desired constraint by clicking on the entry select button located to the immediate left of an individual row in the Constraint Definition spreadsheet.
3. Click the right mouse button to display the pop-up menu.
4. Select Remove.

Create Like Table

To create a new table with parameters that are identical to an existing table:

1. Select a table from the navigator and then select Create Like from the Object menu. The Create Table property sheet appears with all parameters filled in except Table (name), based on the values from the selected table.
2. Click Create to create the table and close the property sheet.

Altering a Table

To alter an existing table:

1. Select a table in the navigator. A property sheet with the table's current parameters displays.
2. Modify the parameters as necessary.
3. Click Apply.

Triggers Folder

The Triggers folder contains a list of triggers defined within the connected database.

A trigger is a stored PL/SQL procedure that is implicitly executed when an INSERT, UPDATE, or DELETE statement is issued against the associated table or view. For more information about using triggers, see *Oracle Server Administrator's Guide* for your server release.

Triggers Multi-column List

The Triggers multi-column list displays when the Triggers folder is selected in the navigator. The list summarizes information for all triggers defined within the connected database.

The columns of the Triggers multi-column list are described as follows:

Schema

Name of the schema containing the trigger. This column is not shown when viewing a particular schema.

Trigger

Name of the trigger.

Type

Trigger type (BEFORE EACH ROW, AFTER EACH ROW, or FOR EACH ROW).

Event

Triggering statement (INSERT, UPDATE, DELETE).

Table Owner

Owner of the table with which the trigger is associated.

Table

Name of the table with which the trigger is associated.

Creating a Trigger

To create a new trigger:

1. Choose Create from the Object menu.
2. Select Trigger from the list in the Create Object dialog box. The Create Triggers property sheet appears. The Create Trigger property sheet appears.
3. Enter the requisite information.
4. Click Create.

The property sheet consists of two pages:

- General
- Timing

Trigger Property Sheet: General Page

The General page allows you to define the primary characteristics of a trigger such as the name of trigger, schema to which the trigger belongs, and the trigger body. This page contains the following information:

Name

Name of the trigger being defined.

Schema

Schema in which the trigger is to be defined. The default schema for a new trigger is SYSTEM.

Trigger On

Table/View: When selected, allows you to create a trigger on a table or a view.

Schema: Drop-down list displaying all available schema in the database. This field specifies the schema that contains the table.

Name: Drop-down list displaying all available tables or views in the selected schema. This field specifies the name of the table on which the trigger is to be created. Triggers on SYS schema tables cannot be created.

Replace if Exists

When selected, indicates that a trigger is to be recreated (if it already exists) without dropping, recreating and regenerating object privileges previously granted.

Enabled

When checked, indicates that a trigger is enabled.

During a Create operation, this control is checked and disabled. It is not checked and enabled if a trigger is disabled.

Trigger Body

Defines the PL/SQL block that Oracle executes to fire the trigger.

Trigger Property Sheet: Timing Page

The Timing page allows you to specify the conditions that initiate a trigger. This page contains the following information:

Triggering Statement

Fire Trigger:

- *Before:* Indicates that Oracle fire the trigger before executing the triggering statement.
- *After:* Indicates that Oracle fire the trigger after executing the triggering statement.
- *Instead Of:* Indicates that Oracle fire the trigger instead of executing the triggering statement. This option is only valid for triggers on views. You cannot specify an *INSTEAD OF* trigger on a table.

Insert: When checked, indicates that Oracle fire the trigger whenever an INSERT statement adds a row to a table.

Delete: When checked, indicates that Oracle fires the trigger whenever a DELETE statement removes a row from a table.

Update of Columns: Indicates that Oracle fires the trigger whenever an UPDATE statement changes a value in a column selected in the multi-select list. It is only enabled and populated when the Update checkbox is selected.

Trigger for Each Row

When checked, designates the trigger to be a row trigger. Oracle fires a row trigger once for each row that is affected by the triggering statement and meets the optional trigger constraint defined in the Condition field. When not checked, the trigger is a statement trigger and is fired only once when the triggering statement is issued (assuming the optional trigger constraint is met).

Referencing

- *Old as:* New name for default correlation name OLD.
- *New as:* New name for default correlation name NEW.

You can use correlation names in the PL/SQL block and WHEN clause of a row trigger to refer specifically to old and new values of the current row. The default correlation names are OLD and NEW.

Condition : Specifies the trigger restriction.

The trigger restriction contains a SQL condition that must be satisfied for Oracle to fire the trigger. This condition must contain correlation names and not a query. You can only specify a trigger restriction for a row trigger. Oracle evaluates this condition for each row affected by the triggering statement. See the Oracle SQL Reference manual for more information.

Create Like Trigger

To create a new trigger with parameters that are identical to an existing trigger:

1. Select a trigger in the navigator.
2. Select Create Like from the Object menu. The Create Trigger property sheet appears with all parameters filled in, except the name, based on the values from the selected table.
3. Click Create.

Altering a Trigger

To alter an existing trigger:

1. Select a trigger in the navigator. The Trigger property sheet with the selected trigger's current parameters displays.
2. Modify the trigger body or triggering statement as necessary.
3. Click Apply.

Altering a trigger only affects values generated after the alter is performed.

Views Folder

The Views folder contains a list of all views defined within the connected database.

A view is a custom tailored presentation of the data in one or more tables and can also be thought of as a *stored query*. Views do not actually contain or store data; rather, they derive their data from the tables on which they are based, called base tables. Base tables can in turn be actual tables or can themselves be views. Views can be queried, updated, inserted into, and deleted from, with restrictions. All operations performed on a view actually affect the base tables of the view. For information about views, see *Oracle Server Concepts*.

Views Multi-column List

The Views multi-column list displays when the Views folder is selected in the navigator. The list summarizes information for tables defined within the connected database.

The columns of the Views multi-column list are described as follows:

Schema

Name of the schema containing the view. This column is not shown when viewing a particular schema.

View

Name of the view.

Status

Current status of the view (Valid or Invalid).

Creating a View

To create a new view:

1. Choose Create from the Object menu.
2. Select View from the list in the Create Object dialog box. The Create View property sheet appears.
3. Enter the requisite information.
4. Click Create.

The Create View property sheet consists of two pages:

- General
- Advanced

View Property Sheet: General Page

The General page contains the following information:

Name

Name of the view being defined. During a Create operation, you can specify a valid Oracle identifier as the name of the view.

Schema

Schema in which the view should be defined. The default schema for a new view is the user's default schema.

Query Text

Identifies columns and rows of the table(s) that the view is based on. A view's query can be any SELECT statement without the ORDER BY or FOR UPDATE clauses. Its select list can contain up to 254 expressions.

View Property Sheet: Advanced Page

The Advance page allows you to specify options or place constraints on the view. This page consists of the following:

Options

Replace if Exists: When checked, specifies that the view is to be recreated, if it already exists, without dropping, recreating, and regenerating object privileges previously granted on it.

Force: When checked, specifies that the view is to be created regardless of whether the view's base tables exist or the owner of the schema containing the view has privileges on them.

Read Only: When checked, specifies that no deletes, inserts or updates can be performed through the view.

Constraint

Specifies the name assigned to the CHECK OPTION constraint.

With Check Option: When checked, specifies that inserts and updates performed through the view must result in rows that the view query can select.

The CHECK OPTION cannot make this guarantee if there is a subquery in the query of this view or any view on which this view is based.

If you omit this identifier, Oracle automatically assigns the constraint a name of the form: SYS_Cn, where n is an integer that makes the constraint name unique within the database.

As Object

Allows you to create an object view.

Schema: Schema containing the object type.

Object Type: List of user-defined object types contained within the selected schema.

Object ID: Allows you to specify the attributes of the object type that will be used as a key to uniquely identify each row in the object view.

Default: Specifies that the intrinsic object identifier of the underlying object table or object view be used to uniquely identify each row.

Specify Attributes: Allows you to specify attributes of an object type from which the object view is to be created. You can specify the unique order of the table columns using the two-column spreadsheet. Table columns correspond to the top-level attributes of the object type.

Create Like View

To create a new view with parameters that are identical to an existing view:

1. Select a view from the navigator.
2. Choose Create Like from the Object menu. The Create View property sheet appears with all parameters filled in except the view name.
3. Enter a new name and modify any parameters (if necessary).
4. Click Create.

Altering a View

To alter an existing view:

1. Select a view in the navigator. A property sheet with the view's current parameters displays.
2. Modify the query text or other property sheet parameters as necessary.
3. Click Apply.

Managing Backup and Recovery

This chapter describes how to use Oracle Enterprise Manager Backup Manager to administer your database backup and recovery environment. Backup Manager consists of three subsystems: Oracle8 Recovery Manager, Operating System Backup, and Enterprise Backup Utility. Subsystem availability depends on the version database you are attempting to back up. For Oracle7 databases, you can select Operating System Backup or the Enterprise Backup Utility. For Oracle8 databases, you can select Operating System Backup or Oracle8 Recovery Manager. All backup subsystems use wizards to simplify and automate backup and recovery tasks. Briefly, the three subsystems are:

- **Backup Manager: Oracle8 Recovery Manager Subsystem:** Provides an intuitive graphical user interface for Recovery Manager, an extremely powerful and flexible command line backup and recover utility designed expressly for the Oracle8 database.
- **Backup Manager: Operating System Backup Subsystem:** Provides automated tablespace backup and recovery using the Oracle Enterprise Manager job system (Oracle7 databases). Operating System Backup also provides limited backup and recover capability for Oracle8 databases.
- **Backup Manager: Enterprise Backup Utility Subsystem:** Provides backup and recovery functionality for Oracle7 databases only. See your Enterprise Backup Utility documentation for further information. **Note:** Enterprise Backup Utility will be available on Windows NT beginning with Oracle 7.3.4.

At any time, you can change to a different backup subsystem by choosing Change Backup Subsystem from the File menu.

For more information on backup and recovery strategies, see your *Oracle Backup and Recovery Guide*. For more information on Recovery Manager, see the *Oracle8 Server Backup and Recovery Guide*.

Starting Backup Manager

Before attempting to start Backup Manager, make sure the Oracle Enterprise Manager Console is running and that the job and event system is fully functional. Select a target database from the Oracle Enterprise Manager Console's Navigator. Then click on the Backup Manager icon in the Applications Launch Palette, or choose Oracle Backup Manager from the Console Tools menu.

Attention: Backup Manager requires connecting to the target database as SYSDBA. By default, this role is selected during login. If your target database is not configured for remote operation, see the *Oracle Enterprise Manager Configuration Guide*.

You can change the database connection with the Change Database Connection option in the File menu. You can also change backup subsystems by choosing Change Backup Subsystem from the Backup Manger File menu.

After Backup Manager has successfully connects to the target database, the Backup Subsystem dialog displays. The following figure shows the backup subsystem options available when connecting to an Oracle8 database.

Figure 12–1 Backup Subsystem Dialog



If you are running this application against an Oracle8 target database for the first time, Backup Manager displays a dialog box asking where backup information will be maintained. You can use either a Recovery Catalog or the target database control file. For information regarding the pros and cons of using a recovery catalog rather than a control file, see the *Oracle8 Backup and Recovery Guide*.

If you are using Backup Manager for the first time, a dialog displays asking whether you wish to use a recovery catalog or a control file. See the online help for that dialog for more information. You must select one option before proceeding.

As shown in the following sections, each backup subsystem, once started, is very similar in appearance to one another. Only menus and options change.

Backup Manager: Oracle8 Recovery Manager Subsystem

This backup subsystem provides an easy-to-use graphical user interface to Recovery Manager, an Oracle utility you use to back up, restore, and recover database files. Recovery Manager starts Oracle server processes on the database to be backed up or restored (the target database). These Oracle server processes actually perform the backup and restore. For example, during a backup, the server process reads the files to be backed up, and writes the files out to your tertiary storage device.

Note: Making backups to tertiary media, such as tape, requires vendor-supplied media management software that can interface with Oracle. Oracle calls the media management software routines to back up and restore datafiles to an from media controlled by the media management software.

Recovery Catalog

Recovery Manager capabilities are extended when used with a recovery catalog, a repository of information that is used and maintained by Recovery Manager.

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.

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. However, Oracle recommends using a recovery catalog. If the controlfile is lost, the database cannot be recovered.

Important: The recovery catalog and the Oracle Enterprise Manager repository should not reside in the target database (database to be backed up) because the database can not be recovered in the mounted state. If you loose any data, you will need the recovery catalog for full recovery. However, 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 backup your recovery catalog regularly.

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

Registering a Database with the Recovery Catalog

1. Enter the requisite login information: user=RMAN, password=RMAN (or any user you create to access the database with the recovery catalog).
2. From the Catalog menu, select Use Recovery Catalog.
3. From the Catalog menu, select Register. You are asked if you wish to register the target database in the recovery catalog. Click Yes.

At this point, the database registration is sent as a job to the Oracle Enterprise Manager Job system.

4. Select the Jobs object in the tree list. The Jobs property sheet displays showing the current status of the registration job.
5. From the Jobs property sheet, make sure that the job completed successfully.

When the job completes successfully, your backup and recovery environment is configured. If the job does not succeed, certain functions (those requiring a recovery catalog) will be unavailable. For more information on recovery catalog maintenance, Recovery Manager, and backup and recover strategies, see the *Oracle8 Backup and Recovery Guide*.

Oracle8 Recovery Manager Subsystem User Interface

After selecting Oracle8 Recovery Manager from the Subsystem selection dialog box, the Backup Manager window appears.

Figure 12–2 Backup Manager: Oracle8 Recovery Manager Subsystem



Oracle8 Recovery Manager Menu Options

Backup Manager using the Oracle8 Recovery Manager subsystem includes four standard menus, File, View, Log, and Help. It also includes the Backup, Recover, Catalog, Database, and Channel menus. For information on the standard menus, see *Application Menus* on page 7-10.

Backup Menu

Backup Set

Allows you to define a backup set for a selected object in the tree or multi-column list. Selected objects must consist of the same type (datafiles or archivelogs).

Image Copy

Allows you to define an image copy of the selected object. An image copy contains a single input file (datafile or control file).

Backup Controlfile to Trace

Allows you to back up the target database control file by writing SQL commands to the target database's trace file to recreate the controlfile if it is lost and no backup of the ifle is available..

Backup Wizard

Simplifies the job of backing up a database, tablespaces, datafiles, or archivelogs using an intelligent Wizard to guide you through the backup process.

Recover Menu

Restore

Restore the selected database objects in the tree or multi-column list. Restoring a datafile or tablespace requires that these objects be offline, or that the database be closed or in a mounted or unmounted state.

Switch

Switch the datafile copy into a current datafile. Media recovery is required after switching the datafile.

Recover

Recover the selected object in the tree list.

Error Status

Displays Needs Recovery status for datafiles.

Restore Wizard

Starts the Restore Wizard that guides you through the restore, rename, and recover processes.

Catalog Menu

Resync

Re-synchronizes the Recovery Catalog for 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 an log is archived. Also, any structural changes to the target database would require re-synchronization of the Recovery Catalog.

Reset

Reset the Recovery Catalog for the target database. You must reset the recovery catalog if the target database had been previously opened with the RESETLOGS option.

Register

Register the target database with the Recovery Catalog. .

Connect string

Allows you to define the connect string to the Recovery Catalog.

Use Recovery Catalog

Use the Recovery Catalog or *VSviews* for database backup and recovery.

Report Wizard

Simplifies the job of creating detailed analysis reports for the target database. Reports include: Files requiring backup, unrecoverable files, obsolete files, and the database schema at a point in time.

Stored Script

Allows you to create or access user-defined scripts consisting of Recovery Manager commands. These scripts are stored in the recovery catalog.

Database Menu

Shutdown

Shuts down the target database.

Startup

Starts a target database instance.

Switch Logfile

Switches the current redo log group.

Force Checkpoint

Triggers a checkpoint. Note: After a checkpoint has been triggered, the redo in the redo log files is no longer useful for crash/instance recovery.

Add Log Group

Creates a new redo log group.

Drop Log Group

Deletes the selected redo log group.

Add Log Member

Adds a logfile to a redo log group.

Drop Log Member

Deletes the selected redo logfile member.

Channel Menu

Create

Allows you to create a new channel definition.

Create Like

Allows you to create a new channel with the same connection properties as the selected channel.

Remove

Deletes the selected channel.

Tree List Objects and Containers

Database Object The database object displays the name of the target database. When selected, the database property sheet displays providing you with information regarding database status and archive mode. This property sheet contains the following information.

Database Property Sheet: Status Page

The Status page contains information about the state of the current database, database version, and any installed options. This page also allows you to change the state of the database by selecting one of the Database State options and clicking Apply.

Database State

Shutdown: Database closed.

Instance Started: Instance started but not mounted.

Database Mounted: Database mounted but not open.

Database Open: Database mounted and open.

Database Version

Database version and any installed options.

Attention: If the database is not running, the message "ORACLE not available" is displayed.

Database Property Sheet: Information Page

The Information page contains displays the current state of redo log archival and the system global area.

Archive Information

Log Mode: Specifies whether the connected database is in ARCHIVELOG mode or NOARCHIVELOG mode. To easily switch log modes, you can use Oracle Instance Manager.

Last SCN Archived: The last system change number that was archived. The SCN uniquely identifies the last committed database transaction.

Archive Destination: Specifies the destination where the archive log files are to be created. If you are archiving to disk, it is recommended that a dedicated disk with sufficient storage be used.

Archive Format: Specified the naming convention for the archived log files. The following variables are appended to form a unique archive file name.

- %S Backup Set Number— This value is a control file counter that is incremented for each backup set. The counter value starts at 1 and is unique for the lifetime of the control file.
- %T Backup Stamp—This value indicates the number of seconds since a fixed reference date/time (currently midnight of January 1, 1988).

Automatic Archival: Allows you to specify that redo log files be automatically archived. To enable Automatic archiving upon instance startup, set the LOG_ARCHIVE_START initialization parameter to TRUE. This option is enabled when the connected database is running in ARCHIVELOG mode.

SGA Information

Database Buffers: Size of the database buffer cache (in bytes).

Fixed Size: Memory allocated to the area of the SGA that contains general information about the state of the database and the instance. No user data stored here (in bytes).

Redo Buffers: Size of the redo log buffer (in bytes).

Variable Size: Memory allocated to variable size data structures (in bytes).

Controlfile Object The controlfile object represents the controlfile of the target database. When selected from the tree-list, the control file property sheet displays. This property sheet displays all available control files used by the target database. The multi-column list displays the name and location of the control file as well as its status.

Tablespace Containers, Multi-column Lists and Property Sheets

Tablespace Container

The Tablespaces object type container contains each tablespace in the database arranged alphabetically in a tree list. A particular tablespace can be expanded to show each of its datafiles.

Multi-column List for the Tablespaces Container

When you select the Tablespaces container in the tree list, a multi-column list displays on the right. This list includes a row for each tablespace in the database. The columns of the list are:

Name

Name of the tablespace.

Status

Backup status of the tablespace: Active (an online backup of this tablespace is currently in progress), Not Active (no online backup of this tablespace is currently in progress); or Offline.

Size (M)

Total size in megabytes of the datafiles that comprise the tablespace.

Used (M)

Amount of actual storage used.

Tablespace Property Sheet: General Page

The Tablespace property sheet allows you to alter the status of the tablespace selected in the tree list. After you select In addition, all datafiles contained within the tablespace are displayed.

Name

The name of the selected tablespace.

Status

Indicates the current status of the selected tablespace and also allows you to change the status. The following options are available:

Online: Places the selected tablespace online.

Offline: Takes the selected tablespace offline. When this option is selected, a dropdown list appears to the right allowing you to select the following priorities:

Normal: Used if no error conditions exist for any of the datafiles of the tablespace.

Temporary: Used if there are error conditions for one or more datafiles of the tablespace.

Immediate: Note: You cannot take a tablespace offline immediately if the database is running in NOARCHIVELOG mode.

Read Only: Makes the selected tablespace read-only to prevent further write operations on the datafiles in the tablespace. A tablespace should be backed up before being made read only.

Datafiles

A scrolling multi-column list displaying all datafiles contained within the tablespace. For each datafile, the full name and path, size (K), date and time of the last backup, whether the datafile needs to be backed up, and any datafile errors are displayed.

Datafiles Container, Multi-column Lists and Property Sheets

The Datafiles object type container contains all datafiles in the target database. Selecting an individual datafile from the tree list displays its property sheet.

Datafiles Multi-column List When you select the datafiles container, a multi-column list displays containing the following columns:

Name

Full path and name of the datafile.

Tablespace

Tablespace to which the datafile belongs.

Error

Recovery Manager error, if any.

Needs Recovery

Status as to whether the datafile requires recovery.

File Status

RECOVER, SYSTEM, ONLINE, OFFLINE

Size (K)

Size of the datafile in kilobytes.

Datafiles Property Sheet: General Page

The General page lets you view general information such as name, location, and size of the selected datafile. This page also allows you to change the online/offline status of the selected datafile.

Name

Full path and name of the selected datafile. This field is read-only.

Tablespace

Tablespace to which the datafile belongs. This field is read-only.

Status

The current datafile online/offline status is selected. You can change the status of the datafile by clicking on the desired option and then clicking the Apply button

Size (K bytes)

Datafile size in kilobytes.

Backups Page

The Backup page lists all backup sets or image copies ever taken for the selected datafile.

Time

Date and time of the backup.

Backup Type

Indicates whether or not the backup is a BACKUP SET or an IMAGE COPY. Note: Datafile backup sets may also contain a controlfile backup.

Level

Indicates whether a backup of the datafile is full or incremental. Levels range from 0 to 4. Levels one to four are user-defined and allow you to set up different levels of incremental backups. For example, you might set up the following backup scheme:

- Level 0-Monthly full backup
- Level 2-Weekly incremental
- Level 1-Daily incremental

Note: Incremental backups can only be performed with the Oracle Enterprise Edition, not the Workgroup Server.

For explicit information on using levels in your backup environment, see the *Oracle8 Backup and Recovery Guide*.

Tag

User-defined name representing one or more backup sets or image copies. The maximum length of a tag is 30 characters.

Name

Full path and name of a datafile.

Redo Log Groups Container, Multi-column Lists, and Property Sheets

Redo Log Group Containers The Redo Log Groups object type container contains each redo log group of the database. A particular redo log group can be expanded to show its redo log members contained within a Redo Log Members object type container.

Redo Log Groups Multi-column List When you select the Redo Log Groups object type container in the tree list, a multi-column list displays on the right. This list includes a row for each redo log group in the database. The columns of the list are:

Group

Group number of the redo log group.

Thread

Thread number used in Parallel Server environments. For single instance environments, the thread number is set to one.

Sequence

Sequence number of the redo log group.

Size (K)

Space allocated in kilobytes to the redo log group.

Status

Logging status of the redo log group: Current, Active, or Inactive.

Archived?

Indicates whether the redo log group has been archived. YES or NO.

Low SCN

System Change Number for the first change recorded in the redo log group.

Time of Low SCN

Date and time of the low SCN.

Redo Log Group Property Sheet: General Page

The Redo Log Group property sheet allows you to create a new redo log group, or view the parameters of an existing redo log group. The Create Redo Log Group property sheet includes these elements:

Group #

Displays the redo log group number that Oracle Enterprise Manager has automatically generated for this new redo log group. (For example, if redo log groups 1 and 2 already exist for the database, 3 appears by default in the Group # box.)

Thread

Used in Parallel Server environments, indicates the thread associated with a redo log group. For single instance environments, this number is set to one.

Files

Use Existing File(s): Allows reuse of an existing file(s).

New File(s) Size: Allows you to specify the size of the new Redo Log Group file(s) (Kilobytes or Megabytes).

Current Members

Displays the current members of this redo log group.

The Current Members box is empty when the Create Redo Log Group property sheet first displays. After you add a new member to this redo log group, the new member then appears in this box.

