

Oracle8i

Getting to Know Oracle8i

Release 2 (8.1.6)

December 1999

Part No. A76962-01

ORACLE

Getting to Know Oracle8i, Release 2 (8.1.6)

Part No. A76962-01

Copyright © 1996, 1999, Oracle Corporation. All rights reserved.

Primary Author: Michele Cyran

Contributing Authors: Ruth Baylis, Paul Lane, Mark Niebur

Contributors: Lance Ashdown, Cathy Baird, Mark Bauer, Eric Belden, D. Yitzik Brenman, Dave Colello, Mary Ann Davidson, Sandy Dreskin, John Frazzini, Joe Garcia, Gilbert Gonzalez, Michael Hartstein, Jeff Hebert, Shelley Higgins, Mark Kennedy, Thomas Kurian, Anna Logan, David Mathews, Colin McGregor, Jack Melnick, Tom Portfolio, Den Raphaely, John Russell, Nasser Shariatpanahy, Deborah Steiner, Mark Townsend, Randy Urbano, V. Vijayanandan, Rod Ward, Steven Wertheimer

Graphic Designer: Valarie Moore

The Programs (which include both the software and documentation) contain proprietary information of Oracle Corporation; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs is prohibited.

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

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

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

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

Oracle is a registered trademark, and ConText, JDeveloper, JInitiator, JServer, LogMiner, Net8, Oracle Call Interface, Oracle COM Cartridge, Oracle Data Migration Assistant, Oracle Database Assistant, Oracle Database Configuration Assistant, Oracle DBA Management Pack, Oracle Designer, Oracle Developer, Oracle Enterprise Manager, Oracle Enterprise Manager Performance Pack, Oracle Exchange, Oracle Expert, Oracle iFS, Oracle *interMedia*, Oracle Lite, Oracle Parallel Server, Oracle Provider for OLE DB, Oracle Spatial, Oracle Virtual Private Database, Oracle Visual Information Retrieval (VIR), Oracle Web Application Server, Oracle WebDB, Oracle7, Oracle8, Oracle8 Enterprise Edition, Oracle8i, Oracle8i Lite, PL/SQL, Pro*Ada, Pro*C, Pro*C/C++, Pro*COBOL, Pro*FORTRAN, Pro*Pascal, Pro*PL/1, SQL, SQL*Loader, SQL*Net, SQL*Plus, and Wallet Manager are trademarks or registered trademarks of Oracle Corporation. All other company or product names mentioned are used for identification purposes only and may be trademarks of their respective owners.

Contents

Send Us Your Comments	xiii
Preface.....	xv
1 Overview	
Introducing Oracle8i	1-2
A Family of Database Products.....	1-4
2 Oracle 8i Release 2 (8.1.6) New Features	
Oracle8i Improvements	2-2
Performance Improvements.....	2-2
Improved Installation and Management	2-3
Improved Quality and Support.....	2-4
Enhanced Support of Internet Computing.....	2-5
Enhanced Java Support.....	2-5
XML Support.....	2-5
PL/SQL Pages.....	2-6
<i>interMedia</i> , Spatial, Time Series, and Visual Image Retrieval	2-6
<i>interMedia</i>	2-6
Oracle Spatial	2-7
Oracle Advanced Security Improvements	2-8
Network Security Enhancements.....	2-8
Single Sign-On.....	2-8
PKI Credential Management	2-8

Enterprise User Management	2-9
User/Schema Separation.....	2-9
Advanced New Functionality	2-10
Data Warehouse Enhancements.....	2-10
Analytic Functions.....	2-10
Materialized View Enhancements	2-11
Enhanced Partitioning	2-11
Query Rewrite	2-12
Object Relational Enhancements	2-12
Other Enhancements	2-13
Release 2 (8.1.6) New Features	2-14
Oracle8i Improvements.....	2-14
Java Enhancements.....	2-15
Enhanced Support of Internet Computing	2-15
Security and Networking	2-15
Advanced New Functionality	2-16
Object Relational	2-16
Other Enhancements	2-16

3 Oracle 8i (8.1.5) New Features

Content Management for the Internet	3-2
WebDB.....	3-2
Oracle Internet File System (iFS)	3-3
<i>interMedia</i> and the Web.....	3-3
Oracle <i>interMedia</i>, Spatial, Time Series, and Visual Image Retrieval	3-5
Oracle <i>interMedia</i>	3-5
Text	3-6
Improved Ease of Use	3-6
Extensible Index.....	3-6
Extensible Query Optimizer	3-6
Extensible Knowledge Base (KB)	3-7
Filters.....	3-7
Hierarchical Query Feedback	3-7
Theme Highlighting.....	3-7
Alternate Spelling.....	3-7

Stoplists.....	3-7
Field Sections	3-7
Audio, Image, and Video	3-8
Summary of Audio, Image, and Video Support.....	3-8
Oracle <i>interMedia</i> Audio, Image, and Video Java Client	3-9
Locator	3-9
Oracle Spatial	3-10
VARRAY Storage Model.....	3-11
Dynamic Window Queries	3-11
Spatial Index Mechanism	3-11
RELATE() Operator.....	3-11
Extended Functionality	3-11
Oracle Visual Information Retrieval.....	3-12
New Extensible Indexing	3-12
Temporary LOB Usage	3-13
Scaling Enhancements	3-13
Proprietary Image Formats	3-13
Oracle Time Series	3-13
Irregular Time Series.....	3-14
Additional Calendar Frequencies	3-14
Flexible Timestamp Precisions	3-14
Flexible Patterns.....	3-14
Time Scaling Enhancements	3-14
Nested Table Support	3-15
Ease of Use.....	3-15
Java	3-15
Oracle8i's Java Virtual Machine	3-16
JServer Accelerator	3-16
Programmatic Interfaces.....	3-17
JDBC Drivers.....	3-17
Server JDBC Driver	3-17
OCI JDBC Driver	3-17
Thin JDBC Driver	3-17
SQLJ Translator.....	3-18
Development Model.....	3-18

Java Stored Procedures	3-18
Enterprise Java Beans	3-19
CORBA	3-19
Utilities and Development Tools	3-20
JPublisher	3-20
JDeveloper	3-20
loadjava and dropjava	3-21
Data Warehousing and Very Large Data Bases (VLDB)	3-21
Summary Management Using Materialized Views	3-21
Transportable Tablespaces	3-23
Direct Path Load API	3-23
ROLAP Enhancement: Extended Aggregate Operations	3-24
Data Mining: SAMPLE Function	3-24
Function-Based Indexes	3-24
Descending Indexes	3-25
Enhancements for Bitmap Indexes	3-25
Parallelization of Aggregate Distinct Queries	3-25
Sort Improvements	3-26
Single-Table Hash Cluster	3-26
Data Warehouse Statistics Gathering	3-26
Automated Parallel Query Tuning	3-27
Statistics Generation for Long-Running Operations	3-27
New Constraint Functionality	3-28
Constraint Modification	3-28
DISABLE VALIDATE Constraint State	3-28
Index-Organized Tables	3-28
General Enhancements	3-29
Logical ROWIDs	3-29
Related Features	3-30
Database Features	3-30
Extensibility Framework	3-30
User-Defined Operators and Extensible Indexing	3-31
Extensible Optimizer	3-32
Statistics Collection Functions	3-32
Selectivity Functions	3-32