Remove

Remove the redo log group member that is selected in the Current Members box.

New Member

Name of the new member.

...

Displays a standard Open dialog box, from which you can select the redo log file you want to enter in the New Member box.

Redo Log Member dialog

The Redo Log Member dialog allows you to view an existing or add a new redo log member to a redo log group in the target database.

Filename

Filename of the new member of the redo log group shown in the Group box. When creating a new redo log member, this field is editable.

Group

Redo log group to which the redo log member belongs. When creating a new redo log member, you can select from all groups available in the target database using the drop-down list.

Use Existing File (Create Only)

When selected, specifies that Oracle Backup Manager reuse an existing file.

New File (Create Only)

When selected, specifies that Oracle Backup Manager create a new file.

Archived Logs Container, Property Sheets, and Multi-column lists

The Archived Logs Container displays the all archived logs for the target database.

Archived Logs Multi-column List When you select the Archived Logs container, a multi-column list displays containing the following columns:

Archived Log Name

Full path and name of the archive log.

First Time

Archive date and time for the first member in the redo log group.

Next Time

Archive date and time for the next member in the redo log group.

Delete

Status as to whether the archived log file will be deleted upon backup.

Archived Logs Property Sheet: General Page

The General page allows you to view parameters for the selected archived redo log file.

Filename

Full path and name of the archived redo log file.

Thread No.

Numeric archived redo logfile identifier used for parallel server environments. Thread numbers are used to identify archived redo logfiles belonging to a specific instance. A default value of 1 is used for single instance environments.

Log Sequence No.

Unique numeric identifier for the selected archived redo log file. During media recovery, Oracle applies redo log files in ascending order by using the log sequence number of the necessary archived and online redo log files.

Size (K Bytes)

Physical size of the redo log file in kilobytes.

Backed Up

Status as to whether the selected redo log file has been backed up

Archived Logs Property Sheet: Record Page

The Record page allows you to view the time stamp and system change number of the selected archived log.

Starts At

Date: Archive date and time for the first member in the redo log group.

System Change No.: Low SCN for the first member of the redo log group.

Ends At

Date: Archive date and time for the last member of the redo log group.

System Change No.: High SCN for the last member of the redo log group.

Channels Container

A channel establishes a connection from the Oracle8 Recovery Manager subsystem to the target database for backup or restore operations. Multiple channels can be created to allow multiple backup sets or file copies to be read or written in parallel by a single job. Important: Backup in parallel can only be performed with Oracle Enterprise Edition and not the Workgroup Server.

Note: At least one channel must exist before performing backup or restore operations.

Channels Multi-column list When you select the Channels folder, a multi-column list displays the following columns:

Channel Name

Name of the user-defined channel.

Type

DISK or TAPE.

Default Device

Status as to whether the channel is selected as the default channel for the target database.

Channels Property Sheet: General Page

The General page allows you to create, view, or modify a channel.

Note: At least one channel must exist before performing backup or restore operations.

Channel Name

User-specified name of the channel.

Default Channel for this database

When checked, specifies that the current channel be used by default whenever the currently connected target database is accessed.

Channel available for all databases

When checked, makes the current channel available to any target database.

Channel Type

Drop-down list allowing you to select from one of the following device types:

- **Disk:** Backup sets written to disk.
- **Tape:** Backup sets written to a tape device. The Destination field is disabled when this channel type is selected.

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

Destination (active when Channel Type is set to Disk): Drive and path where backup sets are stored.

Channels Property Sheet: Limits Page

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 the its value. The middle number on the scale 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.

Jobs Object

The jobs object contains all backup, restore, and recovery jobs that have either been submitted, have already been run, or are currently stored in the job library.

When the Jobs object is selected in the tree list, Jobs property sheets are displayed.

Jobs Property Sheets

Active Jobs The Active Jobs page contains a summary of the active jobs that have been submitted to the Oracle Enterprise Manager job system and are not yet completed. You cannot modify these jobs.

Job Name

Name of the job.

Status

Execution status of the job. Status may be one of the following:

- Submitted: The job has been submitted to the agent at the job destination.
- Scheduled: The job has been successfully delivered to the agent and is scheduled for execution.
- Starting: 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 Submitted. Unless you view the Active Jobs page at the exact time that the job is executing, you would not see the Starting 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.
- Completed: The job has been run successfully.

Next Execution Time

Date and time of the job start time.

Job Type

Backup or Restore.

Job History The Job History page contains a list of previous job activities, such as multiple executions of a job. These are jobs that have been submitted to an agent and have executed successfully or unsuccessfully. This page also lists deleted jobs. You cannot modify these jobs.

Job Name

The name of the job.

Status

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.

Execution Start Time

Time when the job started or was deleted. Because the order of agent notifications may vary, it is possible that the Console receives a completed or failed notification before a running notification. If this happens, the start time displays Unknown.

Execution End Time

Time when the job finished, failed, or was deleted.

Job Library The Job Library page displays jobs created and saved using either the Backup or Restore Wizard. This is useful if you want to submit the same job at different times or if you want to make minor modifications to an existing job without having to redefine it from scratch using one of the Wizards.

Double-click on an entry in the Job Library list to start the associated Wizard with the stored parameters already set. You can either submit the job, or modify the settings and save it again.

Job Name

Application generated name of the job.

Job Description

Application generated description of the job.

Oracle8 Recovery Manager Wizards

The Oracle8 Recovery Manager wizards guide you through complex tasks, allowing you to concentrate on high-level strategies for managing your backup and recovery environment.

Backup Wizard

The Backup Wizard allows you to select from a variety of backup operations. To start the Backup Wizard, select Backup Wizard from the Backup menu. The type of backup you select determines which procedure the Backup Wizard guides you through. The following backups types are available:

- Database
- Tablespaces
- Datafiles
- Archivelogs Only

For more information on backups and related strategies, see the *Oracle8 Server Backup and Recovery Guide*. Refer to the Oracle Backup Manager online help for a detailed user interface information.

Restore Wizard

The Introduction page of the Restore Wizard allows you to select from the following types of restore operations:

- Database
- Tablespaces
- Datafiles

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

Depending on the status of the target database (ARCHIVELOG, NOARCHIVELOG, mounted and open, or mounted only), some options will be disabled.

Refer to Oracle Backup Manager online help for a detailed information.

Report Wizard

The Report Wizard allows you to submit a job that generates backup reports. The results of the report are viewed by double clicking on the job entry in the Job History page of the Jobs property sheet. The Job Output dialog displays with the report results. You can use the Report wizard to generate reports for the following backup issues:

- Files requiring backup: Lists all datafiles that are in need of backup. This report assumes that the most recent backup would be used in the event of a restore operation.
- Files which are unrecoverable: Lists all datafiles that are considered unrecoverable (an UNRECOVERABLE operation has been performed against an object residing in the datafile since the last backup of the datafile).
- Files obsolete-can be deleted: Lists redundant backup sets and datafile copies that can be deleted.
- Database Schema at a point-in-time: Lists physical schema of a database (datafiles and tablespaces) at a specified point in time, or at the current time.

Note Files requiring backup and Files which are unrecoverable reports should be generated on a regular basis to ensure that the necessary backups are available to perform recovery, and to ensure that the recovery can be performed within a reasonable length of time.

Manual Backup Operations

In addition to the automated backup procedures you can perform using the Backup Wizard, Oracle8 Recovery Manager subsystem also allows you to manually perform specific backup tasks such as creating a backup set or an image copy.

Creating a Backup Set A backup set is a backup of one or more Oracle files, where files are multiplexed together. These files are called backup pieces. Files within a backup set must be extracted using a restore operation.

To create a backup set:

1. Select a tablespace or datafile from the tree list.
2. Choose Backup Set from the Backup menu. The BackupSet property sheet appears.

BackupSet Property Sheet: General Page

The General page allows you to specify the backup level, a tag, and whether or not the control file is backed up. To backup a range of datafiles, use the Backup Wizard.

Choose an incremental backup level

Click on the up or down arrow to specify a backup level. Backup Manager allows you to specify up to five levels of backup: 0: full backup, 1-4: incremental backup

Enter or choose a tag

Optionally, you can specify a tag by entering a character string of less than 30 characters. You can also use the drop-down list box to select an existing tag.

Defining a tag gives you a symbolic way to refer to a collection of image copies or backup sets. You can also use tags to specify specific input files to a restore or switch.

Include control file with this backup

Select this option if you want the control file backed up with the backup set.

Backup Set Property Sheet: Channels Page

The Channels page allows you to select the channel(s) used for the current backup set.

Choose channel type

Drop-down list allowing you to select one of the following channel types:

- Disk: Displays all available channels that write to disk.
- Tape: Displays all available channels that write to a sequential I/O device or use an OS-specific device access method.

Selected channels

List of channels selected for use by the current backup.

To remove an entry from the Selected channels list, double-click on the entry.

Available channels

List of available channels for use by the current backup.

To add an entry to the Selected channels list, double-click on the Available channels entry.

Note: Multiple selection of channels is only supported with the Oracle Enterprise Edition.

Backup Set Property Sheet: Parameters Page

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

Max. Files per Backup Set

When selected, 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.

Backup Set/Image Copy Property Sheet: Schedule Page

The Schedule page allows you to schedule the execution of a backup or image copy.

Execute

Drop-down list allowing you to select the frequency of backup or image copy execution. 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 setup process. The task executes only one time.
- Once schedules the task only one time at the date and time you choose.
- On Interval allows you to schedule a specific time interval between task executions. The interval can be a combination of hours and minutes, or number of days. Select the value you want to change and click on the scroll buttons. You can also type in a new value.
- On Day of Week allows you to schedule the task on one or multiple days (Sunday, Monday, etc.) of the week. Click on the days of the week to select the days you want the task scheduled.
- On Date of Month allows you to schedule the task on one or multiple days (1 - 31) of the month. Click on the dates of the month to select the dates you want the task scheduled.

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

Start Execution

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

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

End Execution

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

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

Time Zone

Select the time zone from the pull-down list. The choices are Agent, Console, and GMT.

Note: Only the Agent time zone is available with this release.

- **Agent:** The agent schedules the task execution at each destination based on the actual system time of each agent. Tasks are not necessarily run simultaneously.
- **Console:** The agent schedules the task execution simultaneously on all destinations based on the system time of the Console.
- **GMT:** The agent schedules the task execution simultaneously on all destinations based on Greenwich Mean Time.

Creating an Image Copy An image copy contains a single file (datafile, archivelog, or control file) that you can use as-is to perform recovery.

To create an image copy:

1. Select the desired object in the tree list.
2. Choose Image Copy from the Backup menu. The ImageCopy property sheet appears.

Image Copy Property Sheet: General Page

The General Page allows you to enter the requisite information for a complete image copy of a datafile, or control file. An image copy is a direct copy of the selected input file that contains no header or trailer blocks.

Image copy backups can be written only to disk. No compression is used. Restoring a file with an image copy involves substituting the image copy for the actual file.

Destination

The drive and/or full path to the directory where you store image copies. The filename of the selected input file is used by default. The default backup directory is defined by the default channel for the target database.

Channel

Use the drop-down list to select an available channel.

Note: Before performing an image copy, you must define at least one channel.

Enter or choose a tag

Optionally, you can specify a tag by entering a character string of less than 30 characters. You can also use the drop-down list box to select an existing tag.

Defining a tag gives you a symbolic way to refer to a collection of image copies or backup sets. You can also use tags to specify specific input files to a restore or switch.

Image Copy Property Sheet: Schedule Page

See *Backup Set/Image Copy Property Sheet: Schedule Page* on page 12-28 for a detailed description.

Managing Restore and Recover Operations

In addition to the automated restore and recover procedures you can perform using the Restore Wizard, you can manually restore and recover database objects.

To restore an individual tablespace or datafile:

1. Select a tablespace or datafile from the tree list.
2. Choose Restore from the Recover menu. The Restore property sheet displays.

Restore Property Sheet: General Page

The General page allows you to specify the type of restore operation performed by Backup Manager.

Select restore method

Allow Rman to choose the best available backup set.: Selected by default, this option allows Backup Manager to select automatically the most recent backup set or file copy available.

Restore to Tag: When selected, allows you to select the backup set or file copy using the files' tag. Select the desired tag using the drop-down list. If backup sets or file copies are found with the same tag, the tag with the most recent timestamp is selected.

Recover object after restore

When selected, recovery is automatically performed after the restore operation.

Restore Property Sheet: Rename Page

The Rename page allows 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 pop-up menu.

Restore Property Sheet: Channels Page

The Channels page allows you to select the channel(s) used for the current restore operation.

Choose channel type

Drop-down list allowing you to select one of the following channel types:

- Disk: Displays all available channels that write to disk.
- Tape: Displays all available channels that write to a sequential I/O device or use an OS-specific device access method.

Selected channels

List of channels selected for use by the current restore operation.

Arrow buttons

Use the left arrow button to add channels selected in the Available channels box to the list of channels in the Selected channels box.

Use the right arrow button to remove channels selected from the Selected channels box to the list of channels in the Available Channels box.

Double-clicking on a channel in either box performs the same operation the arrow keys perform—it transfers the channel from one box to the other.

Available channels

List of available channels for use by the current restore operation.

To convert a copy of a datafile into a currently used datafile:

1. Select a datafile from the tree list.
2. Choose Switch from the Recover menu. A Switch prompt appears. Click OK. The Switch Datafile property sheet appears.

Switch Property Sheet: General Page

The General page allows you to convert a copy of a datafile into a datafile used by the target database. When a switch is performed, the datafile copy is registered with the controlfile, thus making it a datafile that can be accessed by the database.

Tag

When selected, a drop-down list allows you to select the datafile copy you wish to make current via the datafiles' tag.

Datafile

When selected, allows you to select from a list of available datafile copies. The multi-column list displays the following file attributes:

- *CPDate* Date and time of the last datafile backup or image copy.
- *Tag* Tag associated with the datafile, if any.
- *Name* Actual name of the datafile.

To recover a tablespace or datafile:

1. Select a tablespace or datafile from the tree list.
2. Choose Recover from the Recover menu. The Recover property sheet appears. The Oracle8 Recovery Manager subsystem recovers the selected database object to the most recent date and time possible.

Recover Property Sheet: General Page

The General page allows you to select the channel for recovery of the selected object in the tree list. This property sheet consists of the following:

Choose channel type

Drop-down list allowing you to select one of the following channel types:

- **Disk:** Displays all available channels that read from disk.
- **Tape:** Displays all available channels that read from a sequential I/O device or use an OS-specific device access method.

Selected channels

List of channels selected for use by the current recover operation.

To remove an entry from the Selected channels list, double-click on the entry.

Available channels

List of available channels for use by the current recover operation.

To add an entry to the Selected channels list, double-click on the Available channels entry.

Displaying Datafile Errors: If any errors exist within the target database's datafiles, you can display their current recovery status by performing the following actions:

1. Select the datafile in question from the tree list.
2. Choose Error Status from the Recover menu. The Error Datafiles dialog box appears.

Error Datafiles Dialog Box The Error Datafiles dialog box appears automatically if a problem is found with one or more datafiles.

Name

The full path and name of the datafile.

Tablespace

Name of the tablespace to which the datafile belongs.

Error

Error, if any, generated by the Recovery Manager utility. See the *Oracle8 Backup and Recovery Guide* for more information.

Needs Recovery

Datafile status as to whether the datafile needs recovery. Depending on the type error, the column may be blank.

File Status

Recovery status for the datafile. In most cases, RECOVERY will be the default entry.

Maintaining the Recovery Catalog

The recovery catalog must contain current information about the target database at all times. To facilitate recovery catalog maintenance, the Catalog menu consolidates all tasks required to keep the recovery catalog up-to-date. See *Catalog Menu* on page 12-8

Database Operations

Primary database operations such as startup, shutdown, switching a logfile, adding or deleting a log group or member are consolidated in the Database menu. For more information on these functions see *Database Menu* on page 12-9.

Creating a Channel

A channel establishes a connection from the Oracle8 Recovery Manager subsystem to the target database for backup or restore operations. Multiple channels can be created to allow multiple backup sets or image copies to be read or written in parallel by a single job.

To create a new channel: Choose Create from the Channel menu. The Create Channel property sheet appears.

Channels Property Sheet: General Page

The General page allows you to create, view, or modify a channel. A channel establishes a connection from Backup Manager to the target database for backup or restore operations. Multiple channels can be created to allow multiple backup sets or file copies to be read or written in parallel by a single job.

Note: At least one channel must exist before performing backup or restore operations.

Channel Name

User-specified name of the channel.

Default Channel for this database

When checked, specifies that the current channel be used by default whenever the currently connected target database is accessed.

Channel available for all databases

When checked, makes the current channel available to any target database.

Channel Type

Drop-down list allowing you to select from one of the following device types:

- Disk: Backup sets written to disk.
- Tape: Displays all available channels that write to a sequential I/O device or use an OS-specific device access method.

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

Destination (active when Channel Type is set to Disk): Drive and path where backup sets are stored.

Channels Property Sheet: Limits Page

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 the its value. The middle number on the scale 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.

Creating Stored Scripts

The stored script capability of Oracle8 allows you to store a sequence of Recovery Manager commands within the the recovery catalog for execution at a later time. This allows you to plan, develop, and test a set of commands for backing up, restoring, and recovering the database. Each stored script relates to only one database.

To create a stored script:

1. Choose Stored Script from the Catalog menu. The Open Stored Script dialog box displays.
2. Click on the Create New button. The RMAN script dialog displays.
3. Enter the desired RMAN commands or click Import to import the text from an ASCII file. Conversely, clicking Export will export the contents of this dialog to an ASCII file for backup purposes, or use with other databases.

Backup Manager: Operating System Backup Subsystem

This backup subsystem provides you with an easy-to-use backup tool for backing up and recovering Oracle7 databases. You can also use the Operating System Backup subsystem for limited backup and recovery operations on Oracle8 databases.

Operating System Backup User Interface

After selecting Operating System Backup from the Subsystem selection dialog box, the Backup Manager window appears.

Figure 12-3 Backup Manager: Operating System Backup



Backup Manager Menu Options

Backup Manager includes four standard menus, File, View, Logfile, and Help. It also includes the Backup, Recover, and Logfile menus. For information on standard application menus, see *Application Menus* on page 7-10

Backup Menu

The Backup menu includes the following items:

Begin Online Backup

Prepares the selected tablespace for an online backup.

When you choose Begin Online Backup, two columns of the Tablespaces multi-column list indicate that fact: the Backup Status column changes to "Active," and the Online Backup Started column shows the data and time the backup was started.

The Begin Online Backup menu item is enabled when an individual tablespace is selected, that tablespace is online, and the backup status of that tablespace is not currently active.

End Online Backup

Notifies the database that the online backup is complete.

When you choose End Online Backup, two columns of the Tablespaces multi-column list indicate that fact: the Backup Status column changes to "Not Active," and the Online Backup Started column no longer contains any information for that tablespace.

The End Online Backup menu item is enabled when an individual tablespace is selected, that tablespace is offline, and that tablespace is currently being backed up.

Place Tablespace Online

Places the selected tablespace online.

The Place Tablespace Online menu item is enabled when an individual tablespace is selected, and that tablespace is offline.

Take Tablespace Offline

Places the selected tablespace online.

The Place Tablespace Offline menu item is enabled when an individual tablespace is selected, and that tablespace is online.

Backup Tablespace Wizard

Allows you to create and submit jobs for backing up the tablespaces of the database. Jobs can be saved to the job library for future execution.

Backup Control File

Displays the Backup Control File dialog box, from which you can enter the name and directory location of the backup control file for the database.

Recover Menu

The Recover menu contains the following items:

Recover

Recover the selected object in the tree list.

Recover Wizard

Starts the Recovery Wizard that guides you through the recovery processes.

Database Menu The Database menu contains the following items:

Shutdown Database

Shuts down the database.

Startup Database

Starts up the database.

Restrict Sessions

Makes the database accessible only to users with the RESTRICTED SESSION system privilege. Users already connected are not affected.

Allow All Sessions

Makes the database accessible to all users with the CREATE SESSION system privilege.

Switch Logfile

Allows you to switch redo log groups. Choosing Switch Logfile automatically enables the next redo log group to become the current redo log group.

Force Checkpoint

Forces a checkpoint. During a checkpoint, all modified database buffers are written to the appropriate datafiles.

Add Log Group

Creates a new redo log group for the database.

Drop Log Group

Allows you to remove a redo log group from the database.

The Drop Log Group menu item is enabled when an individual redo log group container is selected.

Add Log Member

Allows you to create a new member for an existing redo log group.

Drop Log Member

Allows you to remove a redo log file from a redo log group.

The Drop Log Member menu item is enabled when an individual redo log file is selected.

Database Object The database object displays the name of the target database. When selected, the database property sheet displays providing you with information regarding database status and archive mode. This property sheet contains the following information.

Database Property Sheet: Status Page

The Status page contains information about the status of the current database, database version, and any installed options. This page also allows you to change the state of the database by selecting one of the Database State options and clicking Apply.

Database State

Shutdown: Database down.

Instance Started: Instance started but not mounted.

Database Mounted: Database mounted but not open.

Database Open: Database mounted and open.

Database Version

Database version and any installed options.

Attention: If the database is not running, the message “ORACLE not available” is displayed.

Database Property Sheet: Information Page

The Information page contains displays the current state of redo log archival and the system global area.

Archive Information

Log Mode: Specifies whether the connected database is in ARCHIVELOG mode or NOARCHIVELOG mode. For information on switching database archive modes, see Oracle 7 Server Administrator's Guide.

Last SCN Archived: The last system change number that was archived. The SCN uniquely identifies the last committed database transaction.

Archive Destination: Specifies the destination where the archive log files are to be created. If you are archiving to disk, it is recommended that a dedicated disk with sufficient storage be used.

Archive Format: Specified the naming convention for the archived log files. ARC, appended with the backup set number (%S) and the backup stamp number (%T), is used to form a unique filename for the backup set.

%P Backup Piece Number—This value starts at 1 for each backup set and is incremented by 1 as each backup piece is created.

%S Backup Set Number— This value is a control file counter that is incremented for each backup set. The counter value starts at 1 and is unique for the lifetime of the control file.

%T Backup Stamp—This value indicates the number of seconds since a fixed reference date/time (currently midnight of January 1, 1988).

Automatic Archival: Allows you to specify that redo log files be automatically archived. To enable Automatic archiving upon instance startup, set the LOG_ARCHIVE_START initialization parameter to TRUE. This option is enabled when the connected database is running in ARCHIVELOG mode.

SGA Information

Database Buffers: Size of the database buffer cache (in bytes).

Fixed Size: Memory allocated to the area of the SGA that contains general information about the state of the database and the instance. No user data stored here (in bytes).

Redo Buffers: Size of the redo log buffer (in bytes).

Variable Size: Memory allocated to variable size data structures (in bytes).

Controlfile Object and Property Sheet

The controlfile object represents the controlfile of the target database. When selected from the tree-list, the control file property sheet displays. This property sheet displays all available control files used by the target database and also displays the name and location of the control file as well as its status.

Tablespace Container, Multi-column List, and Property Sheets

The Tablespaces object type container contains each tablespace in the database arranged alphabetically in a tree list. A particular tablespace can be expanded to show each of its datafiles.

When you select:

- The Tablespaces object type container, a multi-column list displays a row of summary information for each tablespace in the database.
- An individual tablespace, a multi-column list displays a row of summary information for each datafile in the tablespace.

Tablespaces Multi-column List When you select the Tablespaces object type container in the tree list, a multi-column list displays on the right. This list includes a row for each tablespace in the database. The columns of the list are:

Tablespace

Name of the tablespace.

Size (M)

Total size in megabytes of the datafiles that comprise the tablespace.

Backup Status

Backup status of the tablespace: Active (an online backup of this tablespace is currently in progress), Not Active (no online backup of this tablespace is currently in progress); or Offline.