Cost Functions	3-32
Cartridge Services	3-32
Object Relational Enhancements	3-33
LOB Enhancements	3-33
Temporary LOBs	3-33
Other Enhancements.....	3-33
User-Defined Object Identifiers.....	3-34
Nested Table Enhancements.....	3-34
Varray Enhancements.....	3-35
Collection Unnesting	3-35
Enhancement of Object Views.....	3-35
Triggers on Nested Table View Columns.....	3-35
Object Cache Enhancements.....	3-36
ANALYZE Statement, Dangling REFs.....	3-36
Subqueries in VALUES Clause of INSERT Statement	3-36
SQL*Loader Support for Objects, Collections, and LOBs	3-36
Other Object Relational Enhancements.....	3-37
Partitioning of Tables Containing Objects.....	3-37
Parallel Query Support.....	3-37
Optimizer Plan Stability	3-37
Advanced Queuing	3-37
AQ-Based Publish/Subscribe	3-38
Database Event Publication	3-39
Enhanced Propagation.....	3-39
Miscellaneous Enhancements.....	3-39
DBMS_REPAIR Package	3-40
Redo Log Analysis Using LogMiner.....	3-40
Drop Column	3-41
Locally Managed Tablespaces	3-41
Online Index Creation, Rebuild, and Defragmentation.....	3-41
Nonpartitioned Table Reorganization.....	3-42
Online Read-Only Tablespaces.....	3-42
Temporary Tables.....	3-42
National Language Support (NLS)	3-43
General Enhancements	3-43

Changing Character Sets with ALTER DATABASE	3-43
TRIM Function	3-44
External Routines.....	3-44
The Call Specification.....	3-44
External Routine DLL Caching.....	3-45
External Routine Object Support.....	3-45
Database Security.....	3-45
Partitioning Enhancements	3-45
Range Partitioning Enhancements	3-45
Merging Partitions.....	3-45
Updatable Partition Keys	3-46
New Partitioning Methods.....	3-46
Additional Performance Gains	3-47
Enhanced Partition Elimination	3-47
Partition-Wise Join.....	3-47
Partitioned Table LOB Support	3-47
Partitioning of Index-Organized Tables.....	3-47
System Management	3-47
Database Resource Management.....	3-48
Recoverability and Availability in the Database.....	3-49
Multiple Remote Archive Destinations.....	3-50
Multiple Archive Processes.....	3-50
Automated Standby Databases	3-50
Automatic Archival.....	3-50
Read-only Databases.....	3-51
Fast-Start Fault Recovery.....	3-51
Fast-Start Checkpointing.....	3-51
Fast-Start On-Demand Rollback.....	3-51
Fast-Start Parallel Rollback	3-52
Database SUSPEND/RESUME	3-52
Control File Character Set	3-52
Recovery Manager (RMAN)	3-52
Recovery Catalog No Longer Required	3-53
Media Management API, Version 2.0.....	3-53
Proxy Copy.....	3-53

Media Pool Selection.....	3-53
Vendor Identification and Improved Error Messages.....	3-53
Send Statement	3-54
Crosscheck Catalog	3-54
List and Report Statements	3-54
Recovery Catalog Maintenance Statements	3-54
CREATE CATALOG.....	3-54
UPGRADE CATALOG.....	3-54
DROP CATALOG	3-55
Database Maintenance Statements	3-55
STARTUP/SHUTDOWN	3-55
ALTER DATABASE.....	3-55
DUPLICATE Statement	3-55
Node Affinity Detection	3-55
Backups	3-55
Duplexed Backup Sets	3-55
Naming	3-56
Is No Longer Overwritten.....	3-56
Performing TSPITR Without A Recovery Catalog	3-56
Backup Performance Views	3-56
SQL*Loader Enhancements	3-56
Export and Import Utilities	3-57
Operational Simplicity	3-57
Oracle Software Packager and Oracle Universal Installer	3-58
Oracle Software Packager	3-58
Oracle Universal Installer	3-58
Configuration Improvements.....	3-59
Oracle Database Configuration Assistant.....	3-59
Oracle Data Migration Assistant.....	3-60
Oracle Enterprise Manager (EM)	3-60
Oracle Parallel Server.....	3-61
New Architecture	3-61
Easier System Administration	3-62
Diagnostic Enhancements	3-62
Oracle Parallel Server Management (OPSM)	3-62

Parallel Server Installation and Database Configuration	3-62
Instance Affinity for Jobs	3-63
Recoverability and Availability Improvements	3-63
Networking Improvements	3-64
Distributed Systems	3-64
Advanced Replication Enhancements	3-64
Performance Improvements.....	3-64
Internal Apply Packages	3-64
Faster Snapshot Refresh	3-64
Improved Mass Deployment Support.....	3-65
Parameterized Snapshot Deployment Templates	3-65
Column Level Snapshot Subsetting.....	3-65
Offline Instantiation	3-65
Improved Security	3-66
Replication Manager	3-66
Improved Oracle Lite Integration	3-66
Heterogeneous Services	3-66
Improved Management Interface: Fixed Views.....	3-67
Miscellaneous Improvements.....	3-67
Remote Join Enhancements	3-67
Networking, Security, and Oracle Advanced Security	3-68
Securing Data in the Oracle Database Server	3-68
Fine-Grained Access Control	3-69
Application Context	3-69
Secure Network Computing Through the Oracle Advanced Security Option	3-70
SSL	3-70
RADIUS.....	3-72
Integrated Security and Directory Services	3-72
Oracle Wallet Manager.....	3-73
Certificate Authority	3-74
Oracle Internet Directory	3-74
Oracle Enterprise Security Manager.....	3-74
Networking Improvements.....	3-75
Service Naming.....	3-75
Service Registration.....	3-76

Client Load Balancing.....	3-76
Connect-Time Failover	3-76
Connection Load Balancing.....	3-77
Net8 Configuration Assistant	3-77
Native Authentication on Windows NT.....	3-77
Multi-Tier Authentication and Authorization	3-78
Program Interfaces.....	3-78
PL/SQL	3-78
Autonomous PL/SQL Blocks	3-78
Invoker Rights.....	3-79
PL/SQL Bulk Binds.....	3-79
Dynamic SQL in PL/SQL.....	3-79
Parameter Passing by Reference	3-80
PL/SQL Procedures for Supporting REF-based Operations	3-80
Monitoring and Analysis of Program Execution.....	3-80
DBMS_TRACE.....	3-80
DBMS_DEBUG	3-80
DBMS_PROFILER.....	3-80
Increased Package Body Size	3-81
Purity Rules Relaxed.....	3-81
Precompiler Enhancements.....	3-81
Pro*C/C++	3-81
Pro*COBOL	3-82
Java Interfaces	3-84
OCI Enhancements	3-84
DML Returning Enhancements.....	3-84
Enhanced Array DMLs.....	3-84
OCI Callback Registrations	3-84
OCIDescribeAny Enhancements.....	3-85
Asynchronous Event Notification.....	3-85
Non-Blocking OCI.....	3-85
Windows NT Integration	3-85
Oracle Application Generator for Microsoft Visual Studio	3-85
Oracle Objects for OLE	3-86
Microsoft Transaction Server Integration.....	3-87

COM Cartridge	3-87
---------------------	------

4 Release 8.0 New Features and Options

Partitioned Tables and Indexes	4-2
What is Partitioning?	4-2
Greater Control Over Data	4-3
Easy-To-Use Administration.....	4-3
Improved Data Warehouse Performance	4-3
Enhanced Star-Query Processing	4-3
New Parallel Operations.....	4-4
Increased Database Size	4-5
Improvements for OLTP Applications	4-5
Extended Backup/Recovery Subsystem	4-5
Very Large User Populations	4-6
Advanced Queuing	4-6
Parallel Server Improvements.....	4-7
Transparent Application Failover (TAF).....	4-8
Improved TP Monitor Support.....	4-9
Better Security Administration	4-9
Replication	4-9
Object-Relational Technology	4-10
Object Types and Views.....	4-11
Calling External Procedures From Within the Database	4-12
Client-Side Support for Objects	4-12
Evolution of Relational Environments	4-12
Development Tools for Object Modeling	4-13
Multimedia Data	4-13
Java.....	4-13
Extensibility	4-14
Migration and Interoperability	4-14
Other Enhancements	4-14
Index-Organized Tables.....	4-14
Reverse Key Indexes.....	4-14
Improved Constraint Processing	4-15
Two Character Sets in One Database	4-15

Miscellaneous Improvements.....	4-15
5 Oracle8i Feature and Option Availability	
Options	5-2
Features.....	5-4
The V\$OPTION Table	5-13
6 Oracle8i Documentation	
Online Documentation Structure	6-2
Documentation Titles.....	6-3
Oracle8i Master Index	6-3
Oracle8i Server	6-3
Application Development	6-5
Online-Only Books	6-6
Networking and Security	6-6
Parallel Server	6-6
SQL*Plus	6-7
<i>interMedia</i> , Spatial, Time Series, and Visual Image Retrieval	6-7
Oracle Enterprise Manager	6-8
WebDB	6-8
Online-Only Books	6-8
7 Deprecated and Desupported Features	
Strings of Zero Length Are Not Equivalent To a NULL.....	7-2
The SELECT Privilege	7-2
Date Format Strings Are Stricter.....	7-2
SERIALIZABLE=TRUE Is No Longer Supported.....	7-2
Non-Deferred Linking.....	7-3
Applications Using Oracle6 OCI Libraries	7-4
Applications Using Oracle7 OCI Libraries	7-4
Applications Using Oracle8 OCI Libraries	7-5
Single-Task Linking	7-5
CONNECT INTERNAL.....	7-6
Partition Views	7-6

V6 Compatibility Behavior	7-6
Use of "THE(subquery)" Expression	7-7
Server Manager Desupport	7-7
The SGADEF File.....	7-7
LONG Column Support	7-7
The Oracle Security Server and Cryptographic Toolkit.....	7-8
Dynamic Views Used for Monitoring Parallel Execution Performance.....	7-8

Glossary

Index

Send Us Your Comments

Getting to Know Oracle8i , Release 2 (8.1.6)

Part No. A76962-01

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:

- E-mail - infodev@us.oracle.com
- FAX - (650) 506-7228. Attn: Server Technologies Documentation Manager
- Postal service:
Oracle Corporation
Server Technologies Documentation Manager
500 Oracle Parkway
Redwood Shores, CA 94065
USA

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

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

Preface

This book introduces the features and functionality of Oracle8i. It describes new features, options, and enhancements; and it includes feature and option matrixes to identify what is available with each specific product configuration (Oracle8i, Oracle8i Enterprise Edition, or Oracle8i Personal Edition). It references the documentation that is available for Oracle8i and identifies deprecated or desupported features.

This preface contains the following sections:

- [Intended Audience](#)
- [How this Book is Organized](#)
- [How to Use This Book](#)
- [What's New](#)
- [Related Documents](#)
- [Your Comments Are Welcome](#)

Intended Audience

This book is addressed to anyone who would like to become familiar with Oracle8i. While it provides an overview of Oracle8i, it is intended for those who want to identify and develop an understanding of its enhancements, new features, and options.

How this Book is Organized

This book is organized as follows:

Chapter 1, "Overview"	This chapter contains an introduction to Oracle8i.
Chapter 2, "Oracle 8i Release 2 (8.1.6) New Features"	This chapter contains summaries of the new features provided in Oracle8i, release 2 (8.1.6). Some additional products that complement or enhance the functionality of Oracle8i are also discussed.
Chapter 3, "Oracle 8i (8.1.5) New Features"	This chapter contains summaries of the new features provided in Oracle8i, release 8.1.5. Some additional products that complement or enhance the functionality of Oracle8i are also discussed.
Chapter 4, "Release 8.0 New Features and Options"	This chapter summarizes Oracle8 features. This summary is important because many of the Oracle8i features enhance or extend functionality introduced in Oracle8.
Chapter 5, "Oracle8i Feature and Option Availability"	This chapter contains feature matrixes that present the availability and packaging of features and options contained in Oracle8i.
Chapter 6, "Oracle8i Documentation"	This chapter describes the contents of the Oracle8i documentation set.
Chapter 7, "Deprecated and Desupported Features"	This chapter lists the deprecated and desupported features that a user should be aware of when migrating to Oracle8i.
Glossary	This chapter contains definitions of terms, most of which are new to Oracle8i.

How to Use This Book

[Chapter 2, "Oracle 8i Release 2 \(8.1.6\) New Features"](#) describes new features and enhancements. For more information, see the documentation that is cross-referenced in the description.

Oracle8i contains a subset of the features and options that are available in Oracle8i Enterprise Edition and Oracle8i Personal Edition. You can refer to the matrixes in [Chapter 5, "Oracle8i Feature and Option Availability"](#) to determine whether a particular option or feature is available for your product configuration.

What's New

The following books are new to the release 2 (8.1.6) generic documentation set:

- *Oracle Internet Directory Administrator's Guide*
- *Oracle8i Application Developer's Guide - Object-Relational Features*
- *Oracle8i Data Warehousing Guide*

In addition to the new data warehouse features, this books contains the data warehouse and parallel execution material that was in the release 8.1.5 *Oracle8i Tuning* manual.

- *Oracle8i Supplied Java Packages Reference*

This book describes some of the Oracle-supplied Java packages. It complements the 8.1.5 *Oracle8i Supplied Packages Reference*, which has been renamed in release 2 to *Oracle8i Supplied PL/SQL Packages Reference*.

The 8.1.5 *Oracle8i Backup and Recovery Guide* has been divided into three books for release 8.1.6:

- *Oracle8i Backup and Recovery Guide*
- *Oracle8i Recovery Manager User's Guide and Reference*
- *Oracle8i Standby Database Concepts and Administration*

The 8.1.5 *Oracle8i Parallel Server Concepts and Administration* has been divided into two books for release 8.1.6:

- *Oracle8i Parallel Server Concepts*
- *Oracle8i Parallel Server Administration, Deployment, and Performance*

The 8.1.5 *Oracle8i Tuning* book has been renamed in 8.1.6 to *Oracle8i Designing and Tuning for Performance*.

Documentation for Oracle Enterprise Manager is now included as part of the generic documentation set. For book titles in this documentation set, see [Chapter 6, "Oracle8i Documentation"](#).

In release 2 (8.1.6), the *interMedia*, Spatial, Time Series, and Visual Image Retrieval product names start with Oracle, not Oracle8i. For example, in release 8.1.5, the spatial product was called Oracle8i Spatial; however, in release 8.1.6, it is called Oracle Spatial.