Online Backup Started

Date and time the online backup of this tablespace started. If this column is empty, no online backup of this tablespace is currently in progress.

Tablespace Property Sheet: General Page

The Tablespace property sheet allows you to alter the status of the tablespace selected in the tree list. In addition, all datafiles contained within the tablespace are displayed.

Name

The full path and name of the selected tablespace.

Status

Indicates the current status of the selected tablespace and also allows you to change the status. The following options are available:

- **Online:** Places the selected tablespace online.
- **Offline:** Takes the selected tablespace offline. When this option is selected, a drop-down list appears to the right allowing you to select the following priorities:
 - **Normal:** Used if no error conditions exist for any of the datafiles of the tablespace.
 - **Temporary:** Used if there are error conditions for one or more datafiles of the tablespace.
 - **Immediate:** Note: You cannot take a tablespace offline immediately if the database is running in NOARCHIVELOG mode.
- **Read Only:** Makes the selected tablespace read-only to prevent further write operations on the datafiles in the tablespace.

Datafiles

A scrolling multi-column list displaying all datafiles contained within the tablespace. For each datafile, the full name and path, size (K), date and time of the last backup, whether the datafile needs to be backed up, and any datafile errors are displayed.

Redo Log Group Container, Multi-column List, and Property Sheet

The Redo Log Groups object type container contains each redo log group of the database. A particular redo log group can be expanded to show its redo log members contained within a Redo Log Members object type container.

When you select:

- The Redo Log Groups object type container, a multi-column list displays a row of summary information for each redo log group in the database.
- An individual redo log group, a property sheet for that redo log group displays.
- The Redo Log Members object type container for a particular redo log group, a multi-column list displays a row of summary information for each of its redo log members.
- An individual redo log member, a property sheet for that redo log member displays.

For information about managing redo log groups, see the *Oracle Server Concepts*, the *Oracle Server Administrator's Guide*, and the *Oracle Server SQL Reference*.

Redo Log Groups Multi-column List When you select the Redo Log Groups object type container in the tree list, a multi-column list displays on the right. This list includes a row for each redo log group in the database. The columns of the list are:

Group

Group number of the redo log group.

Sequence

Sequence number of the redo log group.

Size (K)

Space allocated in kilobytes to the redo log group.

Status

Logging status of the redo log group: Current, Active, or Inactive.

Archived?

Indicates whether the redo log group has been archived. YES or NO.

Low SCN

System Change Number for the first change recorded in the redo log group.

Time of Low SCN

Date and time of the low SCN.

Redo Log Members Multi-column List

When you select the Redo Log Members object type container in the tree list, a multi-column list displays on the right. This list includes a row for each redo log member in the redo log group. The columns of the list are:

Member Filename

Filename and directory path for the redo log member.

Group

Group number of the redo log group to which this redo log member belongs.

File Status

Status of this redo log member: In Use, Invalid (the file is inaccessible), Stale (the contents of the file are incomplete), or Deleted (the file is no longer in use).

Group Status

Status of the redo log group to which this redo log member belongs: Current, Active, or Inactive.

Archived?

Indicates whether this redo log member has been archived: YES or NO.

Redo Log Group Property Sheet

When you select a particular redo log group in the tree list, a property sheet for this redo log group displays on the right. The items of this property sheet are:

Group #

Displays the redo log group number for this redo log group.

Use Existing File(s)

Click this button to allow Oracle Enterprise Manager to use existing file(s).

New File(s) Size

Click this button in order to specify the size of the new Redo Log Group file(s).

KM Bytes

If you click the New File(s) Size button, enter the new file(s) size in the Bytes box and click K (default) to specify that the new file(s) size is in kilobytes, or M to specify that the new file(s) size is in megabytes.

Current Members

Displays the filenames and directory paths for the redo log members of this redo log group.

Remove

Disabled in this property sheet.

New Member

Enter a new member for this redo log group in the New Member box.

....

Click the Browse button to display the Open dialog box, from which you can select the redo log file you want to enter in the New Member box.

Add

Adds a new redo log member to this redo log group.

The Add button is enabled when you enter a new redo log member in the New Member box.

Redo Log Member Property Sheet

When you select a particular redo log member in the tree list, a property sheet for this redo log member displays on the right. The items of this property sheet are:

Filename

Enter the new filename of the redo log group member.

Group

Displays the redo log group to which the redo log member belongs.

Selected Backup Manager Tasks

This section covers the following topics:

- *Shutting Down the Database*
- *Starting Up the Database*
- *Backing Up Tablespaces of the Database*
- *Adding a New Redo Log Group*
- *Adding a New Member to an Existing Redo Log Group*
- *Performing Recovery*

Shutting Down the Database

To shut down the database, choose Shutdown Database from the Backup menu. The Shutdown property sheet appears.

Attention: Before shutting down a release 7.1 or later database, you must connect as SYSDBA or SYSOPER. Before shutting down a release 7.0 database, you must connect as INTERNAL.

The items of the Shutdown property sheet are described below:

Shutdown Options

Normal: Shuts down the database in normal mode.

Immediate: Shuts down the database in immediate mode. (This is the default.)

Abort: Shuts down the database in the abort mode.

Transactional (Oracle8): Provides a specified length of time (Timeout) in which to complete transactions before shutdown occurs. If the Timeout field is left blank, shutdown does not occur until the last database transaction is complete.

Shut Down

Click this button to shut down the database in the mode you have selected.

Other applications create separate connections when you start them. When performing a shutdown in normal mode, remember to close these windows, or the shutdown will not complete.

Starting Up the Database

To start up the database, choose Startup Database from the Backup menu. The Startup property sheet appears.

Before starting up a release 7.1 or later database, you must connect as SYSDBA or SYSOPER. Before starting up a release 7.0 database, you must connect as INTERNAL.

The items of the Startup property sheet are described below:

Startup Options

Force: Forces an instance to start regardless of the operating circumstances.

Restrict: Starts an instance in restricted mode. Connections are limited to those users who have been granted the RESTRICTED SESSION system privilege.

Parallel: Allows multiple instances to access a single database concurrently.

No Mount: Starts up the instance and mounts the database. The database is accessible only to database administrators.

Mount: Mounts a database upon instance startup. The database is accessible only to database administrators.

Mount and Open: Starts up the instance, and mounts and opens the database.

Don't Use Stored Parameters (checkbox)

This checkbox only appears when initialization parameters are stored in the repository. If visible and not checked, a new INIT.ORA file will be created.

Parameter File

Name of the parameter file used to start the instance.

If you do not specify a parameter file, Oracle Enterprise Manager looks for the parameter file in the default location on your local machine. For information about the default location for the parameter file, see your operating system-specific Oracle documentation.

Browse: Displays the standard Open dialog box, allowing you to locate the parameter file.

Startup

Click this button to start up the database.

Backing Up Tablespaces of the Database

To create a backup job script in order to back up one or more tablespaces of the database, choose Backup Tablespace Wizard from the Backup menu. The Backup Tablespace Wizard displays.

It is best to create a separate backup job script for each database you plan to back up, as it is highly unlikely that the data structures of any two databases would match.

The Backup Tablespace Wizard includes four pages that you complete in order to create a backup job script, which you can then run from the Job subsystem of the Console. For more information about the Backup Wizard, see Backup Manager online help.

Performing Recovery

The Recovery Wizard automates the recovery process. The Introduction page of the Recovery Wizard allows you to select from the following types of recovery operations:

- Database
- Tablespaces
- Datafiles

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

Depending on the status of the target database (ARCHIVELOG, NOARCHIVELOG, mounted and open, or mounted only), some options will be disabled.

Refer to Oracle Backup Manager online help for a detailed information.

Adding a New Redo Log Group To create a new redo log group, choose Add Log Group from the Logfile menu. The Create Redo Log Group property sheet appears.

Group #

Displays the redo log group number that Oracle Enterprise Manager has automatically generated for this new redo log group. (For example, if redo log groups 1 and 2 already exist for the database, 3 appears by default in the Group # box.)

You can change this default group number if you want.

Use Existing File(s)

Click this button to allow Oracle Enterprise Manager to reuse existing file(s).

New File(s) Size

Click this button in order to specify the size of the new Redo Log Group file(s).

KM Bytes

If you click the New File(s) Size button, enter the new file(s) size in the Bytes box and click K (default) to specify that the new file(s) size is in kilobytes, or M to specify that the new file(s) size is in megabytes.

Current Members

Displays the current members of this redo log group.

The Current Members box is empty when the Create Redo Log Group property sheet first displays. After you add a new member to this redo log group, the new member then appears in this box.

Remove

Click this button to remove the redo log group member that is selected in the Current Members box.

New Member

Enter a new member for this redo log group in the New member box.

....

Click this button to display the Open dialog box, from which you can select the redo log file you want to enter in the New Member box.

Add

Click this button to add the new member for this redo log group that you entered in the New Member box.

Create

Click this button to create the new redo log group.

Adding a New Member to an Existing Redo Log Group To create a new member for an existing redo log group, choose Add Log Member from the Logfile menu. The Create Redo Log Member property sheet appears.

Filename

Enter the filename of the new member of the redo log group shown in the Group box.

Group

From this box, select the redo log group to which you want to add a new member.

Use Existing File

Click this button to specify that Oracle Enterprise Manager reuse an existing file.

New File

Click this button to specify that Oracle Enterprise Manager create a new file.

Backup Manager: Enterprise Backup Utility Subsystem

The following information pertains to Enterprise Backup Utility installation and configuration.

The user interface for Enterprise Backup Utility subsystem is similar to Oracle8 Recovery Manager and Operating System Backup subsystems discussed earlier in this chapter. Hence, many of the operational characteristics are identical. For specific information on Enterprise Backup Utility installation, features, and functionality, see the Enterprise Backup Utility documentation.

Figure 12–4 *Enterprise Backup Utility*



Attention: The current version of Enterprise Backup Utility is not compatible with Oracle8 databases. This backup subsystem is comes standard with Oracle Release 7.3.4. It can also ordered separately.

You will also find EBU 2.1.0.1 in the 7.3.3 bundle on UNIX platforms.

When installation is complete, there will be an 'obackup' directory in your oracle home.

Creating an Enterprise Backup Utility Catalog

A catalog database is analogous to an Oracle 8 recovery catalog: it is a set of tables on a separate database.

1. Make sure the SQLNET connect string is setup in your TNSNAMES.ORA and verify you can connect to this catalog database from your target database and from Oracle Enterprise Manger. You can use sqlplus system/manager@connectstring. You can also use Server Manger Line Mode.
2. Create the Enterprise Backup Utility catalog schema with OBackup on the target database system.

Attention: EBU_HOME replaces OBK_HOME beginning with version 2.2.

```
setenv OBK_HOME $ORACLE_HOME/obackup
```

Where ORACLE_HOME is the home you installed Enterprise Backup Utility when you ran oraInst.

3. Add the following lines to your ~/.login (or target database system):

```
setenv OBK_HOME $ORACLE_HOME/obackup
set path=($path $OBK_HOME/bin)
setenv LD_LIBRARY_PATH ${LD_LIBRARY_PATH}:$OBK_HOME/lib
```

4. Enter the following at the OS command line (or target database system):

```
source ~/.login
cd $OBK_HOME/lib
cd $ORACLE_HOME/obackup/bin
obackup -upgrade
```

Attention: Beginning with version 2.2, you can specify 'ebu' in place of 'obackup', although 'obackup' may still be used for backwards compatibility. If you specify 'obackup', OBK_* environment variables are used. If you specify 'ebu', EBU_* environment variables are used.

Note: Specifying an Enterprise Backup Utility upgrade (obackup -upgrade) installs the tables if not already present.

You will be prompted for the catalog connect string (username/password@connect, which you can enter one at a time or use the above syntax, just like sqlplus).

If you created your own catalog database, enter something like obk/obk@your_connect_string. It will then ask you for the SYS password so it can create the user you entered. Once it creates the user, it will connect to the catalog db as that user and create the tables and other requisite schema objects.

For more information on installation, see the Enterprise Backup Utility Administration Guide.

Configuring the PC for Enterprise Backup Utility

1. From the OEM console, click on the target database where you just installed Enterprise Backup Utility and then select Backup Manager from the Tools menu or from the Tools palette.
2. Select Enterprise Backup Utility from the Subsystem Selection dialog box.
3. Then from the Settings menu, select Catalog Connect String and enter the same catalog connect string you gave to Enterprise Backup Utility. This connect string should be configured in the TNSNAMES.ORA file on this PC the same way you configured it on the target database.
4. Now you can hit the Backup toolbar button. The first time you do this, it will detect that the target database is not registered with EBU, and it will pop up a Register dialog box.
5. Enter the target database's Parameter File (init.ora), or let it default (for UNIX systems), then press OK.
6. Click on the Jobs node in the tree on the left side and check the job status of the REGISTER job.
7. Once it completes successfully, you can click on the Backup toolbar button again and this time it should come up with the Backup Wizard.

Using the SQL Worksheet

This chapter describes how to use the SQL Worksheet, which is part of Oracle Enterprise Manager. It explains how to:

- [Overview of the SQL Worksheet](#)
- [Executing Commands in the SQL Worksheet](#)
- [Using the Command History](#)
- [Saving Your Work](#)
- [Running Scripts from the SQL Worksheet](#)

Overview of the SQL Worksheet

The SQL Worksheet allows you to enter, edit, and execute SQL and PL/SQL code. You can also execute the Server Manager DBA commands described in Appendix B, DBA Command Reference. You can also run scripts from the SQL Worksheet. The SQL Worksheet maintains a history of the commands you have issued, so you can easily retrieve and re-execute previous commands.

Starting a SQL Worksheet

You can start a SQL Worksheet from the Enterprise Manager Console window by choosing SQL Worksheet from the Tools menu. The new worksheet is connected to the database you have selected in the Navigator tree list or in the Map window.

If you start the SQL Worksheet from the Windows Program Manager, a database Login dialog box appears.

Note: If you connect as INTERNAL to a database that is currently shut down, the following error message appears in the SQL Worksheet output pane upon startup: `Cannot obtain ORACLE version number.` This message is caused by a minor system interface discrepancy and will not affect SQL Worksheet operation or the database connection.

Using the SQL Worksheet

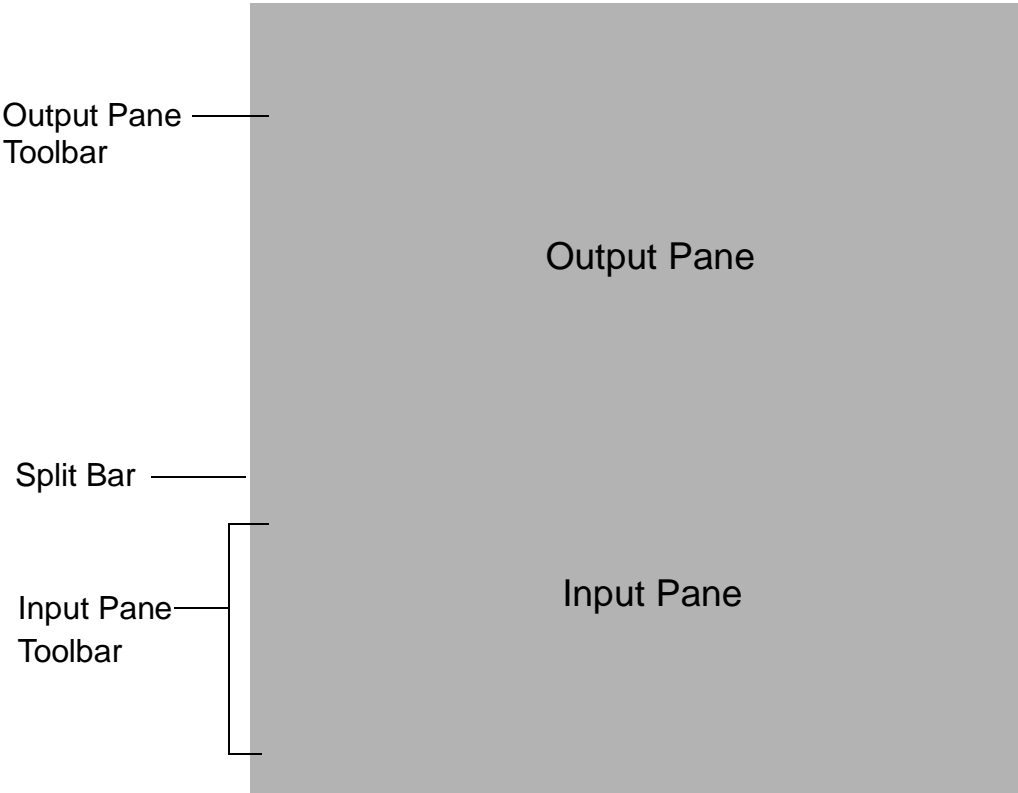
A SQL Worksheet window (Figure 13–1) consists of the following elements:

- Input pane
- Output pane
- Split bar
- Input pane toolbar
 - New Worksheet
 - Open worksheet
 - Save worksheet as
 - Execute
 - Command history
 - Previous command
 - Next command

- Output pane toolbar
- Worksheet pull-down menu

The following figure illustrates a SQL Worksheet window.

Figure 13-1 SQL Worksheet Window



The SQL Worksheet window is described below:

Output Pane

Upper window pane. Displays recently executed commands and output. SQL Worksheet stores the last 32 KB of output for display in the output pane. The actual amount stored varies depending upon the amount of system memory available.

Input Pane

Lower window pane. Enter and edit commands in the input pane.

Split Bar

Adjusts the relative sizes of the panes. Drag the split bar vertically to resize the panes.

Output Pane Toolbar

Icon provides access to the Save Output As menu option.

Input Pane Toolbar

Provides access to the following operations:

Creates a new worksheet.

Imports the contents of an existing SQL script.

Saves the contents of the input pane to a file

Executes the command in the input pane. See *Executing Commands in the SQL Worksheet* on page 13-8 for information about executing commands in a SQL Worksheet.

Displays the Command History dialog box, which displays previously executed commands. See *Using the Command History* on page 13-12 for information about using the command history.

File Menu

The File menu items are described below:

Change Database Connection

Displays the Connect dialog box to connect to an instance. To connect to multiple databases concurrently, you need to launch the application multiple times from the console.

New

Clears the contents of the input pane. You are prompted to save your existing work.

Open

Displays the Open Worksheet dialog box to allow you to open an existing SQL or ASCII text file. The contents of the file appear in the input pane.

Save

Saves the contents of the input pane to the current file.

Save As

Saves the contents of the input pane to a file other than the one that is currently open.

Print

Prints the contents of the input pane.

Save Output As

Saves the contents of the output pane to a file.

Print Output

Prints the contents of the output pane.

Print Setup

Displays the Print Setup dialog box allowing you to set print parameters and choose a local printer.

Exit

Quits the SQL Worksheet.

Edit Menu

The Edit menu items are described below:

Cut

Cuts selected text.

Copy

Copies selected text.

Paste

Pastes selected text at the location of the cursor.

Select All

Selects all the text in a document or selects all objects in a list.

Clear All

Clears all text in a window pane.

Worksheet Menu

The Worksheet menu items are described below:

Execute

Executes the commands in the input pane. Equivalent to the Execute button or the Enter key.

See *Executing Commands in the SQL Worksheet* on page 13-8 for information about executing commands in a SQL Worksheet.

Run Script

Allows you to select and execute a script. See *Running Scripts from the SQL Worksheet* on page 13-15 for information about running scripts from a SQL Worksheet.

Spool

Writes output to a specified file. See *Saving Your Work* on page 13-14 for information about spooling the output of a SQL Worksheet.

Command History

Displays the Command History dialog box, which displays previously executed commands. Equivalent to the Command History button. See *Using the Command History* on page 13-12 for information about using the command history.

Previous Command

Retrieves the previous command from the command history and enters it in the input pane. The retrieved command replaces the entire contents of the input pane.

You can execute Previous Command repeatedly to work back through the command history. Previous Command is unavailable if you have currently retrieved the oldest command from the command history.

Next Command

Retrieves the next command from the command history and enters it in the input pane. The retrieved command replaces the entire contents of the input pane.

Next Command is unavailable if you have currently retrieved the most recent command from the command history.

Suggestion: Use the keyboard shortcuts to quickly access and execute commands from the command history.

Executing Commands in the SQL Worksheet

Enterprise Manager executes commands you have entered in the input pane of a SQL Worksheet.

Entering Commands in the Input Pane

You can enter SQL, PL/SQL, and Server Manager DBA commands in the input pane of a SQL Worksheet. A command entered in the input pane can contain tabs or carriage returns.

All standard text editing capabilities are available in the input pane of a SQL Worksheet. You can edit text in the input pane using the mouse, keyboard functions, and the Edit menu. In addition, you can select text from the output pane and paste it into the input pane.

Executing Commands in the Input Pane

To execute the command you have entered in the input pane, click the Execute icon in the input pane toolbar. You can also execute the command by choosing Execute from the Worksheet menu.

When you execute a command, the command and its output appear in the output pane. However, if you set the parameter ECHO to off, only the output of the command is displayed. For information about setting SQL Worksheet parameters, see the SET command in Appendix B, DBA Command Reference.

The following figure illustrates executing a SQL statement in a worksheet.

Figure 13–2 *Executing a Command in the SQL Worksheet*



After you execute a command, the command remains in the input pane, but is selected, as shown in [Figure 13–2, "Executing a Command in the SQL Worksheet"](#). You can then type in the next command, and it replaces the previous command.

Executing Server Manager DBA Commands

In a SQL Worksheet you can execute the Server Manager DBA commands described in [Appendix B, "DBA Command Reference"](#). Some of these commands behave differently when executed in a SQL Worksheet than when executed in line mode. These commands are summarized below:

CONNECT

Displays the Connect dialog box when you issue the CONNECT command without specifying complete connect information. For information about connecting to an instance of a database, see *Connecting to an Instance* on page 1-23.

RECOVER

If recovery is needed, displays the Recover dialog box when you issue the RECOVER command without arguments. Otherwise, an error message is generated.

EXIT

Closes the worksheet.

Note: The line mode command HOST is not available in the SQL Worksheet.

Connecting from a Worksheet

When you start a SQL Worksheet from the Enterprise Manager console, you are automatically connected to the same database as the console. You can change this connection by issuing the CONNECT command from the input pane of the worksheet. The title of the SQL Worksheet window also changes to reflect your new connection.

If you execute the CONNECT command from the input pane without any arguments or select Change Database Connection from the File menu, the Connect dialog box appears. You can use the Connect dialog box to specify a new connection for that worksheet. For information about connecting to an instance of a database, see *Connecting to an Instance* on page 1-23. You can also specify the connection using arguments in the CONNECT command. For a description of the CONNECT command and its syntax, see [Appendix B, "DBA Command Reference"](#).

Performing Recovery from the SQL Worksheet

When you issue the `RECOVER` command from a SQL Worksheet, the Recovery dialog box appears, if recovery is needed. The following figure illustrates the Recovery dialog box.

Figure 13–3 *Recovery Dialog Box*



From the Recovery dialog box, you can apply the next log file, apply all the log files, or cancel recovery. You can specify the name of the log file to apply or apply the default.

Using the Command History

A SQL Worksheet maintains a history of the last 20 command executions you have issued in that worksheet. To bring up the Command History dialog box, click Command History or choose Command History from the Worksheet menu.

Note: Command executions larger than 50K are not recorded in the Command History.

The Command History dialog box contains a scrolling list of the commands in the worksheet's command history. The following figure illustrates the Command History dialog box.

Figure 13–4 *Command History Dialog Box*



The Command History dialog box is described below:

Command List

Displays an abbreviated list of the last 20 command executions. Explicit SQL commands associated with each command execution are displayed in the Command Text pane.

Command Text Pane

Displays all SQL commands associated with the selected entry in the Command List.

Close

Closes the dialog box without retrieving any text.

Get

Retrieves the selected item in the Command List command. See *Retrieving a Selection from the Command History* on page 13-13 for more information about retrieving commands.

Help

Displays help information for the Command History dialog box.

Retrieving a Selection from the Command History

To retrieve a command from the Command History dialog box, select the desired command execution entry in the Command List. All SQL commands associated with this entry are displayed in the Command Text pane. Click Get. The contents of the Command Text pane appear in the input pane replacing any existing text.

You can also drag and drop commands directly into the input pane by selecting an entry in the Command List and dropping it into the input pane. The text associated with the entry appears in the input pane.

To retrieve only a portion of a command execution from the Command Text pane, select the text you wish to retrieve and press Ctrl-C (copy). Move the cursor to the worksheet's input pane, place the cursor at the position where you want the text inserted and press Ctrl-V (paste).

If no text is selected in the input pane, the copied text is inserted at the cursor position. If you have selected text in the input pane, then the retrieved text replaces the selected text.

After you have retrieved text from the command history, you can edit or execute the command, as desired.

Retrieving a Command Using the Worksheet Menu

You can also retrieve command executions without using the Command History dialog box. The Previous Command and Next Command menu items and input toolbar icons retrieve commands from the command history and enter them in the input pane, completely replacing the existing contents of the input pane. For more information on the Previous Command and Next Command menu items, see *Worksheet Menu* on page 13-6.

Saving Your Work

The Spool menu item, Save Output As icon, and Save Worksheet As icon allow you to save the work you have performed in a SQL Worksheet. Each of these Worksheet menu commands displays the standard file selection dialog box, which you can use to specify the file in which to save your work.

The Spool menu item allows you to dynamically save all output of your worksheet session to a specified file. SQL Worksheet continues to save the output of your worksheet until you choose Stop Spooling from the menu or when you exit the SQL Worksheet.

The Save Output As allows you to write the current contents of your worksheet's output pane to a specified file. The output pane normally contains the last 32 KB of output and the last 100 lines of text. The actual amount may vary according to the amount of available system memory. Any text that has scrolled out of the output pane is also written to the file.

The Save Worksheet As menu item and icon (located in the input pane toolbar) writes the contents of the worksheet's input pane to a specified file.

Note: The SQL Worksheet appends a file extension, such as .log or .txt, to the name of each file created by the Spool, Save Output As, and Save Worksheet As commands.

Running Scripts from the SQL Worksheet

From a SQL Worksheet, you can run any script that contains SQL, PL/SQL, or Server Manager DBA commands. There are four ways to run scripts from a SQL Worksheet:

- Clicking the Open Worksheet Icon from the input pane toolbar.
- Choosing Run Script from the Worksheet menu.
- Dragging and dropping a SQL script from the Windows NT file manager into the worksheet's input pane.
- Executing the Server Manager DBA @ command from the input pane.

Clicking the Open Worksheet icon in the input pane toolbar or selecting the Open menu item displays the standard file selection dialog box. After you select the file and click OK, the content of the SQL script appears in the worksheet's input pane. Click the Execute icon or select the Execute menu item to run the script.

Selecting the Run Script menu item also displays the Run Script dialog box. Select the desired script and click OK. The Server Manager DBA @ command appears with the full path and filename of the selected script in the input pane. The script is run automatically.

To drag and drop scripts from the Windows file manager, select the desired script in the file manager and drop it into the input pane. The content of the script appears in the input pane. Click the Execute icon to run the script.

You can also run a script from the input pane of a worksheet by entering an @ command followed by the script name. The following figure illustrates running a script from the input pane of a worksheet.

Note: In a script you can run other scripts. In the original script you can include the line @@*second_scriptname* to run another script called *second_scriptname*. The @@ command indicates that the second script is in the same directory as the original script.

Managing and Moving Data

This chapter describes how to use Oracle Data Manager wizards to transfer data to and from an Oracle database. It also includes a discussion of the Data Manager tasks that can be submitted with the Job Scheduling system. The topics in this chapter are:

- [Starting Data Manager](#)
- [Exporting Data to a File](#)
- [Importing Data from a File](#)
- [Loading Data from an External File](#)
- [Data Manager Tasks](#)

This chapter assumes that you have read [Chapter 7, "Overview of the Database Tools"](#), and that you are familiar with the interface elements of the database tools.

For more information on the tasks Data Manager performs, see the sections on the Export, Import and SQL*Loader utilities in *Oracle Server Utilities*. For information about standard elements of your operating system, see your operating system-specific documentation.

Starting Data Manager

To start Data Manager, select a database object from the Navigator tree list. Then click on the Data Manager icon in the Applications Launch Palette, or choose Applications and then Oracle Data Manager from the Tools menu.

- To use Data Manager's remote import capability, the Oracle Enterprise Manager Console must be running before starting up Data Manager. Data Manager uses the Console job system to perform remote import operations.
- If you have already connected to a database, the Oracle Data Manager main window displays.
- If a valid database connection has not been made, the Connect dialog box displays. This dialog box also appears when you launch Data Manager in standalone mode. See **Connecting to an Instance** on page 1-23.

Note: You can change the database connection with the Change Database Connection option in the File menu. For information, see **Application Menus** on page 7-10.

Data Manager Window

After Data Manager has successfully connected to a database, the Users container displays in a tree list on the left side of the Data Manager window. The Users container is located in the database that Data Manager is connected to. An individual user container can be expanded to show the tables that the user owns.

The display on the right side of the window is determined by the object selected on the left side of the screen. The right side may contain a multi-column list, property sheet, or other information.

- If the User container is selected, a multi-column list of all users displays.
- If a specific user is selected, the user property sheet displays.
- If a specific table is selected, the table property sheet displays.
- If a table partition is selected, the table property sheet displays.

For general information, see:

- **<Emphasis>**Application Objects on page 7-3.
- **<Emphasis>**Multi-Column Lists on page 7-14
- **<Emphasis>**Dialog Boxes on page 7-15.
- **<Emphasis>**Property Sheets on page 7-17.

An example of a Data Manager window is shown in [Figure 14-1, "Oracle Data Manager"](#).

Figure 14–1 Oracle Data Manager



Data Manager Menu

Oracle Data Manager includes the standard menus, File, View, Log, and Help, plus the Data menu. For information on the standard menus, see <Emphasis>Application Menus on page 7-10.

The Data menu includes the Export, Import, and Load wizard options for:

- *Exporting Data to a File* which transfers data from an Oracle database to an Oracle binary operating system file.
- *Importing Data from a File* which transfers exported data in a file to an Oracle database.
- *Loading Data from an External File* which loads operating system files to an Oracle database.

Context-sensitive menus may also be active when you press the right mouse button to select a specific object from the navigator or the multi-column list. This feature provides quick access to a subset of the menu options provided in the menu bars.

Data Manager Icons

The objects in the tree list are identified by various icons. In the listing:

- A folder icon identifies an object type container.
- A single person icon depicts a user.
- A grid icon signifies a table.
- A broken grid signifies a table partition.

Exporting Data to a File

Data Manager Export is used to transfer data from an Oracle database to an operating system file in Oracle-binary format. Files in this proprietary format can only be read by using the Import component of Data Manager. Export files can be used to transfer data between databases or used as backups.

To export data from a database with the Data Manager Export wizard:

- Select the objects you want to export in the tree displayed on the left side of the Data Manager window. You can also add or remove objects while using the Export wizard.
- From the Data menu, select the Export option to start the Export wizard.
- Complete the pages of the Export wizard to define how the data is exported. In the pages of the wizard, you can specify options such as:
 - The associated objects you want to export
 - Buffer and record length sizes
 - Increment and statistics types

See the Data Manager Export online help for information on the individual pages of the Export wizard.

Remote File Option

This option specifies that the file is on the remote machine that contains the connected database. This option uses the Job Scheduling system and is only available when the Console is running.

BUFFER and RECORDLENGTH Defaults

When running the Export wizard, note that the default values for certain parameters are operating-system specific. Specifically, the BUFFER parameter

defaults to a value of 4K under Windows NT and the RECORDLENGTH parameter defaults to 2K.

Importing Data from a File

Data Manager Import is used to transfer data from an export file to an Oracle database. Import can only read data that has been created with the Export component of Data Manager.

To import data to a database with the Data Manager Import wizard:

- From the Data menu, select the Import option to start the Import wizard.
- Complete the pages of the Import wizard to define how the data is imported. In the pages of the wizard, you can specify options such as:
 - The associated objects you want to import
For user-level imports, expand the available user tree list of the Importable Objects and drag the object to the Selected Objects/Available Objects tree list.
 - Buffer and record length sizes
 - Increment type

See the Data Manager Import online help for information on the individual pages of the Import wizard.

Remote File Option

This option specifies that the file is on the remote machine that contains the connected database. This option uses the Job Scheduling system and is only available when the Console is running.

Note: When the .DMP file resides on a remote machine, Data Manager uses the Console job and event system to retrieve the file before displaying the data in through the Import Wizard. The Remote Import page of the Import Wizard provides a status line at the top of the page indicating the progress of the data retrieval job. There are three conditions that can be displayed: Job Submitted, Job Started, and Job Completed. Important: the job must complete successfully before proceeding with the import operation.

BUFFER and RECORDLENGTH Defaults

When running the Data Manager, note that the default values for certain parameters are operating-system specific. Specifically, the BUFFER parameter defaults to a

value of 4K under Windows NT and the RECORDLENGTH parameter defaults to 2K.

Loading Data from an External File

Data Manager Load is used to transfer data to an Oracle database from files that have a different format from export files. Often these are delimited text files. Using a control file, you specify how the data is stored externally and how it is to be loaded into the database.

To load data into a database with the Data Manager Load wizard:

- From the Data menu, select the Load option to start the Load wizard.
- Complete the pages of the Load wizard to define how the data is loaded. In the pages of the wizard, you can specify options such as:
 - The control, data, and log, files. **Important:** You must specify complete paths when entering filenames.
 - The records to load, records to skip, and rows per commit

See the Data Manager Load online help for information on the individual pages of the Load wizard.

Note: You must supply a control filename to load data. For information on control files, see the *Oracle Server Utilities Guide*.

Remote File Option

This option specifies that the file is on the remote machine that contains the connected database. This option uses the Job Scheduling system and is only available when the Console is running.

Data Manager Tasks

The Export, Import, and Load tasks are set up and submitted with the Oracle Enterprise Manager Job Scheduling system and executed by Oracle Enterprise Manager Intelligent Agents residing on managed hosts.

The Data Manager tasks are selected in the Task page of the Create Job property sheet. The parameters are defined in the Parameters page of the Create Job property sheet. These parameters are similar to the options available when using the Data Manager wizards. For information on the specific parameters for the Export, Import, and Load job tasks, see the online help for the Data Manager tasks.

For information on using the Job Scheduling system, see [Chapter 4, "Job Scheduling"](#). For information on specific job tasks, see the online help for Predefined Job Tasks.

Managing Software Distribution

Oracle Software Manager, which is bundled with Oracle Enterprise Manager, provides key services that allow server software to be managed from a central location. This product addresses the problems faced by application software administrators by offering:

- An integrated toolset for managing software in a distributed fashion
- A next generation architecture that offers network scalability through an advanced client-agent-server paradigm
- An easy-to-use graphical user interface that allows administrators to easily run software configuration jobs and track assets across the entire network

From the Software Manager interface, software packages can be created, distributed and installed on servers spread over Local and Wide area networks.

Software Manager uses the Enterprise Manager discovery cache to recognize nodes on the network and is integrated with the Enterprise Manager Job system, allowing you to submit the software distribution tasks to Oracle Enterprise Manager intelligent agents residing on selected hosts.

This chapter provides an overview of Oracle Software Manager and a conceptual look at software distribution.

Software Manager Introduction

Oracle Software Manager supports

- Windows NT and Windows 95 as its distribution host where packages are created.
- distribution of creation of packages from your hard drive or a shared drive where you CD-ROM is mounted for both 16 bit and 32 bit products as well as UNIX products.
- UNIX and NT servers as target hosts where the Intelligent Agent and Oracle Installer are running and packages can be created, distributed or installed.
- Installing and distributing OSM-compliant products installable with Oracle Installer 32-bit 3.2.2.1 and higher, Installer 3.1.4.1.5 (16 bit) or Installer 4.0.x and higher on UNIX. Note that when distributing/installing 16 bit packages, you must have a 16-bit Oracle home (default is ORAWIN) available, with the Installer 3.1.4.1.5 installed. See the Enterprise Manager README for a complete compatibility matrix, OSM Console vs. Intelligent Agent versions, as well as specific limitations that apply to the version of OSM you are running.

Note: Before using Software Manager, you must set the environment variable `OSM_PACKAGE_DIR`. The software packages will be stored in this location on agent machines. Depending on the number of packages you create on a host, you may need significant disk space, and you should make sure that the host has ample storage space. If you do not configure this variable, the default will be `%ORACLE_HOME%\NET80\AGENT\PACKAGES` on Windows NT or `%ORACLE_HOME%/NETWORK/agent/packages` on UNIX.

Note: Oracle File Packager can be used to guide you through creating a staging area, and it generates the required installations scripts. These stage areas can be used by Oracle Software Manager to create Releases and Packages that can be distributed via Oracle Software Manager. Please refer to the Oracle File Packager online help for more information.

The Software Manager Interface

The first time you launch the Oracle Software Manager application, it displays the *Navigator window* and the *Catalog window*.

See [Figure 15-1, "Oracle Software Manager Application"](#) for an illustration of the application screen.

Figure 15-1 Oracle Software Manager Application



You can minimize, resize, or reposition these windows and you can open multiple copies of the same type of window. You can open additional Catalog windows by selecting “Another Catalog Window” or “Another Navigator Window” from the Window menu.

Use **Save Layout Now** or **Save Layout on Exit** to preserve the state of your application windows from one session to another.

Navigator Window

The Navigator window is divided into a left and right pane. After you have refreshed, the Navigator shows all objects currently managed by Oracle Software Manager.

The left pane of the Navigator window shows the network tree with expandable contents displayed at the summary level.

Table 15–1 The Navigator Objects

Objects	Functions
World	Displays all objects known to Oracle Software Manager
Hosts	Displays all hosts known to Oracle Software Manager. Hosts that are illuminated in green are up and running; you can connect to these hosts with Oracle Software Manager. Hosts that are grey are currently unavailable. Installed Products, Releases, and Packages are displayed under their individual hosts. You may need to refresh before viewing the hosts. Each host has an Oracle Homes Directory. The Installed Products are under specific Oracle Home directories.
Software Packages	Shows all software packages that are available for installation or distribution in the Oracle Software Manager network.
Releases	Shows all of the releases that have been created through the Oracle Software Manager application and where they are housed in the network.
Oracle Homes	Shows all Oracle Homes on a per-machine basis in the Oracle Software Manager application. It shows all products installed on all homes. NT hosts show NT and Windows 3.1 homes.
Installed Products	Displays all products installed on all homes. Installed Products are assigned to homes which are assigned to Hosts.
Host Folders	Displays customized groupings of hosts, created using the Create Host Folder command.

Creating a Host Folder

The Host folders object in the Navigator window stores host folders that you create using the Create Host Folder command on the Edit menu. This enables you to group hosts in any way you choose. For example, you might wish to group hosts by region.

1. On the Edit menu, click Create Host Folder. The Host Folder Properties dialog box appears.
2. Type a name and description for the folder.
3. Select one or more hosts in the Available Hosts list, then click Add.
4. Click OK. The new host folder(s) appear inside the Host Folders object in the Navigator window. You can drag and drop the icon for any host in the Navigator window to the new folder. See [Figure 15-2, "Host Folder Objects"](#) for an illustration of the host folder objects.

Figure 15–2 Host Folder Objects

Viewing the Oracle Home Properties

To view the oracle home properties, double-click the home icon for the home which you want to view. The Oracle Home Properties dialog box appears.

Table 15–2 Oracle Home Properties

Fields	Function
Directory	The name of the directory
Language	The language used. For example: American
Operating System	NT hosts show NT and Windows 3.1 homes.
Status	The status distinguishes an active home (home on which the agent is running) from a non-active home.

The Oracle Homes Directory shows all Oracle Homes on a per-machine basis in the Oracle Software Manager application. It shows all products installed on all homes. NT hosts show NT and Windows 3.1 homes.

Viewing the Product Properties

To view its properties, double-click the product which you want to view. The Product Properties dialog box appears.

Table 15–3 Product Properties

Options	Function
Name	The name of the product.
Operating System	The operating system for which the product is intended.
Size	The size of the product, in bytes.
Version	The internal version number of the product.
Shareable	A sharable product may be installed on a shared file system and accessed by multiple users via the Oracle Client Software Manager (OCSM).
Description	Description of the product.

Viewing the Release Properties

On the Products menu, click Show Releases. The Releases Catalog window appears. Click the release you want to review. On the View menu, click Properties. The Release Properties information box appears, open to the Properties page.

Click the appropriate tab to view the Properties and the Products pages of this dialog box.

The Properties page displays the properties of the selected software release. These properties identify the listing for this software release in the Releases Catalog window.

The Products page lists all the products in a specified software release, displaying the name, operating system, size and version number of each.

Table 15–4 Properties Page

Options	Functions
Name	The name given to the software release when it was first created.
Description	A description of the software release's contents.
Host	Your hostname.
Location	The location where the parent files for this release reside.
Media Type	The type of medium upon which the parent software files are stored, if known.
Part Number	The part number for the parent software, if known.

Table 15–5 Products Page

Options	Functions
Product Name	The name of the product.
Operating System	The operating system for which the product is intended.
Size	The size of the product, in bytes.
Version	The internal version number of the product.

Catalog Windows

The *Catalog* shows all *releases* and *software packages* known to Oracle Software Manager.

Table 15–6 *Catalog Views*

View	Function
Releases view	Shows all of the releases that have been created through the Oracle Software Manager application and where they are housed in the network.
Software Packages view	Shows all software packages that are available for installation or distribution in the Oracle Software Manager network

A *release* is a collection of Oracle software products distributed as a unit. Releases may consist of several operating-system-specific collections.

A *software package* is a user-defined collection of products for a single operating system, intended for distribution across the Oracle Enterprise Manager network.

This Catalog View section describes the following procedures

- [Choosing the Catalog View](#) on page 15-10
- [Setting Preferences](#) on page 15-10

Choosing the Catalog View

When Oracle Software Manager is started for the first time in a session, the Catalog window shows Releases. Use the two icons at the top of the Catalog window to display either the Releases view or the Software Packages view.

Table 15–7 *Catalog Views*

Task	Method	Result
To display software packages	Click Show Software Packages on the Products menu.	The Software Packages Catalog window appears. Software packages are also shown in the Navigator window under their individual hosts. You may need to refresh before viewing the hosts.
To display releases	Click Show Releases on the Products menu.	The Releases Catalog window appears. Releases are also shown in the Navigator window under their individual hosts. You may need to refresh before viewing the hosts.

Setting Preferences

When you create releases or software packages, Oracle Software Manager automatically lists these objects in the Catalog window in descending order: the object created most recently appears first.

However, you can change the order in which these objects appear to ascending order. You control this sort order through the Preferences dialog box.

To change the sort order of your objects:

1. On the Options menu, click Preferences.
2. To change the sort criteria so that objects appear with the oldest date listed first, click Ascending Order.
3. Click Done to close the Preferences dialog box.

These preferences pertain only to the order in which objects appear in Oracle Software Manager windows. Preferences do not affect objects in the Navigator window.

Toolbar

Below the menu bar is the toolbar, which can be displayed or hidden through the view menu. The toolbar icons represent items in the menus. Move the mouse cursor over a toolbar icon to display the description of the icon's function in the status bar at the bottom of the application.

Drag and Drop

The Oracle Software Manager supports dragging and dropping of certain window objects into other objects or into other windows. For example, software packages can be dragged from the Catalog window and dropped into certain objects in the Navigator; host objects can be dragged from one part of the Navigator window to another.

The following table shows the effects of drag and drop behavior in the Oracle Software Manager application.

Table 15–8 Drag and Drop Behavior

Objects	Can be dragged and dropped into...	Result
Product	Create Package dialog box Selected Products list	Product added to package.
Host icons	Host Folders in the Navigator windows	Host Folder containing this host icon.
Package icons	Host icon	Package distributed and/or installed on the host.

The Software Distribution Process

You can use Oracle Software Manager to distribute software to one or more hosts across the network. The software distribution process has four steps.

Table 15–9 The Software Distribution Process

Step	Task	Process	See Page
1	Specify Distribution Hosts	<p>Using Oracle Software Manager, designate one or more distribution hosts to house software releases and software packages.</p> <p>The software packages are later distributed over the network.</p>	15-14
2	Create Releases	<p>Using Oracle Software Manager, create software releases on the distribution host(s).</p> <p>These releases form the basis for the construction of software packages. You can view the status of the release creation under the current jobs window in Oracle Enterprise Manager.</p> <p>The list of releases is in the Oracle Software Manager Catalog window.</p> <p>The release properties describes products contained in release.</p>	15-15
3	Create Software Package	<p>Using Oracle Software Manager, select products from releases on a distribution host and create a software package that can be installed over the network.</p> <p>Note: If a product needs a custom installation, you must modify the package response file before distribution and/or installation. You can view the status of the package creation under the current jobs window in Oracle Enterprise Manager.</p> <p>The list of created packages is in the Oracle Software Manager Catalog window.</p> <p>Package properties have the list of products it contains.</p> <p>Refer to "Response File Generator" on page 15-37 for instructions.</p>	15-18

Table 15–9 The Software Distribution Process(Cont.)

Step	Task	Process	See Page
4	Distribute and Install Software Packages	Using Oracle Enterprise Manager job system, install, deinstall, delete, or distribute the software package. You could also use Oracle Software Manager's drag and drop feature to distribute and install software packages.	15-21

Once the package is created, it can be copied, using the Distribute Software Package task, to other distribution hosts. These packages are designated as copies and cannot be altered. Copying a package to other distribution hosts allows the distribution load to be spread across multiple servers, which can be spread geographically over a wide-area network.

Specify Distribution Hosts

The first step in the software distribution process is to specify which hosts will be distribution hosts.

A *distribution host* is a server that has been designated as a creation, distribution and storage center for software packages and releases. Like all hosts, a distribution host runs the Oracle Intelligent Agent and appears in the Navigator window.

You can designate any host as a distribution host and add or delete distribution hosts from the application, but there must be at least one distribution host designated on each network managed by an Oracle Software Manager application.

Any host with adequate storage space can be a distribution host, including the application machine, if the application machine is also running the Oracle Intelligent Agent software.

To designate a distribution host:

1. On the Products menu, click Distribution Hosts. The Distribution Hosts dialog box appears.
2. Select the host(s) you want to serve as distribution hosts from those listed in the All Hosts list.
3. Click Add to add the hosts to the Distribution Hosts list, then click OK. The icon changes in the Navigator window to reflect its new status.

A distribution host can only create packages whose operating systems are the same as its operating system.

You can only delete a distribution host if it is not the source host for a software package.

Removing distribution host status does not delete the host itself.

Create Releases

The second step in the software distribution process is to create releases. A release is an Oracle-defined collection of software whose location is registered with the Oracle Software Manager application.

Releases can be a stage area on a CD, a stage area on a network or local drive, or a stage created using the Oracle File Packager. For more information, see the Oracle File Packager section on page 15-38.

Releases can be created only on distribution hosts. If the application machine is also a distribution host, you can create releases on the application machine.