Related Documents

While this book identifies enhancements, new features, and options for Oracle8i, release 2 (8.1.6), it does not attempt to identify every new or changed initialization parameter, data dictionary table, view, command, package, etc. Such information should be available from the cross-referenced documentation.

Some general references for finding such information include:

<i>Oracle8i Migration</i>	For a complete list of all new Oracle8i initialization parameters, static data dictionary tables, and dynamic performance views (v\$ views). Many initialization parameters have been added, changed, or dropped in this release.
<i>Oracle8i Supplied PL/SQL Packages Reference and Oracle8i Supplied Java Packages Reference</i>	New packages are included, or cross-referenced, in these books.
<i>Oracle8i Reference</i>	This is the complete Oracle8i reference for the following: <ul style="list-style-type: none">■ Initialization Parameters■ Static Data Dictionary Views■ Dynamic Performance (v\$) Views■ Database Limits■ SQL Scripts■ Oracle Wait Events■ Enqueue and Lock Names■ Statistics Descriptions

Oracle8i SQL Reference This is the complete reference of SQL statements. It identifies all new Oracle8i statements.

Your Comments Are Welcome

We value and appreciate your comments as an Oracle user and reader of our references. As we write, revise, and evaluate, your opinions are the most important input we receive. At the front of this reference is a Reader's Comment Form that we encourage you to use to tell us both what you like and what you dislike about this (or other) Oracle manuals. If the form is missing, or if you would like to contact us, then please use the following address or fax number:

Server Technologies Information Development Manager
Oracle Corporation
500 Oracle Parkway
Redwood Shores, CA 94065
FAX: 650-506-7228

You can also e-mail your comments to the Information Development department at the following e-mail address: infodev@us.oracle.com

1

Overview

This chapter presents an overview of the Oracle8i database server. This chapter contains the following sections:

- [Introducing Oracle8i](#)
- [A Family of Database Products](#)

Introducing Oracle8i

Oracle8i, the database for Internet computing, changes the way information is managed and accessed to meet the demands of the Internet age, while providing significant new features for traditional online transaction processing (OLTP) and data warehouse applications. It provides advanced tools to manage all types of data in Web sites, and it also delivers the performance, scalability, and availability necessary to support very large database (VLDB) and mission-critical applications.

Oracle8i is much more than just a simple relational data store. A new option, Oracle8i *interMedia*, allows businesses to manage and access multi-media data, including image, text, audio, video, and spatial (locator) data. The *interMedia* Clipboard and Web Agent work together to Web-enable *interMedia*. WebDB is a new HTML-based development tool for building HTML Web pages with content based on data stored in Oracle databases.

Oracle8i introduces new support for Java by including a robust, integrated, and scalable Java Virtual Machine (VM) within the server, named Oracle JServer. This expands Oracle's support for Java into all tiers of applications, allowing Java programs to be deployed where they perform best—in the client, server, or middle tier—without recompiling or modifying the Java code.

Not only does Oracle8i introduce breakthrough technology, such as a Java VM, it also has significant new features and functionality for traditional OLTP and data warehouse applications. For data warehouses, Oracle8i now provides sophisticated summary management features to store aggregates that are commonly queried, greatly reducing query processing.

For OLTP applications, Oracle8i introduces the optimizer plan stability feature, allowing users to ensure that the optimizer selects the same tuned plan every time the same query is executed. Advanced Queuing has been enhanced to support a publish/subscribe paradigm. OLTP applications benefit from the many new and enhanced features that improve the recoverability and availability of the database during routine operations, such as index rebuilds, and in disaster situations where a standby database may be activated.

Oracle8i introduces database resource management. This gives DBAs the ability to control the processing resources allocated to a user or group of users. Two new partitioning methods, hash and composite, complement the established range partitioning method to provide a rich set of partitioning methods from which DBAs may choose the best method to fit an application's profile and workload.

Oracle Parallel Server features a new architecture in Oracle8i. Cache fusion is a new "diskless" ping architecture that greatly improves inter-instance communication.

New networking features improve the ease of use for Oracle Parallel Server, and system administration is made easier through the enhancement of Oracle Parallel Server Management, the new Oracle Universal Installer, and the Oracle Database Configuration Assistant.

Oracle8i extends the functionality of advanced replication, focusing on mass-deployment applications. Data can be replicated to servers that are closer to users and have only the data those users need, providing better performance. Security has been improved. The replication manager has been rewritten in Java and is no longer constrained to run on a Windows-based machine.

While Oracle has expanded its support of multimedia data through *interMedia*, and other complex data types through the separately packaged Visual Information Retrieval, Time Series, and Spatial options, it has also provided a new extensible architecture for independent software vendors (ISVs) and other developers to exploit when developing support for new types of data or specialized processing. There is significant enhancement to Oracle's object relational technology, and Oracle8i makes it available in both the standard edition and the Enterprise Edition.

Language improvements for PL/SQL, Pro*C/C++, Pro*COBOL, and the Oracle Call Interface (OCI) include significant new functionality. A new book provides information on the enhanced features with National Language Support.

Fine-grained access control and application context features build row-level application security into the database, rather than leaving it at the application level. Stronger security is provided throughout the enterprise by expanding or enhancing methods of authentication and authorization, centralizing user management, and supporting standards.

There is greater Windows NT integration in Oracle8i. It provides full, native integration with Microsoft Transaction Server (MTS) in the Windows NT environment. Application development is made simpler by the Oracle Application Wizard (AppWizard) for Visual Studio, which provides developers with a GUI tool for creating a Visual C++, Visual InterDev, or Visual Basic application accessing data in an Oracle database. The generated code framework includes Oracle Objects for OLE (OO4O) classes for connectivity and data access to Oracle databases.

Finally, Oracle8i includes Oracle Enterprise Manager, which is the comprehensive management framework for managing the Oracle database and application environment. Oracle Enterprise Manager presents an easy-to-use centralized console, a rich set of management tools, and the extensibility to detect and solve any problems that may arise. It also includes several administrative applications for performing day-to-day tasks for databases and applications, such as scheduling backup routines.

Oracle8i is designed to access and manage all your data using the style and infrastructure of the Internet. Oracle8i is the most complete and comprehensive platform for building, deploying, and managing Internet and traditional applications.

- Oracle8i simplifies the development of applications.
- Oracle8i simplifies the management of Internet content.
- Oracle8i simplifies the deployment of applications.

Oracle8i provides the lowest cost platform for developing and deploying applications on the Internet.

A Family of Database Products

Oracle8i is available in four standalone versions:

- Oracle8i, also referred to as the standard edition delivers unprecedented ease-of-use, power, and database functionality for workgroup or department-level applications. It includes core features for most Windows NT and UNIX applications.
- Oracle8i Enterprise Edition adds several high-end features and options for mission-critical OLTP and data warehousing applications.
- Oracle8i Personal Edition is a single-user development database for Windows NT, Windows95, or Windows98, that now includes all Enterprise Edition options and features.
- Oracle8i Lite is a lightweight mobile Java database that easily synchronizes laptops and hand held devices with corporate databases. It is not discussed in this book.

The Oracle8i database is available as Oracle8i, Oracle8i Enterprise Edition, and Oracle8i Personal Edition. Not all features and options are available with all editions of the Oracle8i database.