The Create Releases section describes the following procedures:

- [Creating a Release](#) on page 15-15
- [Displaying the Properties of a Release](#) on page 15-17
- [Removing a Release from the Releases Catalog Window](#) on page 15-17

Creating a Release

Note: Before creating a release, check that you have completed the following configuration steps through the Oracle Enterprise Manager console:

- Check that you have registered with the up/down host event, accept 3rd party events.
- Check that you have set up your preferences for the node that you will be distributing the package to.
- Check that you have set the environment variable OSM_PACKAGE_DIR. The software packages, which may be large, will be stored in this location on agent machines. A directory with sufficient space is required. If you do not configure this variable, the default will be %ORACLE_HOME%\NET80\AGENT\PACKAGES on Windows NT or %ORACLE_HOME/NETWORK/agent/packages on UNIX.
- Check that you have refreshed the host that you will be distributing and creating the packages to. In the Oracle Software Manager application, select `Refresh Hosts` from the toolbar, menu, or F5.

Refer to the *Oracle Enterprise Manager Configuration Guide* for detailed information.

To create a release:

1. From the Oracle Software Manager, click Create Release on the Products menu. The Create a Release dialog box appears.

See [Figure 15–3, "The Create a Release Dialog Box"](#).

Figure 15–3 The Create a Release Dialog Box



2. Type a name to identify the release. The name is limited to 10 characters.
3. Specify the distribution host where the release information resides. This is the host on which the Oracle CD or staging area is made available.
4. Type the directory on the distribution host where PRODINFO.ORA is stored for this release (for example, the directory on the CD or network drive where the release is stored).
5. Click OK. The Creating Release message box appears, then the Release Properties information box, which shows the information you just specified.

Software Manager looks for PRODINFO.ORA in the location you have specified in Create Release dialog. Software Manager also recursively searches for *.PRD files under that location. The contents of PRODINFO.ORA and each PRD file are read and returned as a TCL list to the Oracle Software Manager application. The application stores the information about the release and the products it contains in its repository.

6. Click Save. The new release is created and catalogued in the Releases Catalog Window.

When Oracle Software Manager creates a release, it creates pointers to the location and identifying characteristics of that release. The contents of the release are not copied anywhere on the distribution host's or the application workstation's hard disk.

Displaying the Properties of a Release

To display the properties of a release:

1. On the Products menu, click Show Releases. The Releases Catalog window appears.
2. Click the release you want to review.
3. On the View menu, click Properties. The Release Properties information box appears, open to the Properties page. Examine both the Properties and the Products pages of this dialog box.

Releases are marked with three types of icons.

Type of Icon	Indicates
normal	The release is ready to be used.
green light	The release is being created.
red light	The release was not created because of an error. You can look at the properties of the release to see the error information.

Removing a Release from the Releases Catalog Window

To remove a release from the Releases Catalog window:

1. In the Releases Catalog window, select the release(s) you want to remove.
2. On the Products menu, click Delete Release. The Delete dialog box appears, displaying a confirmation message.
3. Click Yes, or Yes to All to delete the release(s) from their designated hosts.

Create Software Packages

The third step in the software distribution process is to create software packages. A *software package* is a collection of products from one or more releases that can be distributed to/installed on hosts in the network.

You create software packages by selecting from the list of all products from created releases for a particular operating system on a particular distribution host. You might combine products from the same or multiple releases (as long as they are from the same distribution host) to form unique software packages.

You can then create a job in Oracle Enterprise Manager to distribute or install these packages on any host in the network.

The Create Software Packages section contains the following topics.

- [Creating a Software Package](#) on page 15-19
- [Deleting a Software Package](#) on page 15-21

Creating a Software Package

To create a software package from the Oracle Software Manager:

1. On the Products menu, click Create Software Package. The Define New Software Package dialog box appears. The default name for the package includes a number based on the date and time.

Figure 15–4 *Define New Software Package*



2. If desired, type a different name (than the one provided as a default) for the software package in the Name field. The name is limited to 10 characters.
3. Select the appropriate operating system from the list, if desired.
4. Click the Products tab.
5. In the Create Package on Host list, select the host on which you wish to build a package.

6. In the Build Package From list, choose one of the following options to display all products available for the specified operating system, or a subset of them.
 - a. To display all products in all releases on the selected host, select All Products on this Host. Then continue to Step 7.
 - b. To display all releases on the selected host, select Product by Release. The Select Release list appears. Select a release to display all the products in that release. Then continue to Step 7.
7. In the Available Products list, select the products to add to the package. Oracle Software Manager always shows one default operating system. You can only add products that belong to the same operating system.
8. Click Add to add the selected products to the package. The products are added to the Products in the Package list.
9. Click Save. Packages are saved under OSM_PACKAGE_DIR. If you did not configure the OSM_PACKAGE_DIR environment variable, the default will be %ORACLE_HOME%\NET80\AGENT\PACKAGES on Windows NT or %ORACLE_HOME/NETWORK/agent/packages on UNIX.

When Oracle Software Manager creates a software package, all the files representing the products in the package are copied to the distribution host, in compressed form.

Software packages are shown in the Software Packages view of the Catalog window and in the Navigator window under the host's packages.

Note: If a product needs a custom installation, you must specify the response file before the installation of the package.

A *response file* is a text file with the extension .RSP that contains the default responses for a package's installation. Oracle Software Manager uses the information in response file to install each product in a software package.

Software packages are marked with three types of icons.

Type of Icon	Indicates
normal (package only)	The package is ready to be used.
green light	The package is being created.
red light	The package was not created because an error occurred. You can look at the properties of the package to see the error information.

Deleting a Software Package

To delete a software package:

1. In the Software Packages Catalog window, select the software package you want to remove from the list of available packages.
2. On the Products menu, click Delete Software Package. The Remove Software Packages dialog box appears.
3. Click Delete.

Distribute and Install Software Packages

The fourth step in the software distribution process is to use the Oracle Enterprise Manager job system to distribute and install software packages on specified hosts.

For detailed information, please refer to [Software Manager Tasks](#) on page 15-24.

The Oracle Software Manager Network

The Software Manager network is comprised of Oracle Software Manager, the Oracle Enterprise Manager Console, one or more Windows NT and UNIX hosts running Enterprise Manager Intelligent Agents.

Setting up a Software Distribution Network

Before creating a software package that can be installed via the Enterprise Manager Job system, you must first designate at least one managed host as a distribution host. The distribution host acts as a clearinghouses for all software staging activities. Any host with adequate storage space can be used as a distribution host.

Establishing Communication in the Network

Assuming that all machines have the required software installed and configured, communication is established in the network.

The Oracle Software Manager network looks for and registers host(s). It initiates actions and the Intelligent Agent is responsible for executing them.

During installation of the Oracle Software Manager application, configuration information about each host is stored in the Oracle Enterprise Manager repository database.

The application operator sets Oracle Enterprise Manager to regularly query the Intelligent Agents to determine status. The application then registers the host, which appears in the Navigator window.

See the *Oracle Enterprise Manager Configuration Guide* for more information about registering hosts.

Flow of Information

The Oracle Software Manager network initiates actions and the Intelligent Agent is responsible for executing them.

The process works as follows:

1. The operator registers software and creates packages and jobs using Oracle Enterprise Manager and Oracle Software Manager.
2. Upon submission, the application sends the jobs to the target hosts through the communications layer.
3. The hosts receive the jobs, and the Intelligent Agents interpret the instructions contained within, and run the tasks in the job.
4. The Agents send status reports back to the application.

Oracle Software Manager uses a lightweight store-and-forward messaging system that keeps network traffic to a minimum. In the event of network problems, any message that is waiting to be sent, either from the application or the agent, is held until the network is back online. Communications between the application and agents are not resource intensive; they require only one connection per agent. If a host becomes unavailable, the connection is freed.

Software Manager Tasks

After using Software Manager to create software packages and/or release, you use the Oracle Enterprise Manager Job system to implement the administrative functions of software distribution. The following tasks are defined in the Oracle Enterprise Manager Job system and executed by the Oracle Enterprise Manager Intelligent Agent residing on managed hosts.

The Distribute and Install Software procedure consists of the following subsets

- [Determining the Job Name, Destination Type, and Job Destination](#) on page 15-24
- [Choosing a Task](#) on page 15-26
- [Setting Parameters](#) on page 15-28
- [Saving and Submitting the Job](#) on page 15-36

Determining the Job Name, Destination Type, and Job Destination

To set up job tasks with the Oracle Enterprise Manager Job Scheduling system:

1. Select Create from the Job menu to create a new job. The Create Job property sheet appears.
2. Enter a name for the new job in the Job Name field of the General page. You may also enter a description for the job in the Description field
3. Select “Node” from the Destination Type pull-down list.

See [Figure 15–5, "General Property Sheet"](#).

Figure 15–5 *General Property Sheet*



4. From the Available Destinations list, select a destination node and click the << (Add) button to add the destination to the Selected Destinations list. Only network objects that have an agent running are included in the list of available destinations.

If you need to remove a destination from a job, select the destination in the Selected Destination list and click the >> (Remove) button. You can remove one node per task for deinstall and install.

Choosing a Task

To choose the tasks that you want to perform:

1. Click on the Tasks page tab of the Create Job property sheet. The Tasks property sheet appears.

See [Figure 15–6, "Tasks Property Sheet"](#).

Figure 15–6 *Tasks Property Sheet*



2. From the Available Tasks list, select a task and click the << (Add) button to include the task in the job. You can add multiple tasks to the job from the Available Tasks scrolling list.

To remove a task from this list, select the task from Selected Tasks and click the >> (Remove) button.

3. 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.

You can make a task conditional on a previous task. Select a task and press the right arrow button to indent the task. A conditional, or indented, task is executed only if the previous task located further to the left successfully executes.

The following tasks can be used to complete the software distribution process:

Job Tasks	Function
Install Package	Install Package Task is used to specify the software package to be installed at the target nodes (one or more hosts in the network). The package must reside on the host, having been created there or placed there using the Distribute Package task.
Deinstall Package	Deinstall Product Task directs the Oracle Installer to remove Oracle products from the target nodes (one or more hosts in the network), automatically reconfiguring system parameters.
Delete Package	Delete Package Task specifies a software package to be deleted from the target node (the host where it resides)
Distribute Package	Distribute Package Task specifies a software package to be distributed to the target node without installing it anywhere else.

Setting Parameters

To set parameters:

1. Click the Parameters page tab of the Create Job property sheet to set the parameters for the selected task.
2. Select the task in the Selected Tasks list. The parameters for the selected task are displayed on the right side of the Parameters Page.

The parameters vary according to the job task. For information on the specific task parameters, refer to the appropriate parameter and function charts.

- [Install Package Task](#) on page 15-29
- [Deinstall Product Task](#) on page 15-31
- [Delete Package Task](#) on page 15-33
- [Distribute Package Task](#) on page 15-34

Install Package Task

Use Install Package Task to specify a software package to be installed at the target nodes.

Figure 15-7 Install Package Task



Refer to [Table 15-10, "Install Package Task Parameters"](#) for descriptions.

Table 15–10 Install Package Task Parameters

Parameters	Functions
Choose a Host	Click “Choose a Host” to view the list of host from the job’s target list
Select Software Package	Lists all packages available for installation, their operating system platforms, their source distribution hosts, and the number of products they contain. Scroll through the list and click the desired package(s).
Install Shared (available for client installation)	Click “Installed Shared” to specify that the package should be installed in shared mode.
Use Separate Response File From Destination Node	Click “Use Separate Response File” to specify a customized response file for use during package installation. To install a package into a specific Oracle Home, you must specify it in a response file. Note: the customized .RSP you specify should be on the target host (destination). In version 1.3 and earlier, the .RSP specified was on the application machine.
Full Pathname	The response file location is relative to the agent machine and not the application.

To create a separate response file, run installer in record mode (using the `/rspdest` command line option). From the DOS prompt, type

```
orainst /install <product label corresponding to the interface label> /prd
<location with path of prd file and the prd file name> /rspdest <location of
where you want your rsp file and the name of the rsp file>
```

The installer records the response which you have specified. You can then modify the response file.

Note: For multiple target jobs, you need to specify parameters for each host.

Deinstall Product Task

Deinstall Product Task directs the Oracle Installer to deinstall Oracle products from the target nodes.

Figure 15–8 Deinstall Product Task



Refer to [Table 15–11, "Deinstall Product Task Parameters"](#) for details.

Table 15–11 Deinstall Product Task Parameters

Parameters	Functions
Choose a Host	Click “Choose a Host” to view the host on which the selected product resides.
Deinstall Product from Oracle Home	Select a home and the products to delete from the home.
All Product(s)	Lists all products already installed on the selected home of the target node. To select a product to be deinstalled, click its listing, then click Add. The product moves to the Products to Deinstall list.
Add	Click “Add” to add a selected product icon to the Selected Products list.
Remove	Click “Remove” to remove a selected product icon from the Selected Products list.
Products to Deinstall	Shows the products that will be removed from the specified home on the specified node. To remove a product from this list, click the product, then click Remove. The product moves back to the All Products list.

Note: Deinstalling an Oracle product for Windows requires that the Oracle Installer for Windows be installed.

Delete Package Task

Delete Package Task specifies a software package to be deleted from the target node(s).

Table 15–12 *Delete Package Task Parameters*

Parameters	Functions
Delete Package from Host	Click “Delete package from Host” to view the host on which the selected package resides. The host that is initially selected (the default) is the source host for the selected package
Select a Software Package to Delete	Shows all software packages, their operating systems, number of products, and the date each was created. To select a software package to delete, click the first column of its listing.

Distribute Package Task

Distribute Package Task specifies a software package to be distributed to the target node.

Figure 15–9 Distribute Package Task



Refer to [Table 15–13, "Distribute Package Task Parameters"](#) for details.

Table 15–13 *Distribute Package Task Parameters*

Parameters	Function
Select a Software Package to Distribute	Click "Select a Software Package to Distribute" to select a software package from the list. It shows all software packages, their operating systems, number of products, and the date each was created.
Distribute Separate Response file	Click "Distribute Separate Response File" to specify a separate response file located on the agent machine that you are distributing from and not the application machine
Full Pathname	The response file location is relative to the agent machine and not the application.
Get Package From	Click "Get Package From" to view the host on which the selected package resides. The host that is initially selected (the default) is the source host for the selected package

Note: You can only distribute one software package per task. To distribute more than one software package at a time, add another Distribute Package task to the job.

When distributing a package from agent #1 to agent #2, you can optionally specify a separate response file located on agent #1 to be distributed to agent #2 as its default response file. This is an advanced feature that can be accessed through the Oracle Enterprise Manager console but not through the Oracle Software Manager application.

If you specify a separate response file, the task copies the response file from agent #1 to agent #2 and renames it so it becomes the default response file.

Saving and Submitting the Job

To save and submit the new job:

1. Click the Save button to save the new job in the Job Library window.
2. Click the Submit button to submit the job to the agents at the selected destinations.

When the agent begins processing the job, the job appears in the Active Jobs page in the Job window.

Custom Configurations

This section contains information on the Response File Generator and the Oracle File Packager.

Response File Generator

A *response file* is a text file with the .RSP extension that contains the default responses for a package installation.

Response files can be edited by hand in a text editor, or they can be edited at the distribution host using the Response File Generator. The Response File Generator runs an installation during which you provide the answers you want to the installation questions. They are recorded in a new response file.

The Response File Generator generates files to a default location for a package. You no longer need to specify input and output filenames in order to generate an RSP file. Instead, the Response File Generator uses the location OSM_PACKAGE_DIR as both input and output files. If you did not configure the environment variable OSM_PACKAGE_DIR, the default will be %ORACLE_HOME%\NET80\AGENT\PACKAGES on Windows NT or %ORACLE_HOME/NETWORK/agent/packages on UNIX.

You can run the Response File Generator on a target machine in order to get a customized response file for your software package.

To generate an .RSP file:

1. Create the package.
2. Run the Response File Generator on your agent machine. This installs the package on that host and generates an RSP file.

Note: In Software Manager, versions 1.3 and earlier, you had to move the RSP to your application machine for distribution and installation to other hosts. In Software Manager, versions 1.4 and later, distribution and installation automatically uses this response file unless you specify another response file. You do not need to move or rename the RSP file.

Modifying a Response File

To modify a response file (.RSP):

1. Open the .RSP file in a text editor. The format looks similar to the following example:

```
*****
#The user is asked which of the two SQL*Net products
#to install. The net2_adapter_content is set to
#whichever product is desired:
#"SQL*Net Client"
#"SQL*Net Server"
*****/
#net2_adapter_content="SQL*Net Client"
```

2. Locate the installation setting(s) you want to change.

For example, the legal settings in the example above are SQL*Net Client and SQL*Net Server. The default setting is SQL*Net Client.)

3. Change the installation setting(s) as required, making sure to surround the new setting(s) with double quotation marks, as shown in the following example:

```
net2_adapter_content="SQL*Net server"
```

4. Save the changes to the .RSP file and close it.

If you are unsure about whether or not a product requires a custom installation, try installing the package using the response file generator. If the dialog asks for your input, you will need a response file. Enter in the responses and the response file is generated for you. Then you can modify it if you want.

Creating Staging Areas with Oracle File Packager

Oracle File Packager allows you to group files into units that can be installed/distributed via Oracle Installer, Oracle Client Software Manager, and Oracle Software Manager.

The Oracle File Packager Wizard guides you through creating a staging area and it generates the required installations scripts. These stage areas can be used by Oracle Software Manager to create Releases and Packages that can be distributed via Oracle Software Manager.

Please refer to the Oracle File Packager online help for more information.

Part III

Reference

- Chapter A, “Using Server Manager in Line Mode”
- Chapter B, “DBA Command Reference”
- Chapter C, “Compatibility with SQL*DBA”

Using Server Manager in Line Mode

This chapter describes how to use Server Manager in line mode. Line mode is useful for performing unattended operations, such as running nightly batch jobs or scripts. In addition, you can use line mode when a graphical device is unavailable.

Note: Server Manager Line Mode is available on Windows NT and Motif platforms.

Starting Server Manager in Line Mode

You start up Server Manager in line mode by typing the appropriate command at the operating system prompt. For example, on some systems the command `svrmgrl` starts up Server Manager in line mode.

You can also start up Server Manager in line mode and execute a script. For example, to execute the script `SMSTART.SQL`, enter the following at the operating system prompt.

```
> svrmgrl command=@smstart
```

Note: The exact command for starting up Server Manager in line mode depends on your platform. For information about starting up Server Manager in line mode, see your operating system-specific Oracle documentation.

Using Server Manager in Line Mode

In line mode you can execute the Server Manager commands described in Appendix B, “DBA Command Reference”. In addition, you can execute SQL statements and PL/SQL code.

Entering Server Manager Commands

Single-line Server Manager commands do not require punctuation or terminators. A Server Manager command may span more than one line if you use a backslash (`\`) at the end of each line to indicate a continuation.

Entering SQL or PL/SQL Code

You can enter single-line or multiple-line SQL statements in line mode. Continuation characters are not needed in multiple-line SQL statements. To end a SQL statement and execute it, either type a semicolon (`;`) at the end of the statement or type a slash (`/`) by itself on the last input line.

To execute PL/SQL code, type a slash (`/`) by itself on the last input line.

Running Scripts

In line mode you can also run scripts using the `@` command. For example, to run the script `TEST.SQL`, issue the statement `@test`. If you do not specify a script name, Server Manager prompts you for one, as in the following example:

```
SVRMGR> @  
Name of script file: test
```

DBA Command Reference

This appendix describes the Server Manager DBA commands available in line mode. These commands can also be used in a SQL Worksheet, although some of these commands behave differently when used in a worksheet.

DBA Commands

The DBA commands described in this appendix are:

- *@ (at symbol)*
- *ARCHIVE LOG*
- *CONNECT*
- *DESCRIBE*
- *DISCONNECT*
- *EXECUTE*
- *EXIT*
- *PRINT*
- *RECOVER*
- *REMARK*
- *SET*
- *SHOW*
- *SHUTDOWN*
- *SPOOL*
- *STARTUP*
- *VARIABLE*

Note: If a reserved word is used as an object name in a DBA command, it must be enclosed in quotes. For example, to start up a database that is named with the reserved word V7, issue the command as follows:

```
STARTUP OPEN 'V7'
```

@ (at symbol)

Purpose Run scripts containing SQL, PL/SQL, and DBA commands.

Prerequisites You must have previously created the script and stored it as an operating system file.

Syntax @ command ::=



where:

scriptname

The name of an operating system file that contains a script of statements recognized by Server Manager.

Line Mode Behavior If you omit the *scriptname* argument, you are prompted for the script name.

Usage Notes This command enables you to run scripts of statements. To insert Comments in the script, use the REMARK command described in *REMARK* on page -17 . To spool the output, use the SPOOL command before running or within the script. The SPOOL command is described in *SPOOL* on page -30.

Do not use the @ command alone within a script unless you specify the full path name of the script.

To run other scripts from within a script, you can use the @@ command. In the original script include the line @@*second_scriptname* to run another script called *second_scriptname*. The @@ command indicates that the second script is in the same directory as the original script.

Additional Information: For information about the standard file selection dialog box for your system, see your operating system-specific documentation.

ARCHIVE LOG

Purpose Start or stop automatic archiving of online redo log files, manually (explicitly) archive specified redo log files, or display information about redo log files.

Note: This command applies only to the current instance. To specify archiving for a different instance or for all instances in a Parallel Server, use the SQL command ALTER SYSTEM.

Prerequisites You must be connected to an open Oracle database as INTERNAL, SYSOPER, or SYSDBA.

Syntax ARCHIVE LOG command ::=



where:

LIST

Requests a display that shows the range of redo log files to be archived, the current log file group's sequence number, and the current archive destination (specified by either the optional command text or by the initialization parameter LOG_ARCHIVE_DEST).

If you are using both ARCHIVELOG mode and automatic archiving, the display might appear like:

```
Database log modeArchive Mode
Automatic archivalEnabled
Archive destinationDISK9:ARCH
Oldest online log sequence30
Next log sequence to archive33
Current log sequence33
```

Because the log sequence number of the current log group and the next log group to archive are the same, automatic archival has archived all log groups up to the current one.

If you are using ARCHIVELOG but have disabled automatic archiving, the last three lines might look like:

```
Oldest online log sequence30
Next log sequence to archive30
Current log sequence33
```

If you are using NOARCHIVELOG mode, the "next log sequence to archive" line is suppressed.

The log sequence increments every time LGWR begins to write to another redo log file group; it does not indicate the number of logs being used. Every time an online redo log file group is reused, the contents are assigned a new log sequence number.

STOP

Disables automatic archival. If the instance is still in ARCHIVELOG mode and all redo log file groups fill, database operation is suspended until a redo log file is archived (for example, until you enter the command ARCHIVE LOG NEXT or ARCHIVE LOG ALL).

START

Enables automatic archiving. Starts the background process ARCH, which performs automatic archiving as required. If ARCH is started and a filename is supplied, the filename becomes the new default archive destination.

ARCH automatically starts on instance startup if the initialization parameter LOG_ARCHIVE_START is set to TRUE.

NEXT

Manually archives the next online redo log file group that has been filled but not yet archived.

ALL

Manually archives all filled, but not yet archived, online log file groups.

integer

Causes archival of the online redo log file group with log sequence number n. You can specify any redo log file group that is still online. An error occurs if the log file cannot be found online or the sequence number is not valid. This option can be used to re-archive a log file group.

'destination'

Specifies the destination device or directory in an operating system. Specification of archive destination devices is installation-specific; see your platform-specific Oracle documentation for examples of specifying archive destinations. On many operating systems, multiple log files can be spooled to the same tape.

If not specified in the command line, the archive destination is derived from the initialization parameter `LOG_ARCHIVE_DEST`. The command `ARCHIVE LOG START 'destination'` causes the specified device or directory to become the new default archive destination for all future automatic or manual archives. A destination specified with any other option is a temporary destination that is in effect only for the current (manual) archive. It does not change the default archive destination for subsequent automatic archives.

Additional Information: For information about specifying archive destinations, see your platform-specific Oracle documentation.

Usage Notes If an online redo log file group fills and none are available for reuse, database operation is suspended. The condition can be resolved by archiving a log file group.