See Also: For feature and option availability, see [Chapter 5](#).

Oracle 8i Release 2 (8.1.6) New Features

Oracle8i release 2 (8.1.6) is the first maintenance release of the Oracle8i database. It includes many bug fixes, performance improvements, and feature enhancements. Moreover, this release also includes several noteworthy new features for designing data warehouses and building Internet applications.

The enhancements and new features in release 8.1.6 can be categorized as:

- [Oracle8i Improvements](#)
- [Enhanced Support of Internet Computing](#)
- [Oracle Advanced Security Improvements](#)
- [Advanced New Functionality](#)
- [Release 2 \(8.1.6\) New Features](#)

Oracle8i Improvements

The main focus of release 2 (8.1.6) is to maintain stability and improve on quality, performance, and installation/configuration and management of the Oracle8i database server.

Performance Improvements

Various components of the Oracle8i server have been optimized to perform better in release 8.1.6.

Oracle Parallel Server has been optimized for a 2-node high availability configuration. In this special high availability mode of Oracle Parallel Server, one node acts as the primary database instance, while the other node can be used for some DBA tasks, report generation, and other non-buffered write database activities. If the primary node fails, then the secondary node becomes the primary instance. In release 8.1.6, this configuration is optimized to provide better throughput by efficiently managing the locks on only the primary node, resulting in less overhead and better response times. This means that now you can migrate from a single instance Oracle database to a highly available Oracle Parallel Server database with the same guaranteed performance.

For details on how to configure Oracle Parallel Server for primary and secondary instances, see *Oracle8i Parallel Server Setup and Configuration Guide*.

The performance of SQL*Loader direct path load has been dramatically improved for loading simple files encoded using single-byte character sets (e.g., ASCII, Latin1). When the input data fields are either terminated by a single byte delimiter (e.g., comma separated values) or positional (or are a combination of the two), the load performance can be up to two times as fast.

For more information, see *Oracle8i Data Warehousing Guide*.

PL/SQL performance has been optimized in several areas, including enhanced performance of C callouts and the manipulation and use of varrays and nested tables.

For more information, see *PL/SQL User's Guide and Reference*.

Applications, including many packaged applications, that use literals instead of bind variables in SQL can see improved performance because release 8.1.6 can automatically substitute bind variables in place of the literal values. This lets multiple users share the parsed SQL regardless of the literal value, dramatically reducing the SQL parse time and saving memory in these applications.

For more information, see *Oracle8i SQL Reference* and *Oracle8i Designing and Tuning for Performance*.

Improved Installation and Management

Installation of Oracle products is now much easier with the consolidation in the Oracle8i installation process of Oracle Enterprise Manager, the DBA Management Pack, and Oracle Internet Directory.

Oracle Enterprise Manager version 2.1 is available with Oracle 8i release 2 (8.1.6). In this release, Oracle Enterprise Manager can manage twelve additional products, including replication, Oracle Parallel Server, *interMedia* Text, Oracle Applications, and Forms Server. Distributed Access Manager is a new component of Oracle Enterprise Manager that manages and administers heterogeneous distributed environments. Also, Oracle Internet Directory now includes Oracle Directory Manager, a configuration and management tool integrated in the Oracle Enterprise Manager framework.

Functionality has been enhanced to support improved target discovery, notification filtering, and administrative notification blackouts. Notification blackouts allow administrators to suspend email and paging notifications for events or jobs that are registered against targets that are being brought down for backups or any other maintenance.

The DBA Management Pack, a set of management tools bundled with Oracle Enterprise Manager, now supports many additional Oracle8i features, including dropping table columns and controlling database resources (such as CPU) by managing resource plans and consumer groups. New wizards are provided for creating dimensions and materialized views which greatly assist data warehousing projects.

DBA Studio combines the functionality of the Schema, Security, Storage, and Instance managers into one tool for improved workflow between the various day-to-day DBA tasks. A new Create View Wizard is now available to easily create views. A new Analyze Wizard is included with Enterprise Manager for gathering statistics on tables, indexes, and clusters. In addition, the Create Table Wizard now provides the capability for creating partitioned tables.

Finally, version 2.1 has numerous improvements to the graphical user interface, such as displaying descriptions for instance parameters, more information about control files and archived logs, and bar charts that readily display tablespace and datafile used space.

For more information, see *Oracle Enterprise Manager Concepts Guide*.

Oracle Universal Installer and the Net8 Configuration Assistant have been enhanced to support Oracle Parallel Server, simplifying the installation and network setup of Oracle8i running in a clustered environment. Also, the Database Configuration Assistant (DBCA), which is used to create or delete databases, now has improved multi-threaded server (MTS) support, supports the use of raw disk devices, and runs scripts during database creation that enable most Oracle database options, such as the Oracle Internet Directory, Oracle Jserver, and XML Parser for Java. The Oracle Data Migration Assistant continues to provide support for migrating Oracle7 databases to Oracle8i and upgrading Oracle8 databases to Oracle8i.

For more information, see *Net8 Administrator's Guide* and *Oracle8i Migration*.

Improved Quality and Support

Release 8.1.6 is the first maintenance release of Oracle8i and contains many bug fixes, but it also has many enhancements to improve the reliability and availability of database operations.

Unplanned downtime due to hardware corruption is reduced in release 8.1.6 with improved block checking functionality to protect the critical `SYSTEM` tablespace from corruption without impacting performance on other tablespaces.

The `SYSTEM` tablespace is extremely critical, because hardware corruption in the `SYSTEM` tablespace causes the database to shut down. In release 8.1.6, block checking and checksums is always be turned on for the `SYSTEM` tablespace, allowing the Oracle server to detect and repair many `SYSTEM` tablespace corruptions before they are written to disk or read by other transactions. Error handling is improved, because Oracle re-reads control file blocks from a different file if all attempts resulted in checksum failure. Release 8.1.6 also re-reads redo blocks from non-multiplexed redo logs if previous reads resulted in checksum failure. These sophisticated corruption detection and repair mechanisms reduce downtime due to possible hardware-induced corruption.

For more information, see *Oracle8i Administrator's Guide*.

Tracing PL/SQL code execution has been enhanced by storing trace data in tables, allowing easy access through standard SQL. Also, more data can be captured, such as SQL text and bind variables.

For more information, see *PL/SQL User's Guide and Reference*.

Enhanced Support of Internet Computing

Oracle8i includes significant enhancements in release 2 (8.1.6) to keep pace with the technological requirements of demanding Internet applications. Oracle JServer, the integrated Java Virtual Machine (VM) introduced with the initial release of Oracle8i, has improved performance, support for "Java2" and XML, and various other enhancements. Security and support of the LDAP standard have been enhanced in release 8.1.6, allowing applications to efficiently implement and manage robust security policies.

Enhanced Java Support

Oracle8i release 8.1.6 includes performance and functionality improvements to Oracle JServer for building and running Java-based applications. These enhancements include the following:

- Java2 support (JDK 1.2)
- JDBC 2.0 core support plus connection pooling and distributed transactions
- Multi-byte charset support
- Improved performance through lower per-session and per-call overhead and through optimized server JDBC and SQLJ execution.
- Remote Java debug support, accessible from many Java debuggers including Oracle JDeveloper
- Performance gains and reduced footprint due to Java support for shared, read-only object memories
- Java Message Service (JMS) support, allowing easy integration with Oracle's Advanced Queuing functionality via the standard JMS API
- Object type support for Advanced Queuing through the Java API

For more information, see *Oracle8i Application Developer's Guide - Advanced Queuing* and the Java documentation set.

XML Support

XML has emerged as the de facto standard for describing business data for business to business commerce on the Internet. Just as HyperText Markup Language (HTML) is the universal language of Web pages on the Internet, businesses have sought a similar language for describing business data. The adoption of eXtensible Markup Language (XML) by the World-Wide Web Consortium (W3C) was the first step

towards identifying a common format for businesses to exchange information over the Internet. Similar to HTML, XML provides the facility to embed business descriptors along with raw business data in a message. Oracle provides several XML parsers that convert XML data into a form that can be easily understood by the Oracle8i database.

Oracle's XML products include the XML Parser for Java - version 1 and version 2 (which includes an XSLT processor), the XML Class Generator, and the XML Parsers for C, C++, and PL/SQL.

Because the Oracle XML Parser for Java is implemented in Java, it runs efficiently in the Oracle8i Java VM, Oracle JServer. Release 8.1.6 includes the Oracle XML Parser for Java, where the Java XML parser classes (DOM/SAX APIs) are pre-loaded into Oracle JServer.

For more information, see *Oracle8i Supplied Java Packages Reference*.

PL/SQL Pages

Oracle 8i release 8.1.6 introduces PL/SQL Server Pages (PSP), which are server-side Web pages (written in HTML or XML) with embedded PL/SQL scripts designated with special tags. This functionality enables fast development and deployment of Web pages with dynamic content, where the 'code' can reside on the server while only the invocation is present in the HTML page.

For more information, see *Oracle8i Concepts* and *Oracle8i Application Developer's Guide - Fundamentals*.

interMedia, Spatial, Time Series, and Visual Image Retrieval

interMedia

Oracle *interMedia* enables Oracle8i to manage text, documents, images, audio, video, and geographic location information in an integrated fashion with structured enterprise information.

Oracle8i *interMedia* Text provides content-based retrieval on free text with both literal (word) predicates and thematic predicates. Several performance improvements have been made to *interMedia* Text release 8.1.6:

- The ABOUT algorithm has been greatly enhanced for better precision and recall.
- With the MULTI lexer, multi-lingual databases can now store documents of different languages.

- Enhanced XML support includes indexing and searching attribute text, using "nest within" for sophisticated queries, and support for doctype-limited tag detection.

Other components of *interMedia* have been enhanced in release 8.1.6, including adding native support in *interMedia* Audio for Real Networks RealAudio and support for all layers of audio formats MPEG1 and MPEG2, including layer 3 (more commonly known as MP3). Similarly, *interMedia* Video now includes native support for Apple Quicktime 3.0, Microsoft AVI, and RealNetworks RealVideo. Also, the `Set_Properties` method has been enhanced across all *interMedia* image, audio, and video formats. When present in these formats, *interMedia* can extract application metadata and store them as an XML string within the media object. Finally, *interMedia* Web agent support has been extended to include the Apache Web server.

For more information, see *Oracle8i interMedia Text Reference*, *Oracle8i interMedia Text Migration*, and *Oracle8i interMedia Audio, Image, and Video User's Guide and Reference*. Also, see the online `README.txt` file in your `ORACLE_HOME` directory. Depending on your operating system, this file is in `ORACLE_HOME/ord/im/README.txt` (for information on *interMedia* Audio, Image, and Video and *interMedia* Locator), or `ORACLE_HOME/ord/vir/README.txt` (for information on Visual Information Retrieval).

Oracle Spatial

Oracle Spatial is an integrated set of functions and procedures that enables spatial data to be stored, accessed, and analyzed quickly and efficiently in an Oracle database. Oracle Spatial includes the following enhancements for release 8.1.6:

- Improved metadata management through system-level views and metadata tables.
- New types (geometry and element) allow for better representation of the geometries in the database. These types also enable better management of spatial geometries in a client-side application.
- More features in the Validate Geometry function. These new features can help validate complex geometries and provide better feedback on invalid geometries.
- Improved performance for all spatial functions, especially for Within Distance and Buffer operations. Also, performance of all index creation operations is improved.

- Enhancement of Boolean operations to operate at the geometry level rather than at the polygon level.

For more information, see *Oracle Spatial User's Guide and Reference*.

Oracle Advanced Security Improvements

Oracle Advanced Security has improved configuration and management tools to simplify security management. Oracle Advanced Security also provides new forms of network encryption, to ensure the security of all protocols accessing the Oracle8i database, and enhanced single sign-on.

Network Security Enhancements

Release 2 (8.1.6) enhances Oracle's support for the SSL (Secure Sockets Layer) standard. SSL encryption for Internet Intra-ORB Protocol (IIOP) communications is now available, enabling secure Enterprise Java Beans (EJBs). Also, JavaASO, a Java version of the Oracle Advanced Security encryption libraries, is now available to secure "thin" JDBC connections. JavaASO provides DES encryption, with anonymous Diffie-Hellman key exchange, in 100% Java. Oracle Advanced Security thus secures *all* protocols into the Oracle8i database, whether IIOP, "thick" or "thin" JDBC, or Net8.

Oracle Advanced Security has also completed the operational testing phase of FIPS-140 (Federal Information Processing Standard) level 2 certification, a United States government standard that relates to the security of cryptographic products. Completion of the FIPS-140 certification is required by many organizations, among them the United States federal government and many financial markets.

Single Sign-On

Oracle Advanced Security already supports many forms of single sign-on for database users, among them Kerberos, SESAME, and DCE. Release 8.1.6 adds support for SSL-based single sign-on.

PKI Credential Management

Oracle Wallet Manager provides secure management of PKI (public key infrastructure)-based user credentials. Oracle Wallet Manager creates a private and public key pair for a user and issues a PKCS#10 certificate signing request which can be fulfilled by a Certificate Authority (CA). After the CA issues an X.509 certificate, the user can load the certificate into his wallet. Oracle Wallet Manager

also manages user trustpoints, the list of root certificates that the user trusts, and is pre-configured with root certificates from PKI vendors such as VeriSign and Cybertrust. Wallets are protected using password-based strong encryption.

In most cases, a user need never access a wallet once it has been configured, but can easily access his wallet using Oracle Enterprise Login Assistant, a very simple-to-use login tool that hides the complexity of a private key and certificate from users. After a user has securely opened his wallet, he can then connect to multiple databases over SSL, without providing additional passwords. This provides the benefit of strong authentication as well as single sign-on.

SSL for single sign-on can be used alone or in conjunction with enterprise user management, described below.

For more information, see *Net8 Administrator's Guide*, *Oracle Advanced Security Administrator's Guide*, and *Oracle8i Distributed Database Systems*.

Enterprise User Management

Enterprises today face tremendous challenges in managing information about users, keeping user information current, and securing access to all the information in an enterprise. Oracle8i release 8.1.6 introduces enterprise user management. Enterprise users and their authorizations are managed in Oracle Internet Directory, an LDAP-based directory service, using Oracle Enterprise Security Manager, a tool accessible through Oracle Enterprise Manager.