Examples ARCHIVE LOG START

Starts up the archiver process and begins automatic archiving, using the archive destination specified in `LOG_ARCHIVE_DEST`.

```
ARCHIVE LOG STOP
```

Stops automatic archiving.

```
ARCHIVE LOG 1001 'DISK9:[TEMPARCH]TEMP'
```

Archives the log file group with the sequence number 1001 to the destination specified. 'TEMP' specifies the prefix of the filename on the destination device; the remainder of the filename is dependent on the initialization parameter `LOG_ARCHIVE_FORMAT`, which specifies the filename format for archived redo log files.

CONNECT

Purpose Connect to a database using the specified username.

Prerequisites Only valid username/password combinations can successfully connect. The AS clause allows users to connect to an instance with a default schema of SYS for database administration. Users connecting as SYSOPER or SYSDBA must have the necessary privileges to access the SYS schema.

Syntax CONNECT command ::=



where:

username

Any valid Oracle username for the current database. It can be a null string. If neither username nor INTERNAL is specified, you are prompted for a username and password.

password

The password corresponding to the specified username. Password can be a null string.

instance-path

A valid specification for an instance/database combination. If an instance is specified, it becomes the current instance for the duration of the connection. It does not become the default instance for subsequent connections.

SYSOPER/SYSDBA

The AS clause allows privileged connections by users who have been granted SYSOPER or SYSDBA system privileges, if using a password file, or their operating system equivalents, if using OS authentication.

/

Uses operating system authentication.

INTERNAL

Connects as username SYS via keyword INTERNAL. Should be used rarely and only by the DBA for certain maintenance purposes; requires operating system privileges. CONNECT INTERNAL is supported for backwards compatibility only.

Line Mode Behavior If you omit the *password*, you are prompted for one. If you omit both the *username* and *password*, you are prompted for both.

SQL Worksheet Behavior If you omit the *username* argument, the Connect dialog box displays.

Usage Notes If only the Oracle username is specified, the password is requested using the prompt "Password:". The entered password is not echoed.

CONNECT can be used without a DISCONNECT to connect to another username.

The connect is always accompanied by the opening of a cursor.

You can specify an instance path for a remote instance when connecting. If you attempt a connection to a remote instance, authentication occurs on the remote node, so you must have the appropriate privileges on the remote node.

Non-Secure Connection Connecting as SYSOPER or SYSDBA over a

To connect to Oracle as a privileged user over a non-secure connection, you must satisfy the following conditions:

- The server to which you are connecting must have a password file.
- You must be granted the SYSOPER or SYSDBA system privilege.
- You must connect using a username and password.

For information about creating a password file, see the *Oracle7 Server Concepts* and the *Oracle7 Server Administrator's Guide*, or **Secure Connection** Connecting as SYSOPER or SYSDBA over a Local

To connect to Oracle as a privileged user over a local or a secure connection, you must satisfy either of the following sets of conditions:

- You can connect using a password file, provided you meet the criteria outlined above.
- If the server is not using a password file, or you have not been granted SYSOPER or SYSDBA and are therefore not in the password file, your operating

system name must be authenticated by the operating system for a privileged connection. This form of authentication is system-specific.

Additional Information: For information about operating system authentication, see your operating system-specific Oracle documentation.

Examples To connect to an instance on the current default node, enter:

```
CONNECT
```

To connect to an instance on the current node as username SCOTT with password TIGER, enter:

```
CONNECT scott/tiger
```

To use a password file to connect to an instance on the current node as a privileged user named SCOTT with password TIGER, enter:

```
CONNECT scott/tiger AS SYSDBA
```

Note that your default schema is now SYS, not SCOTT.

To perform the same connection using OS authentication, enter:

```
CONNECT / AS SYSDBA
```

To connect to an instance on a different node as username SCOTT with password TIGER, enter:

```
CONNECT scott/tiger@instance-path
```

Connecting as INTERNAL Using Server Manager you can connect as INTERNAL to a release 7.0 database. CONNECT INTERNAL is supported for backwards compatibility only. For information on how to use CONNECT INTERNAL, refer to your *Oracle7 Server Administrator's Guide*.

DESCRIBE

Purpose Describe a function, package, package body, procedure, table, or view.

Prerequisites You must be currently connected to a database.

Syntax DESCRIBE command ::=



where:

name

The name of the object to describe.

Usage Notes If you do not specify an object type, Server Manager attempts to determine the object type. For the statement

```
DESCRIBE emp
```

Server Manager describes the table or view if a table or view with the name EMP exists. If no table or view with that name exists in the user's schema, Server Manager will search for a PL/SQL function, procedure, or package with that name and describe it.

The shorthand DESC is now equivalent to the command DESCRIBE. This means that DESC has been added to the list of Server Manager reserved words. If you have a database object named DESC, you need to place quotes around the name to show that it is an object name. For example:

```
INSERT INTO 'DESC' VALUES('onetwothree', 123); )
```

Note: Some versions of PL/SQL allow the user to describe a package by giving the package name, and some versions require the user to specify an object in the package to describe. Server Manager supports package description for those versions of PL/SQL which support this functionality.

Example A command like

```
DESCRIBE scott.addemp
```

produces output similar to the following example:

```
PROCEDURE SCOTT.ADDEMP (EMPNO INTEGER, ENAME VARCHAR2, SAL NUMBER(9,2))
```

The use of the object type is no longer required. If you do not specify an object type, Server Manager attempts to determine the object type. For the statement

```
DESCRIBE emp
```

Server Manager describes the table or view if a table or view with the name EMP exists.

If no table or view with that name exists in the user's schema, Server Manager searches for a PL/SQL function, procedure, or package with that name and describes it. For the statement

```
DESCRIBE payroll.emp
```

Server Manager looks for a table or view in schema PAYROLL with the name EMP.

If there is no PAYROLL schema or no EMP table/view in the PAYROLL schema, Server Manager will search the current schema for a package PAYROLL containing a function or procedure EMP.

If both a table EMP in schema PAYROLL and a package PAYROLL with function/procedure EMP in the current schema exist, the statement

```
DESCRIBE FUNCTION payroll.emp
```

This allows you to specify the type of the object you want to describe.

Server Manager now has the capability to describe both functions and procedures contained in packages. Previously only functions and procedures not contained in packages could be described.

Note: Some versions of PL/SQL allow the user to describe a package by giving the package name, and some versions require the user to specify an object in the package to describe. Server Manager supports package description for those versions of PL/SQL which support this functionality.

DISCONNECT

Purpose Disconnect from an Oracle server.

Prerequisites You must be currently connected to a database.

Syntax DISCONNECT command ::=

Usage Notes Upon disconnection, line mode reverts to the current default host machine after closing all open cursors and committing any uncommitted transactions.

Example DISCONNECT

EXECUTE

Purpose Executes a one-line PL/SQL statement.

Prerequisites You must be currently connected to a database. You must also have privileges to use any stored procedures, packages, package variables, and functions referenced with this command.

Syntax EXECUTE command ::=

Usage Notes You can execute only one line of PL/SQL code with the EXECUTE command. If you wish to execute a PL/SQL block with many lines, you must use the format:

```
BEGIN
[PL/SQL BLOCK]
END;
```

You can reference any command line mode bind variables in PL/SQL statements by preceding the variable name with a colon. The example below illustrates the use of bind variables with the VARIABLE, EXECUTE, and PRINT commands.

Example VARIABLE balance NUMBER

```
EXECUTE :balance := get_balance(34056)
PRINT balance
BALANCE
-----
4678.24
```

EXIT

Purpose Exits Server Manager line mode or closes a SQL Worksheet.

Prerequisites None.

Syntax EXIT command ::=

Line Mode Behavior The EXIT command leaves line mode unconditionally, commits the current transaction, and returns to the operating system prompt.

SQL Worksheet Behavior

The EXIT command commits the current transaction and closes the worksheet.

Example EXIT

PRINT

Purpose Print the value of a variable defined with the VARIABLE command.

Prerequisites None.

Syntax PRINT command ::=

where:

variable

The name of a variable defined by the VARIABLE command.

Usage Notes Bind variables referenced with the print command do not need to be preceded by a colon. You can use the PRINT command to display variables defined only in the current line mode session. The SET CHARWIDTH and SET NUMWIDTH commands can affect the display of the PRINT command.

If no variable name is specified all currently defined variables are printed.

Examples PRINT balance

```
BALANCE
```

```
-----
```

```
4687.24
```

```
SET CHARWIDTH 10
```

```
PRINT ename
```

ENAME

SCOTT

RECOVER

Purpose Performs media recovery on one or more tablespaces, one or more datafiles, or the entire database.

Prerequisites You must be connected to the Oracle server as INTERNAL, SYSOPER, or SYSDBA. You cannot use the RECOVER command when connected via the multi-threaded server.

Syntax RECOVER command ::=



UNTIL clause ::=



PARALLEL clause ::=



where:

DATABASE

Specifies recovering the entire database.

USING BACKUP CONTROL FILE

Specifies that a backup of the control file is being used instead of the current control file.

TABLESPACE *tablespace*

Specifies recovering a particular tablespace. Tablespace is the name of a tablespace in the current database. You may recover up to 16 tablespaces in one statement.

DATAFILE *filename*

Specifies recovering a particular datafile. You can specify any number of datafiles.

UNTIL CANCEL

Specifies an incomplete, cancel-based recovery. Recovery proceeds by prompting you with the suggested filenames of archived redo log files, and recovery completes when you specify CANCEL instead of a filename.

UNTIL CHANGE *integer*

Specifies an incomplete, change-based recovery. The integer is the number of the SCN following the last change to which you wish to recover. For example, if you want to restore your database up to the transaction with an SCN of 9, you would specify UNTIL CHANGE 10.

UNTIL TIME *date*

Specifies an incomplete, time-based recovery. Use single quotes, and use the following format:

```
'YYYY-MM-DD:HH24:MI:SS'
```

PARALLEL

DEGREE specifies the number of recovery processes used to apply redo entries to datafiles on each instance. An integer specified for DEGREE overrides the initialization parameter RECOVERY_PARALLELISM.

DEGREE DEFAULT indicates that twice the number of datafiles being recovered is the number of recovery processes to use.

INSTANCES specifies the number of instances to use for parallel recovery. The number of recovery processes specified with DEGREE is used on each instance, thus the total number of recovery processes is the integer specified with DEGREE multiplied by the integer specified with INSTANCES. INSTANCES is only pertinent for the Oracle Parallel Server.

INSTANCES DEFAULT or not including the INSTANCES keyword causes has operating system-specific consequences. Consult your Oracle operating system documentation for a description of the default behavior of the INSTANCES DEFAULT specification.

NOPARALLEL

Specifies that recovery is to proceed serially. Note that a specification of PARALLEL(DEGREE 1 INSTANCES 1) is equivalent to specifying the NOPARALLEL keyword.

The PARALLEL keyword overrides the RECOVERY_PARALLELISM initialization parameter. The number specified with the PARALLEL keyword is the number of recovery processes used to apply redo entries to datafiles.

SQL Worksheet Behavior In a SQL Worksheet, if you issue the RECOVER command without arguments, Server Manager brings up the Recover dialog box.

Usage Notes To perform media recovery on an entire database (all tablespaces), the database must be mounted EXCLUSIVE and closed.

To perform media recovery on a tablespace, the database must be mounted and open, and the tablespace must be offline.

To perform media recovery on a datafile, the database can remain open and mounted with the damaged datafiles offline (unless the file is part of the SYSTEM tablespace).

Before using the RECOVER command you must have restored good copies of the damaged datafile(s) from a previous backup. Be sure you can access all archived and online redo log files dating back to when that backup was made.

When another log file is required during recovery, a prompt suggests the names of files that are needed. The name is derived from the values specified in the initialization parameters LOG_ARCHIVE_DEST and LOG_ARCHIVE_FORMAT. You should restore copies of the archived redo log files needed for recovery to the destination specified in LOG_ARCHIVE_DEST, if necessary. You can override the initialization parameters by setting the LOGSOURCE variable.

During recovery you can accept the suggested log name by hitting return, cancel recovery by entering CANCEL instead of a log name, or enter AUTO for automatic file selection without further prompting.

If you have enabled autorecovery (that is, SET AUTORECOVERY ON), recovery proceeds without prompting you with filenames. Status messages are displayed when each log file is applied.

When normal media recovery is done, a completion status is returned.

For more information on recovery and the RECOVER command, see the *Oracle7 Server Administrator's Guide*.

Examples RECOVER DATABASE

```
RECOVER DATABASE UNTIL TIME 30-AUG-90:04:32:00
RECOVER TABLESPACE ts_one, ts_two
RECOVER DATAFILE 'data1.db'
```

REMARK

Purpose Enter a Comment, typically in SQL script files.

Prerequisites None.

Syntax REMARK command ::=



Usage Notes Primarily for batch use of line mode. The Comment is ignored by line mode and by Oracle. REMARK can be shortened to REM.

REM must be the first non-blank character string in the line.

Examples Examples of valid Comments embedded in a SQL file follow:

```
REM   This command file is used to create a
REM   database. Edit it to fill in file names
REM   and sizes, and invoke it from line mode.
REM
REM
CREATE DATABASE dbname ....
/* This is a SQL Comment */ ...
...
;
```

REMARKs are recognized by SQL*Plus, as well as by line mode. They are used to put Comments between SQL statements, while SQL Comments (/*...*/) are used to place comments within statements.


SET


Purpose Set or change characteristics of the current command line mode session.

Prerequisites None.

Syntax SET command ::=

See next page for syntax diagram.





where:

APPINFO

Registers the Server Manager application through the database's `DBMS_APPLICATION_INFO` package (Oracle 7.2 or later). By default, the APPINFO function is set to OFF.

Application registration allows DBAs to see what software is currently running to better monitor resource utilization for database tuning.

When APPINFO is turned on, application registration proceeds normally. If a SQL script is not being run, the default registration text string is "Oracle SQL Worksheet." Optionally, you can create a customized registration text string. If a script is being run through Server manager, the script name is used as the registration text string.

Text String Registers as:

```
SET APPINFO abc abc
SET APPINFO abc defabc
SET APPINFO "abc def"abc def
SET APPINFO 'abc def'abc def
SET APPINFO "abc def"Error
SET APPINFO 'abc def'Error
```

AUTO- RECOVERY

Specifying ON causes the RECOVER command to automatically apply the default filenames of archived redo log files needed during recovery. No interaction is needed when AUTORECOVERY is set to ON, provided the necessary files are in the expected locations with the expected names. The filenames used when AUTORECOVERY is ON are derived from the values of the initialization parameters LOG_ARCHIVE_DEST and LOG_ARCHIVE_FORMAT.

OFF, the default option, requires that you enter the filenames manually or accept the suggested default filename given.

CHARWIDTH integer

Sets the column display width for CHAR data. If entered with no argument, returns the setting to the default. The default is 80; the range of values is operating system-specific.

COMPATIBILITY

Sets compatibility mode to V6, V7, or NATIVE. The compatibility mode setting affects the specification of character columns, integrity constraints, and rollback segment storage parameters. NATIVE matches the version of the database.

CHAR Columns: When creating tables in Version 6 compatibility mode, CHAR columns are variable length. In Oracle7, such column definitions are fixed length.

Integrity Constraints: In Version 6 compatibility mode, the Version 6 syntax is still recognized, and the Oracle7 syntax is disabled. For V6 mode, table constraints on CREATE TABLE statements are specified with V6 syntax:

```
CREATE TABLE {UNIQUE | PRIMARY KEY} CONSTRAINT ....
```

and specified constraints are disabled by default. For V7 mode, table constraints are specified with Oracle7 syntax:

```
CREATE TABLE CONSTRAINT .... {UNIQUE | PRIMARY KEY}
```

and they are enabled.

Rollback Segment Parameters: Version 6 compatibility mode allows PCTINCREASE and MAXEXTENTS to be specified for rollback segments, as well as for other segments. Although the specifications are ignored, the syntax is allowed. (Use of these parameters is not recommended. They exist only for backward compatibility.)

Bind Variables: Bind variables of type VARCHAR2 are given type CHAR in Version 6 compatibility mode.

DATEWIDTH integer

Sets the column display width for DATE data. If entered with no argument, the setting returns to the default. The default is 9; the range of values is operating system-specific.

ECHO

ON enables echoing of commands entered from command files. OFF, the default, disables echoing of commands.

FETCHROWS integer

Limits the number of rows returned by a query. Useful with ordered queries for finding the "top ten" items in a category, for example. Also used with unordered queries to find the "first n" records that satisfy a given criteria. integer must be in the range 1-999,999.

In a SQL Worksheet the default is ECHO ON.

INSTANCE instance-path

Changes the default instance for your session to the specified instance path. Does not connect to a database. The default instance is used for

commands when no instance is specified.

Any commands preceding the first use of SET INSTANCE communicate with the default instance.

To reset the instance to the default value for your operating system, you can either enter SET INSTANCE with no *instance-name* or SET INSTANCE LOCAL. See your operating system-specific Oracle documentation for a description of how to set the initial default instance.

This command may only be issued when SQL*Net is running. You can use any valid SQL*Net connect string as the specified instance path. See your operating system-specific Oracle documentation for a complete description of how your operating system specifies SQL*Net connect strings. The maximum length of the instance path is 64 characters.

LOGSOURCE pathname

Specifies the location from which archive logs are retrieved during recovery. The default value is set by the LOG_ARCHIVE_DEST initialization parameter. Issuing the SET LOGSOURCE command without a pathname restores the default location.

LONGWIDTH integer

Sets the column display width for LONG data. If entered with no argument, the setting returns to the default. The default is 80; the range of values is operating system-specific.

MAXDATA integer

Sets the maximum data size. Indicates the maximum data that can be received in a single fetch during a SELECT statement. The default is 20480 bytes (20K). The maximum is operating system-specific.

NUMWIDTH integer

Sets the column display width for NUMBER data. If entered with no argument, the setting returns to the default. The default is 10; the range of values is operating system-specific.

RETRIES integer | INFINITE

Sets the number of tries that are attempted when the RETRY option is

used with the STARTUP command. INFINITE, the default, specifies an infinite number of retries.

SERVEROUTPUT

ON enables debugging output from stored procedures that use DBMS_OUTPUT.PUT and PUT_LINE commands. OFF disables output.

You can specify the size in bytes of the message buffer using the syntax SIZE *n*. That is the total number of bytes of all messages sent that can be accumulated at one time. The minimum is 2,000 bytes. If the buffer fills before calls to the get-message routines make room for additional message bytes, an error is returned to the message-sending program.

STOPONERROR

Specifying ON indicates that if a command file is running and an error occurs, execution should terminate. OFF disables STOPONERROR.

TERMOUT

ON, the default, enables terminal output for SQL commands. OFF disables terminal output. Useful for preventing output to terminal when spooling output to files. Note that nothing appears on the terminal until SET TERMOUT ON is used.

TIMING

ON displays parse, execute, and fetch times (CPU and elapsed) for each SQL statement executed. OFF, the default, does not display timing information.

Examples SET INSTANCE D:DEV-PROD

```
SET TIMING ON
SET LONGWIDTH 132
SET NUMWIDTH 20
SET CHARWIDTH 5
```

Either of the following commands can be used to revert to the initial default host:

```
SET INSTANCE
SET INSTANCE LOCAL
```


SHOW

Purpose Show settings currently in effect.

Prerequisites None.

Syntax SHOW command ::=



where:

ALL

Shows all settings except for ERRORS, PARAMETERS, and SGA.

APPINFO

Shows the current status (ON/OFF) and the application registration text.

CHARWIDTH

Shows the column display width for CHAR data.

COMPATIBILITY

Shows the sessions compatibility mode: V6, V7, or NATIVE.

DATEWIDTH

Shows the column display width for DATE data.

ECHO

Shows whether or not commands from command files are echoed.

ERRORS

Shows the errors generated from the last compilation of a procedure, package, or function, if any. This option displays the line number, column number, and the error message generated. Use the SHOW ERRORS command if line mode error 72 or 73 is returned after compiling a package.

The SET CHARWIDTH command can be used to expand or truncate the display from the SHOW ERRORS command.

FETCHROWS

Shows the numerical limit set for the number of rows returned by a query.

INSTANCE

Shows the connect string for the default instance. SHOW INSTANCE returns the value LOCAL if you have not used SET INSTANCE or if you have used the LOCAL option of the SET INSTANCE command.

LOGSOURCE

Shows the current setting for archive log location. Displays DEFAULT if the default setting is in effect, as specified by the LOG_ARCHIVE_DEST initialization parameter.

LONGWIDTH

Shows the column display width for LONG data.

MAXDATA

Shows the maximum data size.

NUMWIDTH

Shows column display width for NUMBER data.

PARAMETERS

Displays the current values for one or more initialization parameters. You can use a string after the command to see a subset of parameters whose names include that string. For example, if you enter:

```
SHOW PARAMETERS COUNT
you would see:
```

NAME	TYPE	VALUE
-----	-----	-----
db_file_multiblock_read_count	integer	12
spin_count	integer	0

The SHOW PARAMETERS command, without any string following the command, displays all initialization parameters.

RETRIES

Displays the number of retries that will be attempted when restarting an instance in parallel mode.

SERVEROUTPUT

Displays ON if output from stored procedures and functions is enabled. Otherwise, displays OFF. SERVEROUTPUT is reset to OFF if you issue a CONNECT command.

SGA

Displays information about the current instance's System Global Area.

SPOOL

If spooling is enabled, displays the name of the output spool file. Otherwise, displays OFF.

STOPONERROR

Displays whether or not errors encountered during execution of command files should stop execution of the file.

TERMOUT

Shows if output to the terminal is enabled.

TIMING

Shows whether or not the parse, execute, and fetch times (CPU and elapsed) for each SQL statement executed are shown.

Usage Notes SHOW with no arguments is the same as SHOW ALL.

Examples SHOW TIMING

returns a display such as:

```
Timing                OFF
```

SHOW ALL

returns a display like:

```
Instance              local
Spool                 OFF
Timing               OFF
Termout              ON
Echo                 OFF
Stoponerror          OFF
Autorecovery         OFF
Logsource             <default>
Maxdata              20480
Numwidth             10
Charwidth            80
Longwidth            80
Datewidth            9
```

```

Labwidth          32
Compatibility     NATIVE
Retries           infinite
Server Output     OFF

```

```
SHOW SGA
```

returns a display like:

```

Total Shared Global Area      4612820 bytes
Fixed Size                    36376 bytes
Variable Size                 4445372 bytes
Database Buffers             122880 bytes
Redo Buffers                  8192 bytes

```

```
SHOW ERRORS PACKAGE BODY name
```

returns a display like:

```

ERRORS FOR PACKAGE BODY name:
LINE/COLERRORS
-----
...

```

SHUTDOWN

Purpose Shut down a currently running Oracle instance, optionally closing and dismounting a database.

Prerequisites You must be connected to a database as INTERNAL, SYSOPER, or SYSDBA. You cannot be connected via a multi-threaded server.

Syntax SHUTDOWN command ::=



where:

ABORT

Proceeds with the fastest possible shutdown. Does not wait for calls to complete or users to disconnect. Does not close or dismount the

database, but does shut down the instance. Requires instance recovery on next startup. You must use this option if a background process terminates abnormally.

IMMEDIATE

Does not wait for current calls to complete, prohibits further connects, and closes and dismounts the database. Finally, shuts down the instance. Does not wait for connected users to disconnect. Does not require instance recovery on next startup.

NORMAL

Waits for currently connected users to disconnect from the database, prohibits further connects, and closes and dismounts the database. Finally, shuts down the instance. Does not require instance recovery on next startup. **NORMAL** is the default option.

Attention: In Server Manager, you can have several separate connections open at any time in multiple windows. If you have any connections open, remember to close them before performing a shutdown in normal mode. Otherwise, the shutdown will not complete.

Example SHUTDOWN

```
Database closed.  
Database dismounted.  
Oracle instance shut down.
```

SPOOL

Purpose Enable or disable spooling of output to a specified file.