Enterprise users can be assigned *enterprise roles* (which are containers of database-specific *global roles*), that determine their access privileges in databases. For example, the enterprise role CLERK could contain the global role HRCLERK on the human resources database and the global role ANALYST on the payroll database. An enterprise role can be granted or revoked to one or more enterprise users. For example, an administrator could grant the enterprise role CLERK to a number of enterprise users who hold the same job. This information about users and roles is protected in the directory through Access Control Lists, ensuring that only a privileged administrator can manage users, and grant and revoke roles.

User/Schema Separation

In general, users do not need their own accounts – or their own schemas – in a database, they merely need to access an *application* schema. For example, users John, Firuzeh, and Jane are all users of the payroll application, and they need access to the payroll schema on the finance database. None of them needs to create his or her own objects in the database; in fact, they need only access payroll objects.

Release 8.1.6 lets you separate users from schemas, so that many enterprise users can access a single, shared application schema. Instead of creating a user account (that is, a user schema) in each database a user needs to access, you need only create an enterprise user in the directory, and “point” the user at a shared schema that many other enterprise users can also access. For example, if John, Firuzeh, and Jane all access the sales database, you need only create a single schema, e.g. ‘sales_application’ which all three users can access, instead of creating an account for each user on the sales database.

Now, you can truly create an enterprise user once, in the directory, who nonetheless can access multiple databases using only the privileges she needs to perform her job, thus lowering the cost of managing users in an enterprise. Another benefit of schema-independent users is that you can manage many more users than could otherwise be done with users tied to individual database accounts.

Schema-independent users thus enables scalability of user management for the Internet.

Oracle’s LDAP version 3-compliant directory server, Oracle Internet Directory, is fully integrated with Oracle8i and supports “off-the-shelf” enterprise user management. Other LDAP directories, including Novell Directory Service (NDS) and Microsoft’s Active Directory for Windows 2000 will be certified to operate with Release 8.1.6.

For more information, see *Oracle Internet Directory Administrator’s Guide* and *Oracle Advanced Security Administrator’s Guide*.

Advanced New Functionality

Oracle8i release 2 (8.1.6) introduces many new features to improve the functionality and performance of the Oracle database, including several important improvements for data warehouse and object relational applications.

Data Warehouse Enhancements

Analytic Functions

Oracle8i release 8.1.6 introduces a powerful new family of SQL functions for business intelligence and data-warehousing applications. These functions are collectively called ‘analytic functions’, and they provide significantly improved performance for many business analysis queries.

Some examples of the new capabilities are:

- Ranking ("Find the top 10 sales representatives in each region.")
- Moving aggregates ("What is the 200-day moving average of our company's stock price?")
- Period-over-period comparisons ("What is the percentage growth of January 1999 over January 1998?")
- Ratio-to-report ("What is January's sales as a percentage of the entire year's?")
- Other functions include cumulative aggregates and lag/lead expressions

These functions significantly extend the capabilities of Oracle8i for analytic applications. Moreover, these new SQL functions are now being reviewed by ANSI for addition to the SQL standard in 2000, which encourages all third-party analytic tools and applications to leverage the power, scalability, and performance of these new functions.

For more information, see *Oracle8i Data Warehousing Guide* and *Oracle8i SQL Reference*.

Materialized View Enhancements

Materialized views, first introduced in the initial release of Oracle8i, have been an indispensable feature for high-performance data warehousing.

Materialized views include the functionality to create and maintain summary tables, resulting in dramatic performance improvements in aggregation queries. New functionality in Oracle Enterprise Manager supports the creation and management of materialized views and related dimensions and hierarchies through a graphical user interface, greatly simplifying the management of materialized views.

For further improved performance, an `ORDER BY` clause can be used when creating a materialized view and when inserting records using `INSERT-SELECT`. This allows rows in the table or the materialized view to be inserted in the specified desired order, resulting in better performance of query sorts when the sort order is the same as the inserted order.

For more information, see *Oracle8i Data Warehousing Guide* and *Oracle8i Replication*.

Enhanced Partitioning

Many data warehouses and data marts use partitioning to implement a "rolling window" scheme. In a rolling window, new data is added and old data is purged from a data warehouse on a regular periodic basis. For example, a data warehouse

may store the most recent 13 months of historical data. Every month, a new month's data (corresponding to a single range partition) is added to the warehouse, and the 14-month-old data (also corresponding to a single partition) is dropped from the warehouse.

The refresh capabilities of materialized views has been enhanced to allow partition maintenance operations to be executed on base tables without requiring a full refresh of associated materialized views. For example, a materialized view is not marked as 'stale' if partitions are dropped, and a materialized view can be fast-refreshed as a new partition is added.

The administration functionality of composite partitioning, the combination of range partitioning and hash partitioning, has also been improved. It is now much easier to support a rolling-window data application. An entire hash-partitioned table, with all of its partitions, can be exchanged with a composite-partitioned table's range partition and all of its hash subpartitions. Therefore, when a new rolling-window range needs to be added or dropped, multiple subpartitions can be added efficiently.

For more information, see *Oracle8i Data Warehousing Guide*.

Query Rewrite

In release 8.1.6, query rewrite has enhanced capabilities for Oracle DATE without the use of dimensions.

For more information, see *Oracle8i Data Warehousing Guide*.

Object Relational Enhancements

The extensibility framework that was introduced in the initial release of Oracle8i is enhanced in release 8.1.6 by adding Java wrappers to the ODCI (Oracle Data Cartridge Interface) API so that data cartridges can be seamlessly developed in Java. Other improvements to the Oracle8i object relational functionality include:

- Domain indexes, indexes created and maintained using the extensibility framework, are now supported in the Import/Export utilities.
- The SQL RETURNING clause now supports object types, collections, and references.
- Object views can now be updated if the object view is a multi-level object view or if it contains object constructors or a make_ref() construct.

For more information, see *Oracle8i Application Developer's Guide - Object-Relational Features*.

CACHE READS LOBs are supported in release 8.1.6. The CACHE READS option for LOBs specifies that LOB values should be brought into the buffer cache only during read and not write operations.

It is now possible to bind LONG values to LOBs, allowing smoother interoperability between LONG and LOB datatypes.

For more information, see *Oracle8i Application Developer's Guide - Large Objects (LOBs)* and *Oracle8i SQL Reference*.

Other Enhancements

Oracle8i release 8.1.6 includes many improvements to many areas of the Oracle server:

- Recovery Manager now automatically discovers which nodes of an Oracle Parallel Server cluster contain the backups needed for a restore operation.
For more information, see *Oracle8i Recovery Manager User's Guide and Reference*.
- Replication includes the Java RepAPI functionality, which provides improved integration with Oracle Lite, including support for IIOP connectivity and template functionality.
For more information, see *Oracle8i Replication*.
- Oracle now supports native connectivity to heterogeneous systems through the ODBC and OLE DB protocols. The new generic connectivity for OLE DB is integrated into the server and allows scaled down access to any system that supports ODBC and OLE DB.
For more information, see *Oracle8i Distributed Database Systems*.
- Reporting of the archiving of redo logs to a standby database site has been improved by providing a greater level of detail on the progression of the operation.
For more information, see *Oracle8i Standby Database Concepts and Administration*.
- To reduce contention for the free lists of a table, you can now use the ALTER FREELISTS statement to add free lists
For more information, see *Oracle8i Designing and Tuning for Performance*.
- Pro*COBOL now supports multi-threaded applications.
For more information, see *Pro*COBOL Precompiler Programmer's Guide*.