Prerequisites None.

Syntax SPOOL command ::=



where:

filename

Any valid name for a spool file. If not specified, the file type or file extension is LOG. If a file by the specified name exists, an error is generated.

OFF

Closes the previously opened spool file.

Usage Notes If you do not specify a file, Server Manager prompts you for a filename.

Additional Information: The default filename is operating system specific. Refer to your operating system-specific Oracle documentation for the default filename on your operating system. To see whether you are currently spooling, enter SHOW SPOOL.

Note: When long lines of text (usually from a table query) are written to a spool file or to a terminal, newline codes are inserted in the text if the lines exceed the maximum line length for the platform.

Example To create a file named NOV2.LOG, enter:

```
SPOOL NOV2
```

STARTUP

Purpose Start an Oracle instance with several options, including mounting and opening a database.

Prerequisites You must be connected to a database as INTERNAL, SYSOPER, or SYSDBA. You cannot be connected via a multi-threaded server.

Syntax STARTUP command ::=



where:

FORCE

Shuts down the current Oracle instance (if it is running) with SHUTDOWN mode ABORT, before restarting it. If the current instance is running and FORCE is not specified, an error results. FORCE is useful while debugging and under abnormal circumstances. It should not normally be used.

RESTRICT

Only allows Oracle users with the RESTRICTED SESSION system privilege to connect to the database. Later, you can use the ALTER SYSTEM command to disable the restricted session feature.

PFILE=filename

Causes the specified parameter file to be used while starting up the instance.

MOUNT

Mounts a database but does not open it.

OPEN

Mounts and opens the specified database.

NOMOUNT

Causes the database not to be mounted upon instance startup. Cannot

be used with SHARED, EXCLUSIVE, PARALLEL, MOUNT, or OPEN.

RECOVER

Specifies that media recovery should be performed, if necessary, before starting the instance. STARTUP RECOVER has the same effect as issuing the RECOVER DATABASE command and starting an instance. Only complete recovery is possible with the RECOVER option.

Recovery proceeds, if necessary, as if AUTORECOVERY is set to ON, regardless of whether or not AUTORECOVERY is enabled. If a redo log file is not found in the expected location, recovery continues as if AUTORECOVERY is disabled, by prompting you with the suggested location and name of the subsequent log files that need to be applied.

If recovery fails using the RECOVER option, the database remains mounted and closed.

database

The name of the database to mount or open. If no database name is specified, the database name is taken from the initialization parameter DB_NAME.

Examples

To start an instance using the standard parameter file, mount and open the database, enter:

```
STARTUP
```

To shut down the current instance, immediately restart it without mounting or opening, and allow only database administrators to connect, enter:

```
STARTUP FORCE NOMOUNT RESTRICT
```

To start an instance using the parameter file TESTPARM without mounting the database, enter:

```
STARTUP PFILE=testparm NOMOUNT
```

To shut down a particular database, immediately restart and open it in parallel mode, allow access only to database administrators, and use the parameter file MYINIT.ORA. enter:

```
STARTUP OPEN databasename PFILE=myinit.ora FORCE SHARED RESTRICT
```

To start up an instance and mount but not open a database, you can use the following sequence of commands (the system's response is also shown):

```
CONNECT INTERNAL  
Connected.
```

```
STARTUP NOMOUNT  
Oracle instance started.
```

```
ALTER DATABASE MOUNT  
Statement processed.
```

At this point, you could run a maintenance command and then open the database, as shown in the following commands:

```
ALTER DATABASE ARCHIVELOG;  
Statement processed.
```

```
ALTER DATABASE OPEN;  
Statement processed.
```

VARIABLE

Purpose Declare a bind variable for use in the current session with the EXECUTE or PRINT command, or for use with a PL/SQL block.

Prerequisites None.

Syntax VARIABLE command ::=



where:

name

The name of the bind variable you wish to define.

type

The datatype of the bind variable. The valid datatypes are NUMBER, CHAR, and VARCHAR2. You can specify a length from 1 to 255 for

CHAR, and from 1 to 2,000 for VARCHAR2. If you do not specify a length, CHAR and VARCHAR2 default to the current setting of CHARWIDTH.

Usage Notes Bind variables defined with the VARIABLE command exist until the end of the session. Variables defined in your session cannot be accessed or changed by a different session.

Warning:

CHAR variables are fixed length, padded with blanks. VARCHAR2 variables are variable length. You must use CHAR when passing a character variable to a PL/SQL procedure that defines a CHAR parameter. Otherwise, a conversion error results. A PL/SQL procedure that expects a VARCHAR2 variable, on the other hand, will automatically convert a CHAR parameter to the proper form. (In general, it is a good idea to avoid the use of CHAR variables in PL/SQL procedures unless blank-padding is an absolute requirement.)

Note: In V6 compatibility mode, a CHAR variable is variable length, rather than fixed.

Examples VARIABLE balance NUMBER

```
VARIABLE ename CHAR(20)
```

Compatibility with SQL*DBA

This appendix lists the differences between Oracle7 SQL*DBA and Server Manager DBA commands.

Figure 15–10 Functional Differences

Feature	SQL*DBA	Server Manager
DESCRIBE for PL/SQL	Lists information about the procedure argument in a tabular form	Lists the definition of the procedure argument
PRINT <no args>	Does nothing	Prints all defined variables
PRINT (before connect)	SQLDBA> print a DBA-00302: not connected to a database	SVRMGR> print a A ----- 10
SET CHARWIDTH	Displays a character string equivalent to the number of bytes required. For example, if CHARWIDTH is set to 14 and a two-byte character set is used, seven characters are displayed.	Displays the number of characters specified by CHARWIDTH regardless of the number of bytes used by the character string. For example, if CHARWIDTH is set to 14, 14 characters are displayed.
SET COMPATIBILITY NATIVE	Unsupported	Server Manager uses the native SQL of the database to which you are connected
SHOW <no args>	SQLDBA> show show * DBA-00115: unexpected end of command	Prints current Server Manager settings
SPOOL <no args>	SQLDBA> spool File sqldba.log opened... No prompt	SVRMGR> spool Name for the spool file: (svrmgr.log) Prompts for a filename. The default is SVRMGR.LOG.
STARTUP	If any step of STARTUP fails, SQL*DBA backs out and leaves the database in a shutdown state.	If any step of STARTUP fails, Server Manager leaves the database in its current state. For example, if an error is encountered after the mount stage, Server Manager leaves the database mounted.
VARIABLE a CHAR	Default size is 1	Default size is CHARWIDTH
@<no args>	SQLDBA> @ ORA-07225: ... Prints an error	SVRMGR> @ Name of script file: Prompts for the name of the file
Empty line	SQLDBA> garbage 2> SQLDBA> SQLDBA>	SVRMGR> garbage 2> 3> garbage * ORA-00900: invalid SQL statement SVRMGR>

Figure 15–11 Cosmetic Differences

<i>Feature</i>	<i>SQL*DBA</i>	<i>Server Manager</i>
Prompt	SQLDBA>	SVRMGR>
Product Banner	SQL*DBA	Server Manager
Connect to idle instance	SQLDBA> connect internal con- nected.	SVRMGR> connect internal con- nected to an idle instance.
PRINT <i>a</i>	A ----- 10 1 row selected. The variable name is space padded: "A ".	A ----- 10 The variable name is not space padded: "A" Also, no feedback is listed.
SET	SQLDBA> set echo on SQLDBA>	SVRMGR> set echo on Echo ON SVRMGR> Server Manager echoes the new value
Error Message Prefix	DBA-xxxxxx: ...	MGR-xxxxxx: ...
Wrapping text	Writes 80 characters and then wraps	Writes the entire width and relies on the display device to deal with wrapping. This is better for wide terminals.

Figure 15–12 Unsupported Options

<i>Feature</i>	<i>SQL*DBA</i>	<i>Server Manager</i>
RECOVERY prompt		FROM <log source> clause is unsupported
SET FLAGGER...	SET FLAGGER ENTRY	Unsupported: Use the ALTER SESSION SET FLAGGER statement
SHOW FLAGGER	Shows FLAGGER value	Unsupported
SET LABELWIDTH		Unsupported
SET RETRIES		Unsupported
SHOW VARIABLES SHOW VAR	Shows all defined variables	Unsupported: Use the PRINT command with no arguments
SHUTDOWN <i>dbname</i>		Unsupported

Obsolete Functionality

<i>Feature</i>	<i>SQL*DBA</i>	<i>Server Manager</i>
SET ARRAYSIZE	Changes the size of each fetch call	No longer used: Server Manager automatically chooses optimum fetch size.
SET CYCLE	Sets the monitor interval	No longer used
SET/SHOW DISP-WIDTH	Line mode	No longer used
SET HISTORY	Screen mode	No longer used
SET LINES	Screen mode	No longer used
SET/SHOW LOG-WIDTH	Line mode	No longer used
SET TERM	Screen mode	No longer used

Operational Differences

<i>Feature</i>	<i>SQL*DBA</i>	<i>Server Manager</i>
Executing a SQL script from the operating system command line	>sqldba @script.sql	>svrmgrl command = @script.sql
Executing a server command from the operating system command line	Unsupported	>svrmgrl command="connect x/y@z"
Entering a Server Manager command spanning multiple lines	Unsupported	Type a backslash (\) at the end of each line to indicate continuation. SVRMGR> connect \ 2> myrdb/password
Entering a quoted string spanning multiple lines		Type a backslash (\) at the end of each line to indicate continuation. SVRMGR> select * 2> from emp 3> where ename = 'A\ 4> BCDEF';
Exiting the application	SQLDBA> exit	SVRMGR> exit or SVRMGR> quit
Aborting a command	Hit <Return> twice or Enter a period (.) by itself on a line.	Enter a period (.) by itself on a line. This procedure applies to Server Manager 1 Line Mode only. SVRMGR> delete from emp 2> . SVRMGR>
Executing the current SQL statement		Enter a slash (/) by itself on a line. SVRMGR> select * 2> from emp 3> /

Index

A

- ABORT mode, B-29
- acknowledging events
 - Event menu, 5-12
 - Event window, 5-9, 5-23
- Active Jobs page, 4-6
- add a group to a map
 - Map menu, 3-6
- adding privileges to a role, 9-22
- adding privileges to a user, 9-15
- adding roles to a user, 9-15
- adding subroles to a role, 9-22
- additional administration utilities, 1-9
- additional services folders
 - Navigator tree, 2-2, 2-3
- administration task
 - Navigator objects, 2-15
- Administrator Toolbar, 1-7
 - application and float styles, 1-15
 - customize, 1-14
 - database list, 1-14
 - exiting, 1-15
 - menu options, 1-14
 - online help, 1-15
 - status of the default database, 1-15
- administrators
 - setting up in Event system, 5-24
- Advanced Mode
 - View menu, 1-12
- Advanced mode, 7-17
- Advanced Queuing, 11-40
- agent
 - down, 5-4
 - intelligent, 6-2
 - out of sync with Console, 6-11
 - process, 6-2
 - UpDown Node event, 5-4
- Alter Rollback Segment property sheet
 - about, 8-22
 - Storage page, 8-22
- Application Pending Notifications
 - Daemon Manager, 6-6
- apply (register) an event set, 5-19
- applying edits, 10-13
- ARCH background process
 - starting, B-5
- ARCHIVE LOG command, B-6
 - about, B-3
 - ALL option, B-5
 - LIST option, B-4
 - NEXT option, B-5
 - START option, B-5
 - STOP option, B-5
- Archived Logs Property Sheet
 - General Page, 12-19
 - Record Page, 12-19
- archiving redo log files
 - about, B-3
 - archive location, B-4
 - device written to, B-6
 - re-archiving, B-5
 - stopping, B-5
- Array Type
 - about, 11-10
 - create like, 11-12
 - editing, 11-12
- Array Type property sheet, 11-11

- Array Types, about, 11-10
- audience for this guide, xxiv
- Auto Extend page, 8-17
- AUTORECOVERY option, B-21
- Available Destinations list
 - Job system, 4-11

B

- background processes
 - startup after abnormal termination, B-30
- backing up a tablespace, 12-50
- Backup Manager
 - menus, 12-38
 - starting, 12-2
- Backup Manager, about, 7-5
- Backup menu, 12-39
- Backup Set Property Sheet
 - Channels Page, 12-27
 - Parameters Page, 12-27
- Backup Set/Image Copy Property Sheet
 - Schedule Page, 12-28
- Backup Tablespace Wizard
 - about, 12-50
- Backup Wizard, 12-24
- BackupSet Property Sheet
 - General Page, 12-26
- Basic mode, 7-17
- Broadcast Message task, 4-22

C

- Catalog, 15-9
- Catalog window, 15-3, 15-9
- Channels Property Sheet
 - General Page, 12-20, 12-35
 - Limits Page, 12-21, 12-36
- channels, creating, 12-35
- CHAR datatype
 - displaying in line mode, B-21
 - in line mode, B-26
- Cluster Property Sheet
 - General Page, 11-14
- Clusters
 - editing, 11-16

- Clusters folder
 - about, 11-12
- Clusters Multi-column List, 11-13
- Clusters, creating, 11-13
- Column Details, 11-15
- command buttons
 - dialog box, 7-16
 - property sheet, 7-19
- command history
 - about, 13-12
 - and Worksheet menu commands, 13-13
 - retrieving a command from, 13-13
- Command History dialog box, 13-12
- commands
 - @ command, B-2
 - ARCHIVE LOG command, B-3
 - CONNECT command, B-6
 - DESCRIBE command, B-9
 - DISCONNECT command, B-11
 - EXECUTE command, B-12
 - executing, 13-10
 - executing in SQL worksheet, 13-8
 - EXIT command, B-12
 - PRINT command, B-13
 - RECOVER command, B-14
 - REMARK command, B-17
 - SET command, B-18
 - SHOW command, B-25
 - SHUTDOWN command, B-29
 - SPOOL command, B-30
 - STARTUP command, B-31
 - VARIABLE command, B-34
- commands (Server Manager)
 - @ command, in a SQL Worksheet, 13-15
 - @@ command, 13-15, B-3
 - programs, setting characteristics of, B-18
- Comments
 - in scripts, B-17
- commit
 - forcing, 10-21
- common services
 - and agents, 1-6
 - Enterprise Manager, 1-6
- communication
 - Enterprise Manager, 1-6

- Communication Daemon, 6-4
 - Enterprise Manager, 1-6
- communication daemon
 - parameters, 6-7
 - populating Navigator, 2-3
- composite job, 4-3
- configure
 - mail services, 5-14
 - modem, 5-14
 - paging services, 5-14, 5-28
- CONNECT command
 - about, B-6
 - INTERNAL option, B-7
 - user SYS, B-8
- Connect dialog box, 1-23
- connect string
 - Enterprise Manager, 1-19
- connecting
 - as SYSOPER or SYSDBA, 1-24
 - File menu item, 7-10, 13-5
 - Navigator menu item, 2-6
 - service name, 1-24
 - to a database, 1-23
 - to a database instance in the Navigator, 2-6
 - to a repository, 1-19
- connecting to a database, B-6
- connections
 - changing, in a SQL Worksheet, 13-10
 - multiple, 1-25
- Console
 - customize layout, 1-11
 - Enterprise Manager, 1-3
 - File menu, 1-11
 - Help menu, 1-12
 - menus, 1-10
 - user preferences, 1-25
 - using menus, 1-10
- context-sensitive menus
 - Navigator, 1-10, 2-7
 - using, 7-13
 - using in the Navigator, 1-10
- Controlfile Object and Property Sheet, 12-43
- copying tree objects
 - Navigator, 2-16
- create
 - event set, 5-13
 - job, 4-8
 - map, 3-6
 - or modify a job, 4-9
 - or modify an event set, 5-14
 - specified object type in the Navigator, 2-5
- Create Cluster property sheet, 11-13
- Create Datafile
 - General Page, 8-16
- Create Datafile property sheet
 - Auto Extend page, 8-17
- Create Index property sheet
 - Storage page, 11-24
- create like
 - profile, 9-28
 - role, 9-20
 - user, 9-14
- Create Like Cluster
 - Clusters
 - create like, 11-16
- Create package body property sheet
 - General page, 11-34
- Create Procedure property sheet
 - General page, 11-39
- Create Profile property sheet, 9-25
 - profile name, 9-25
- Create Redo Log Group property sheet
 - about, 12-50
- Create Redo Log Member property sheet
 - about, 12-51
- create refresh group property sheet
 - General page, 11-50
 - Snapshot page, 11-51
- Create Role property sheet
 - General page, 9-20
- Create Rollback Segment property sheet, 8-20
 - about, 8-19
- create sequence property sheet
 - General page, 11-53
- Create Snapshot Log property sheet
 - Storage page, 11-56
- Create Snapshot property sheet
 - Cluster page, 11-62
 - General page, 11-61
 - Storage page, 11-60

- Create Synonym property sheet
 - General page, 11-65
 - Create Table property sheet
 - Column Details property sheet, 11-73
 - General page, 11-70
 - Storage page, 11-69, 11-73
 - Create Tablespace
 - Extents page, 8-10
 - General page, 8-8
 - Create Trigger property sheet
 - General page, 11-77
 - Timing page, 11-78
 - Create User property sheet, 9-10
 - about, 9-8
 - adding quotas, 9-13
 - Default Roles page, 9-13
 - General page, 9-9
 - Privileges page, 9-10
 - Roles/Privileges page, 9-10
 - Create View property sheet
 - General page, 11-81
 - Creating a Channel, 12-35
 - cursors
 - disconnect and closing cursors, B-11
 - opened at CONNECT, B-8
 - customize and manage network views
 - Map system, 3-6
- ## D
-
- Daemon Manager
 - configuration parameters, 6-7
 - menu, 6-11
 - Data Manager
 - context-sensitive menus, 14-4
 - Export wizard, 14-5
 - Import wizard, 14-6
 - Load wizard, 14-7
 - menus, 14-4
 - starting, 14-2
 - window, 14-3
 - Data Manager, about, 7-6
 - database
 - changing connections, 13-10
 - mounting, 10-10, B-32
 - opening, 10-10, B-32
 - recovering from SQL worksheet, 13-11
 - shutting down, 10-10
 - database administration applications
 - DBA, 1-12
 - Database Fault Management Events, 5-31
 - database files
 - recovering, B-14
 - Database folder
 - mounting a database, 10-10
 - opening a database, 10-10
 - shutting down a database, 10-10
 - database instance property sheet
 - accessing from a map, 3-5
 - Database Link
 - create like, 11-19
 - database link folder
 - creating a database link, 11-17
 - Database Link Property Sheet
 - General Page, 11-18
 - Database Links folder
 - about, 11-16
 - Database Links Multi-column List, 11-17
 - Database menu, 12-40
 - database name
 - at startup, B-31
 - Database Property Sheet
 - Information Page, 12-11, 12-42
 - Status Page, 12-10, 12-41
 - Database Resource Management Events, 5-32
 - Database Space Management Events, 5-32
 - database tools, 1-7
 - application folders, 7-3
 - folders, 7-3
 - quitting, 7-10
 - database tools and utilities, 1-7
 - about, 7-2
 - dialog boxes, 7-15
 - property sheets, 7-17
 - Databases folder
 - Navigator tree, 2-2
 - datafiles
 - adding, 8-13
 - altering, 8-18
 - creating, 8-16

- creating like, 8-18
 - multi-column list, 8-15
 - renaming, 8-16
- Datafiles folder
 - about, 8-15
 - altering a datafile, 8-18
 - create like datafile, 8-18
- Datafiles Property Sheet
 - General Page, 12-15
- DATE datatype
 - in line mode, B-22, B-26
- DBA
 - database administration application, 1-7
- DBA mode, B-32
- DBA privilege, 1-20, B-32
- DBA tools
 - menus, 7-10
- Deinstall Product Task, 15-31
- delete
 - group, 3-6
- Delete option
 - Daemon Manager, 6-11
 - removing job or event operations, 6-11
- Delete Package Task, 15-33
- DESCRIBE command, B-9
- destination
 - archive, B-6
 - overriding default, B-6
- dialog box
 - about, 7-15
 - command buttons, 7-16
- dialog boxes, 7-15
- differences
 - between SQL*DBA and Server Manager, C-1
- DISCONNECT command, B-11
- disconnecting, B-11
 - Navigator menu item, 2-6
- discovered nodes
 - saving a list, 1-21
- discovering services
 - communication daemon, 1-4, 2-3
 - Navigator, 2-8
 - SQL*Net network, 2-3
- discovery
 - problems, 2-13

- discovery cache
 - Enterprise Manager, 1-6
 - providing network information, 1-6
- Distribute Package Task, 15-34
- distribution host
 - and the software distribution process, 15-14
 - definition, 15-14
- drag and drop
 - Navigator, 2-16
 - Navigator objects, 2-2
- drawers
 - opening, 7-20

E

- Edit menu, 13-6
- elapsed timing, B-24, B-28
- email
 - configuring, 5-27
- EMS
 - Event Management system, 5-1
- Enterprise Manager
 - components, 1-2
 - connect string, 1-19
 - database administration, 1-7
 - installing, 1-19
 - introduction, 1-2
 - launch palettes, 1-13
 - Oracle releases supported by, 1-2
 - privileges needed to use, 1-20
 - quitting, 1-11
 - repository, 1-19
 - starting, 1-19
 - starting up, 1-19
 - toolbars, 1-13
- event
 - destination requirements, 5-20
 - numeric pager event Ids, 5-33
 - set, 5-3
- Event History page
 - Event window, 5-10
- Event Management, 5-1
 - categories of events, 5-6
 - configuring paging services, 5-28
 - creating an event set, 5-14

- defining administrators, 5-24
- Enterprise Manager, 1-5
- fault management, 5-6
- introduction, 5-1
- menu, 5-12
- modifying a registered event set, 5-19
- modifying an event set, 5-14
- Oracle events, 5-31
- parameter settings, 5-17
- performance management, 5-7
- process, 5-2
- registering an event set, 5-19
- resource management, 5-7
- space management, 5-7
- types of events, 5-6
- viewing a registered event set, 5-19
- viewing an event set, 5-11
- window, 5-8
- Event Set Library page
 - Event window, 5-11
 - predefined event sets, 5-11
- Event Set Management property sheet
 - creating an event, 5-13
 - modifying a new event, 5-13
- example
 - creating an event set, 5-17
 - group, 3-13
 - Job Scheduling, 4-15
 - Map, 3-9
 - Navigator, 2-17
 - registering an event, 5-22
- EXECUTE command, B-12
- EXIT command, B-12
- exiting an application, 7-10
- exiting Enterprise Manager, 1-11
- expanding
 - container in Navigator, 2-15
 - map objects, 3-4
 - possible delays with a container, 2-15
- Export
 - predefined task, 14-8
- exporting
 - a full database to a file, 14-8
 - data to a file, 14-5
 - objects to a file, 14-8

- owners to a file, 14-8
- tables to a file, 14-8

F

- Fault Management events, 5-6
- fetching rows in a query
 - setting MAXDATA, B-23
- File menu, 13-5
 - DBA tools, 7-10
- filter
 - Navigator, 2-4
 - Navigator tree, 2-7
- filtering folders
 - Navigator, 2-18
- fixit job, 4-3, 5-1
 - creating, 4-10
 - events, 5-16
 - removing, 5-16
 - submitting, 5-16
 - with events, 5-1
- folder pull-down menus
 - about, 7-5
- folder tabs, 7-20
- folders
 - opening, 7-4
 - pull-down menu, 7-5
- FORCE option, B-32
- forcing the retry of operations
 - Daemon Manager, 6-11
- Function
 - altering, 11-22
 - create like, 11-21
- Function Property Sheet
 - General Page, 11-21
- Functions folder
 - about, 11-19
 - creating a function, 11-20

G

- General page
 - job, 4-10
- group
 - add or delete, 3-11

- bitmap background, 3-12
- create, 3-10
- example, 3-13
- job destination, 4-11
- modify, 3-12
- user-defined, 3-3
- Groups folder
 - Navigator tree, 2-2

H

- Help menu, 7-12
- Help system
 - using, 1-8, 1-12, 7-12
- Host Folders icon, 15-4
- host name, B-7

I

- Image Copy Property Sheet
 - General Page, 12-30
 - Schedule Page, 12-30
- IMMEDIATE mode, B-30
- Import
 - predefined task, 14-8
- importing
 - a full database from a file, 14-8
 - data from a file, 14-6
 - tables from a file, 14-8
 - user objects, 14-8
- Index
 - altering, 11-26
 - create like, 11-26
- Index Property Sheet
 - General Page, 11-25
- Indexes folder
 - creating an index, 11-24
- Indexes Multi-column List, 11-23
- in-doubt transactions
 - forcing a commit, 10-21
 - forcing a rollback, 10-21
- Initialization Parameters
 - about, 10-11
- Initialization parameters
 - displaying, B-27

- initialization parameters
 - editing, 10-13
- INIT.ORA file, B-32
- Install Package Task, 15-29
- installing
 - Enterprise Manager, 1-19
 - Intelligent Agent, 6-2
- Instance Manager, about, 7-6
- instances
 - connecting, 1-23, B-6
 - default, B-26
 - displaying current, B-26
 - shutting down, B-29
 - starting, B-31
- Intelligent Agent, 1-6, 6-2
 - Enterprise Manager, 1-6
 - with events, 5-4
- IP address
 - listening address, 6-8

J

- job
 - alter, 4-8
 - composite, 4-3
 - create, 4-8
 - manage, 4-8
 - modify, 4-8
 - view, 4-8
- Job History page, 4-7, 4-8
- Job Library, 4-7
- job output
 - Progress Page, 4-15
- Job Scheduling, 4-1
 - credential preference, 4-3, 4-12
 - destination requirements, 4-11
 - Enterprise Manager, 1-4
 - example, 4-15
 - Failed to Authenticate User error, 4-3
 - General page, 4-10
 - job destination, 4-10
 - job destinations, 4-11
 - job name, 4-10
 - job tasks, 4-11
 - job type, 4-10

- managing jobs, 4-4
- menu, 4-8
- saving history, 4-8
- scheduling jobs, 4-4
- scheduling process, 4-2
- task parameters, 4-12
- window, 4-5

Job Scheduling system, 15-24

Job system

- and the software distribution process, 15-21
- tasks used with Oracle Software Manager, 15-21

Jobs Property Sheets

- Backup Manager, 12-22

K

keystrokes

- Navigator, 2-19

L

launch palettes, 1-13

launching a database tool, 1-22

- from a map, 3-5
- from the Console, 2-4
- from the Navigator, 1-4
- Navigator, 2-4
- with the right-mouse button, 2-4

line mode

- about, A-1
- entering commands in, A-2
- entering PL/SQL in, A-2
- entering SQL in, A-2
- running scripts from, A-2
- Server Manager, 1-8
- starting, A-2
- using, A-2

list of saved maps

- Map menu, 3-6

LIST option, B-4

Listener Fault Management Events, 5-32

Listeners folder

- Navigator tree, 2-2

Listening Address

- communication parameter, 6-8

loading

- data from an external file, 14-7

Log menu, 7-12

log sequence numbers

- for manual archiving, B-5
- listing, B-4

LOG_ARCHIVE_DEST parameter

- ARCHIVE LOG LIST, B-4
- overriding, B-6

LOG_ARCHIVE_START parameter, B-6

- ARCHIVE LOG command, B-5

Logfile menu, 12-7

login information, 1-25

M

mail

- configuring, 5-27

mail engine

- event system notification, 5-27

Management Information Base (MIB)

- with SNMP, 6-3

manipulation of the objects

- in the Navigator, 2-15

Manual Backup Operations, 12-26

Map, 3-1

- bitmap background, 3-10
- create, 3-2
- creating with dragging and dropping, 3-2
- customized views, 3-1
- example, 3-9
- expanding objects, 3-4
- modify, 3-8
- status of objects, 3-4
- types of views, 3-3
- user-defined views, 3-3
- window, 3-2

MAPI

- mail engine, 5-27

menus

- Console, 1-10
- context-sensitive, 7-13
- context-sensitive, using, 7-13
- File menu, 1-11, 7-10
- folder pull-down, about, 7-5

- Help menu, 1-12, 7-12
- Job, 4-8
- Security Manager, 9-4
- View, 1-11
- View menu, 7-11
- modem
 - changing settings, 5-14
- modify
 - event set, 5-13
 - group, 3-6
 - map, 3-6
 - registered events, 5-13
- monitoring
 - tools, 1-9
 - utilities, 1-9
- mounting a database, B-32
- multiple database connections, 1-25
- multiple tasks
 - submitting a job, 4-9
- multi-selection
 - Navigator, 2-2

N

- Navigator, 2-1
 - connecting to a database, 2-6
 - Enterprise Manager, 1-4
 - example, 2-17
 - filtering, 2-4
 - filtering folders, 2-18
 - keystrokes, 2-19
 - manipulating objects, 2-15
 - menu, 2-5
 - multi-selection, 2-2
 - new view, 2-6
 - objects, 1-4
 - splitting the tree, 2-4, 2-6
 - user-defined map groups, 3-3
 - window, 2-2
- navigator objects
 - manipulating, 7-4
- Navigator window, 15-3, 15-4
- network, 15-22
 - container, 2-2
 - encryption, 6-3

- establishing communication, 15-22
 - flow of information, 15-23
 - specifying service names, 1-24
- Networking Polling Timer
 - communication parameter, 6-8
- new view
 - Navigator tree, 2-6
- NLS environment, 6-3
- Node Fault Management Events, 5-32
- Node Heartbeat Interval
 - communication parameter, 6-9
- Node Heartbeat Timer
 - communication parameter, 6-9
- Nodes folder
 - Navigator tree, 2-2
- NORMAL mode, B-30
- NUMBER datatype
 - in line mode, B-23, B-27
- Number of Cached Connections
 - communication parameter, 6-8
- Number of Worker Threads
 - communication parameter, 6-9

O

- object lists
 - about, 7-4
 - refreshing, 7-11
 - selecting from, 7-13
 - sorting, 7-14
- Object Type property sheet, 11-28
- Object Types
 - about, 11-26
 - adding attributes, 11-29
 - adding methods, 11-30
- online help, 1-8, 1-12, 7-12
 - accessing, 1-8
- online redo log files
 - re-archiving, B-5
- opening a database, B-32
- Operation and Notification Queries folder, 6-6
- Operation Retry Timer
 - communication parameter, 6-9
- Oracle Client Software Manager (OCSM), 15-7
- Oracle Daemon Manager, 6-6

- Oracle DB Fault
 - event profile, 5-12
- Oracle DB Resource
 - event profile, 5-12
- Oracle DB UpDown
 - event profile, 5-12
- Oracle Enterprise Manager, job system, 15-21
- Oracle Enterprise Manager/Oracle Software
 - Manager network, 15-22
 - establishing communication, 15-22
 - flow of information, 15-23
- Oracle events, 5-31
- Oracle Expert, 1-9
- Oracle File Packager, 15-15
- Oracle Host UpDown
 - event profile, 5-12
- Oracle Installer, 15-31
- Oracle Listener UpDown
 - event profile, 5-12
- Oracle Lock Manager, 1-9
- Oracle Multi-Protocol Interchange
 - and Communication Daemon, 6-4
- Oracle Parallel Servers
 - administering, 1-2
- Oracle Performance Pack, 1-9
- Oracle Remote Operations
 - and Communication Daemon, 6-4
- Oracle Secure Network Services (SNS)
 - and Communication Daemon, 6-4
- Oracle Server releases
 - supported by Enterprise Manager, 1-2
- Oracle Software Manager, 15-1
 - and Oracle Enterprise Manager, 15-22
 - Catalog window, 15-10
 - distribution host, 15-14
 - drag and drop behavior, 15-11
 - Host Folder Properties, 15-5
 - Host Folders, 15-4
 - Hosts folder, 15-4
 - network, 15-22
 - Oracle Home Properties, 15-6
 - Oracle Homes folder, 15-4
 - Preferences, 15-10
 - PRODINFO.ORA, 15-16
 - Product Properties, 15-7

- release, 15-9
- Releases folder, 15-4
- response file, 15-20
- software package, 15-9
- Software Packages folder, 15-4
- World folder, 15-4
- Oracle Tablespace Map, 1-9
- Oracle TopSessions, 1-9
- Oracle Trace, 1-9
- oratab, 2-6
- OraTcl
 - event scripts, 5-3
 - job scripts, 4-2
- Output Dialog box
 - Job Scheduling, 4-15
- Outstanding Events page
 - Event system, 5-9

P

- Package
 - create like, 11-37
 - editing, 11-38
- Package Bodies folder
 - about, 11-33
 - creating a package body, 11-34
- Package Bodies Multi-column List, 11-33
- Package Body
 - create like, 11-35
 - editing, 11-35
- Packages folder
 - about, 11-36
 - creating a package, 11-36
- Packages Multi-column List, 11-36
- Packages Property Sheet
 - General Page, 11-37
- page tabs, 7-17
- pages, 7-17
- paging
 - communication settings, 5-30
 - configuration, 5-28
 - GSM protocol, 5-29
 - numeric event Ids, 5-33
 - pager type, 5-29
 - settings, 5-28

- TAP protocol, 5-29
- Paging Services Configuration
 - Event menu, 5-14
- parameter files (INIT.ORA files)
 - specifying alternate, B-32
- parameters
 - communication daemon, 6-7
 - Daemon Manager, 6-7
 - events, 5-16
 - jobs, 4-12
- Parameters Page
 - job, 4-12
- passwords
 - CONNECT command, B-7
- Performance Management events, 5-6
- performance monitors, 1-9
- pinging
 - checking node with pending operations, 6-12
- platform-specific installation documentation, 6-2
- Pre-defined Profiles
 - Event system, 5-12
- preferred credentials
 - connecting in a map, 3-5
 - connecting to a database, 2-15
 - Console, 1-25
- PRINT command, B-13
- privileges
 - needed to use Enterprise Manager, 1-20
 - system, 1-20
- Privileges page, 9-10
- problem
 - job or event operations, 6-11
- problems
 - Navigator discovery, 2-13
- Procedure
 - altering, 11-40
 - create like, 11-40
- Procedures folder
 - about, 11-38
 - creating a procedure, 11-39
- Procedures Multi-column List, 11-38
- profile
 - create like, 9-28
- Profile Container
 - about, 9-24

- Profile Dependencies
 - showing, 9-29
- profiles
 - altering, 9-28
 - assigning to a user, 9-29
 - creating, 9-25
- Profiles folder
 - altering a profile, 9-28
 - creating a profile, 9-25
 - removing a profile, 9-29
- programs
 - setting characteristics of, B-18
- Progress page
 - job, 4-14
- Property sheet
 - Create Tablespace, 8-7
- property sheet
 - command buttons, 7-19
- property sheet, in-doubt transaction, 10-20
- property sheets, 7-17
 - about, 7-17
 - page tabs, 7-17
 - pages, 7-17
- purpose of this guide, xxiv

Q

- Queue Tables
 - about, 11-40
 - creating, 11-41
- Queues
 - adding subscribers, 11-46
 - creating, 11-45
- Queues Folder, 11-40
- Queues folder
 - about, 11-40
- Queues Multi-column List, 11-41
- Quick Edit Tablespace property sheet
 - General page, 8-12
- quitting an application, 7-10
- quitting Enterprise Manager, 1-11
- Quotas page, 9-10

R

- RECOVER command, B-14
- Recover Property Sheet
 - General Page, 12-33
- recovery
 - RECOVER command, B-14
 - SQL worksheet and, 13-11
- Recovery Catalog, 12-4
 - maintaining, 12-34
 - registering, 12-5
- redo log files
 - archiving, B-4
 - listing range of, B-4
 - re-archiving, B-5
- redo log group
 - adding new, 12-50
- redo log group member
 - adding to an existing redo log group, 12-51
- Redo Log Group Property Sheet, 12-46
 - General Page, 12-17
- Redo Log Groups container
 - about, 12-16, 12-45
 - multi-column list, 12-45
 - property sheet, 12-46
- Redo Log Members container
 - property sheet, 12-47
- refresh command
 - using, 7-11
- Refresh Group
 - altering, 11-51
- Refresh Group folder
 - creating a refresh group, 11-49
- Refresh Groups folder
 - about, 11-48
- Refresh Groups Multi-Column List, 11-49
- Register User at Startup
 - communication parameter, 6-9
- registered destination
 - map objects, 3-4
- registering events
 - Event menu, 5-13
- Registrations page
 - Event window, 5-10
- related publications, xxx
- release
 - and the software distribution process, 15-15
 - definition, 15-15
- Releases Catalog window, 15-9
- REMARK command, B-17
- removing
 - a split page in the Navigator, 2-6
 - job with Remove menu item, 4-8
 - problems with a job, 4-8
 - problems with an event, 5-13
 - registered events, 5-13
- removing privileges from a role, 9-23
- removing privileges from a user, 9-17
- removing roles from a user, 9-17
- removing subroles from a role, 9-23
- Report Wizard, 12-25
- repository
 - connecting to, 1-19
 - creating and upgrading, 1-20, 1-21
 - dropping, 1-21
 - Enterprise Manager, 1-5
 - information, 1-6
 - tables, 1-5
 - unique username, 1-20
- Repository Manager Wizard, 1-21
- resetting defaults, B-23
- resetting edits, 10-13
- Resource Management events, 5-6
- Response File Generator, 15-37
- Restore and Recover Operations, 12-31
- Restore Property Sheet
 - Channels Page, 12-32
 - General Page, 12-31
 - Rename Page, 12-31
- Restore Wizard, 12-24
- right mouse button
 - Navigator tree, 2-7, 2-16
 - used in the Navigator, 1-10
- right-mouse button
 - Related Tools menu, 2-4
- role
 - create like, 9-20
- Role Container
 - about, 9-18
- roles

- altering, 9-21
- creating, 9-19
- DBA and RESOURCE, 9-16
- granting UNLIMITED TABLESPACE, 9-16

Roles folder

- adding privileges or roles to a role, 9-22
- altering a role, 9-21
- creating a role, 9-19
- removing a role, 9-21
- removing privileges or subroles from a role, 9-23

rollback

- forcing, 10-21

Rollback folder

- about, 8-19
- altering a rollback segment, 8-22
- creating a rollback segment, 8-19
- object list, 8-19
- pending offline, 8-23
- placing rollback segment online, 8-23
- removing a rollback segment, 8-22
- taking rollback segment offline, 8-23

rollback segments

- altering, 8-22
- creating, 8-19
- placing online, 8-23
- removing, 8-22
- taking offline, 8-23

rolling forward, B-14

Run DBA Script task, 4-19

Run OS Command task, 4-22

Run SQL*Plus task, 4-18

Run Tcl Script task, 4-22

running scripts

- @ command, 13-15, B-3
- @@ command, 13-15, B-3

S

saving

- event history, 5-10
- job history, 4-7, 4-8
- SQL worksheet, 13-14

schedule

- and manage jobs, 4-4

- execution of a job, 4-13

- job, 4-4

- Schema Manager, about, 7-7

- Schema Menus, 11-4

- Schema Objects folder

- about, 11-6

scripts

- running, B-3

- running from line mode, A-2

- running from SQL worksheet, 13-15

- Security Manager, about, 7-8

- SELECT command

- MAXDATA, B-23

- Send Us Your Comments, xxi

Sequence

- altering, 11-55

- create like, 11-54

Sequences folder

- about, 11-51

- creating a sequence, 11-53

- Sequences Multi-Column List, 11-52

Server Manager

- line mode operation, 1-8

- starting, 1-8

Server Manager Line Mode

- platforms, A-1

Server Manager/LineMode

- about, A-1

SERVEROUTPUT

- SET, B-24

- SHOW, B-27

- Service Discovery, 1-4, 2-3

- Navigator menu option, 2-6

Service Discovery Timer

- communication parameter, 6-8

- Service Name, 1-25

service name

- in window titles, 1-25

- specifying, 1-24

- Service Type, 1-25

- services and agents, 6-1

Sessions folder

- about, 10-17

- disconnecting user session, 10-18

- restricting sessions, 10-19

- SET command
 - about, B-18
 - INSTANCE option, B-26
 - SERVEROUTPUT, B-24
- setting a bitmap background
 - Map menu, 3-7, 3-8, 3-10, 3-12
- setting up administrators, 5-24
- short-cut menu, 1-10
- SHOW command, B-25
- Show SQL button, 7-19
- showing SQL statements, 7-19
- shrinking a rollback segment, 8-22
- SHUTDOWN command, B-29
- Shutdown Database task, 4-20
- Shutdown Listener task, 4-24
- shutting down a database, 10-10, 12-48
- signal flag
 - on map objects, 3-4
- SMTP
 - mail engine, 5-27
- Snapshot
 - create like, 11-62
 - editing, 11-63
- Snapshot folder
 - creating a snapshot, 11-60
- Snapshot Log
 - altering, 11-58
 - create like, 11-57
- Snapshot Log folder
 - creating a snapshot log, 11-56
- Snapshot Log property sheet
 - General page, 11-57
- Snapshot Logs folder
 - about, 11-55
- Snapshot Logs Multi-column List, 11-56
- Snapshots folder
 - about, 11-59
- Snapshots Multi-column List, 11-59
- SNMP
 - and agents, 6-3
- software distribution process, 15-12
 - about, 15-12
 - creating releases, 15-15
 - creating software packages, 15-18
 - specifying distribution hosts, 15-14
 - using the Oracle Enterprise Manager job system, 15-21
- Software Manager, 15-1
- Software Manager, about, 7-8
- software package
 - and the software distribution process, 15-18
 - definition, 15-18
- Software Packages Catalog window, 15-9
- Space Management event, 5-6
- specifying service names, 1-24, 1-26
- split (new) view
 - Navigator, 2-6
- splitting the Navigator tree, 2-4, 2-6
 - new view, 2-6
- SPOOL command, B-30
 - long lines of text, B-31
- spooling multiple log files, B-6
- spreadsheets, 7-18
 - DBA tools, 7-18
- SPX
 - listening address, 6-8
- SQL
 - showing generated, 7-19
- SQL error
 - job task, 4-19
- SQL statements
 - showing, 7-19
- SQL Worksheet
 - about, 13-2
 - connecting from, 13-10
 - Edit menu, 13-6
 - entering commands, 13-8
 - executing commands, 13-10
 - executing commands in, 13-8
 - File menu, 13-5
 - input pane, 13-8
 - recovering a database, 13-11, 13-13
 - running scripts from, 13-15
 - starting, 13-2
 - using, 13-2
 - Worksheet menu, 13-6
- SQL worksheet
 - saving work, 13-14
- SQL Worksheet, about, 7-9
- SQL*DBA and Server Manager

- differences, C-1
- SQL*Net network
 - services, 2-3
- starting Enterprise Manager, 1-19
- starting up a database, 10-9, 12-49
- STARTUP command, B-31
- Startup Database task, 4-20
- Startup Listener task, 4-25
- status
 - of a map object, 3-4
 - of a registered event set, 5-3
- Storage Manager, about, 7-8
- stored configurations, 10-14
 - adding parameters, 10-15
 - creating, 10-14
 - creating a file from, 10-16
 - deleting, 10-15, 10-16
 - editing, 10-15
- stored configurations property sheet, 10-14
- Stored Scripts, creating, 12-37
- summary of the pending jobs
 - Active Jobs page, 4-6
- supported Oracle releases, 1-2
- Switch Property Sheet
 - General Page, 12-33
- Synonyms folder, 11-22
 - about, 11-64
 - creating a synonym, 11-64
- Synonyms Multi-column List, 11-64
- SYS username
 - CONNECT INTERNAL, B-8
- SYSDBA
 - role, 3-5
- SYSDBA roles
 - connect as, 1-24
- SYSOPER
 - role, 3-5
- SYSOPER role
 - connect as, 1-24
- System Global Area (SGA)
 - displaying, B-28

T

Table

- altering, 11-75
- create like, 11-75
- dropping constraints, 11-74
 - Removing a Constraint Definition, 11-75
- Table Property Sheet
 - Cluster Columns, 11-74
- Table Types
 - about, 11-66
 - creating, 11-66
- tables
 - DESCRIBE command, B-10
- Tables folder, 11-68
 - creating a table, 11-69
- Tables Multi-column List, 11-68
- Tablespace
 - editing, 8-12
 - menu, 8-4
- Tablespace container
 - multi-column list, 8-7, 12-16
- Tablespace containers
 - about, 12-12, 12-43
- Tablespace Dependencies
 - displaying, 8-11
- Tablespace folder
 - about, 8-7
- Tablespace Property Sheet
 - Extents Page, 8-10
 - General Page, 8-8, 12-13, 12-44
- Tablespaces
 - dropping, 8-12
- tablespaces
 - adding a datafile to, 8-13
 - adding rollback segments, 8-13
 - backing up, 12-50
 - creating, 8-7
 - making read-only, 8-14
 - making writeable, 8-15
 - placing online, 8-14
 - recovering, B-14
 - taking offline, 8-14
- Tablespaces container
 - multi-column list, 12-12, 12-43
- Tablespaces folder
 - taking tablespace offline, 8-14
- tape drive

- spooling multiple log files, B-6
- task parameters
 - Job system, 4-12
- Tcl
 - event scripts, 5-3
 - job scripts, 4-2
 - script examples, 4-23
 - scripts executed by agents, 6-3
- TCP/IP
 - listening address, 6-8
- TCP/IP Port
 - communication parameter, 6-8
- terminals
 - redirecting output, B-24
- third-party applications
 - notifications, 6-6
- third-party events
 - creating, 5-15
 - with job task, 4-24
- timing
 - CPU and elapsed, B-24, B-28
- tnsnames.ora, 2-6, 2-10, 2-12, 2-13, 2-14
- toolbar
 - icons, 1-13
 - options, 1-13
- Tools menu, 1-12
- topology file
 - topo_ops.ora, 1-4, 2-3
- transactions
 - committing at disconnect, B-11
- Transactions folder
 - about, 10-19
 - forcing a commit, 10-21
 - forcing a rollback, 10-21
 - object list, 10-19
- Transparent Network Substrate (TNS)
 - and Communication Daemon, 6-4
- Trigger
 - altering, 11-79
 - create like, 11-79
- Triggers folder
 - about, 11-75
 - creating a trigger, 11-76
- Triggers Multi-column List, 11-76
- troubleshooting

- job and event operations, 6-11

U

- UpDown (Database) event, 5-32
- UpDown (Listener) event, 5-32
- UpDown (Node) event, 5-32
- User
 - create like, 9-14
- user preferences
 - jobs, 4-3, 4-12
 - when running jobs, 4-3, 4-12
- user-defined groups
 - in the Navigator, 2-4
 - job destinations, 4-11
 - Navigator tree, 2-4
- user-defined map views
 - creating, 3-3
- Username, 1-25
- usernames, B-7
- Users
 - menu, 9-4
- users
 - altering, 9-14
 - assigning profile to, 9-29
 - creating, 9-8
- Users Folder
 - about, 9-7
- Users folder
 - adding privileges or roles to a user, 9-15
 - altering a user, 9-14
 - creating a user, 9-8
 - dropping a user, 9-15
 - removing privileges or roles from a user, 9-17, 9-18

V

- VARIABLE command, B-34
- View
 - altering, 11-82
 - create like, 11-82
- View menu, 7-11
- Views folder
 - about, 11-79

creating a view, 11-80
Views Multi-column List, 11-80

W

window titles, 1-25

windows

 Catalog, 15-9

 Navigator, 15-4

 Releases Catalog, 15-9

 Software Packages, 15-9

Worksheet menu, 13-6

World folder icon, 15-4