- CASE expressions are now supported. Also, domain indexes (indexes created and maintained using the extensibility framework) are now supported in the Import/Export utilities.

For specific information about the extensions to SQL that support this, see *Oracle8i SQL Reference*.

- WebDB now provides extensive National Language Support (NLS).
- National Language Support (NLS) now supports two new languages (Hindi and Tamil), a new territory (India), and a new character set (UTFE for EBCDIC platforms based on UTF-EBCDIC).
- Oracle Internet Directory now has full National Language Support (NLS).
For more information, see *Oracle8i National Language Support Guide*.

Release 2 (8.1.6) New Features

Oracle8i Improvements

- Improved corruption prevention and error handling
- PL/SQL tracing enhancements
- TCP/IP and SMTP support in PL/SQL
- SQL*Loader performance enhancements
- Improved performance of PL/SQL collections
- PL/SQL C callout performance improvements
- Oracle Parallel Server two-node failover performance improvement
- Multi-threaded server performance improvements
- Auto-conversion of literals to bind variables for improved performance
- Native OLE/DB driver
- Oracle Enterprise Manager enhancements
- Distributed Access Manager
- Oracle Enterprise Manager DBA Studio
- Oracle Universal Installer multiple CD support
- Improved Oracle Parallel Server support in Oracle Universal Installer

- Enhanced Net8 Configuration Assistant to support Oracle Parallel Server
- Raw device support in Database Configuration Assistant

Java Enhancements

- Lower per-session/call overhead
- Faster server-side JDBC and SQLJ
- Multi-byte support in Java
- Java2 support (JDK 1.2)
- JDBC 2.0 support
- Remote Java debug support through Sun's debug protocol
- Java object memory improvements
- Java Message Service (JMS) support
- Object support for Advanced Queuing through Java API
- XML Parser for Java

Enhanced Support of Internet Computing

- PL/SQL Server Pages
- *interMedia* enhancements
- Oracle Spatial Enhancements

Security and Networking

- Single sign-on over SSL
- JavaASO for "thin" JDBC
- LDAP support for global and enterprise users, roles, and service identifiers
- Net8 directory naming
- Oracle Names to LDAP migration
- Oracle Enterprise Login Assistant, Oracle Wallet Manager, and Oracle Enterprise Security Manager
- Schema-independent users

- SSL for IIOP
- New application context primitives for access control
- Integration with Active Directory
- Current user database links

Advanced New Functionality

- Data warehouse enhancements
- Analytic SQL functions (rank, moving average, percentile, etc.)
- Summary Management Manager
- Extended usage of `ORDER BY` when creating tables and materialized views
- Materialized view refresh enhancement for efficient partition additions
- "Rolling window" support for composite partitioning
- Query rewrite has enhanced capabilities for Oracle `DATE`

Object Relational

- Java support in ODCI (Oracle Data Cartridge Interface)
- Import/Export of domain indexes
- Object support in `SQL RETURNING`
- Object view enhancements
- `CACHE / NOCACHE / CACHE READS LOB` storage options in `CREATE` or `ALTER TABLE`
- Enhanced `LONG` and `LOB` interoperability

Other Enhancements

- Pro*COBOL support for multi-threaded applications
- Automatic location of backups in an Oracle Parallel Server configuration
- Java RepAPI support in replication
- Improved progress reporting while archiving redo logs to a standby database
- Import/Export utilities enhanced to support domain indexes

- To reduce contention for the free lists of a table, you can now use the `ALTER FREELISTS` statement to add free lists
- `SQL CASE` expression
- WebDB provides extensive national language support

Oracle 8i (8.1.5) New Features

This chapter describes new features, options, and enhancements available in Oracle8i release 8.1.5.

This chapter contains the following sections:

- [Content Management for the Internet](#)
- [Oracle interMedia, Spatial, Time Series, and Visual Image Retrieval](#)
- [Java](#)
- [Data Warehousing and Very Large Data Bases \(VLDB\)](#)
- [Database Features](#)
- [Partitioning Enhancements](#)
- [System Management](#)
- [Oracle Parallel Server](#)
- [Distributed Systems](#)
- [Networking, Security, and Oracle Advanced Security](#)
- [Program Interfaces](#)

Content Management for the Internet

The Web and the Internet bring many new challenges to organizations. Today's Web sites are often a series of Web pages that take too much time and too many human resources to develop, deploy, and manage. Additionally, their infrastructure may not be stable or robust enough to support mission-critical applications. Everything necessary to transform data into powerful Internet applications and content-driven Web sites is included with Oracle8i. It provides a Web development environment that allows software developers to easily build dynamic, data-driven Web sites with a standard Web browser and the Oracle8i database.

WebDB

WebDB is an HTML-based development tool for building HTML Web pages with content based on data stored in the Oracle database. Using a browser-based interface, Web pages can be created containing reports, charts, calendars, menus, and forms. Oracle provides extensive NLS support for components and Web sites created in WebDB.

WebDB provides build wizards for step-by-step guidance through the process of creating these components. The wizards include steps for building the PL/SQL statement that selects data used in the components, and steps for setting look and feel attributes of the component, such as color and size.

There is no need to know PL/SQL to create a component. Selecting options in the build wizard guides WebDB to automatically write the PL/SQL code for building the component.

WebDB also provides a wizard to create a Web site within the Oracle database and assign a site administrator to control its development of the site. The Web site administrator then assigns owners who are responsible for providing content to the site. Content can include Web pages containing WebDB components, links to other URLs, and files uploaded to the Web site.

Once a site is created, it is maintained using options within the site itself. Content owners and end users never need to see the WebDB user interface. For example, the look and feel of the site can be updated using options provided within the site. End users perform searches of the site's content using options within the site.

For more information, see *Oracle WebDB Release 2.1: Getting Started - Installation and Tutorial*.

Oracle Internet File System (iFS)

Oracle iFS combines the power of Oracle8i with the ease of use of a file system. Completely integrated with Oracle8i, the Internet File System is a Java application that runs within the Oracle8i Java Virtual Machine. It enables the database to become an Internet development and deployment platform.

Oracle iFS makes it possible to treat the database as if it were simply a shared network drive. Users can store and retrieve files managed by the database as if they were files managed by a file server. Because iFS supports a variety of standard protocols, users have universal access to their data. Whether a user accesses the contents of iFS through Windows Explorer, a Web browser, an FTP client, or an e-mail client, the files appear the same.

Not only can relational data appear as files, but hybrid documents that combine relational and non-relational data into single objects can be created and viewed. Files and folders in iFs are stored in a repository. They are indexed on content and file properties, allowing for intelligent text searches and queries.

From a developer's standpoint, iFS is the single data store containing the data for many different applications. For the system administrator, iFS provides a single system for file storage and messaging, rather than several separate systems to maintain and administer. In short, iFS is a simpler way of storing different types of files in the database.

Note: For Oracle8i, release 8.1.5, an SDK version of iFS is available to selected sites, including online documentation. Oracle iFS is planned for general availability with a later release.

interMedia and the Web

Web applications require advanced data management services that support the rich datatypes used in Web repository, e-commerce, and other Internet applications. Oracle8i *interMedia* adds support that enables the database server to manage multimedia content, both for Internet and traditional applications that require access to image, audio, video, text, and location information. Oracle8i *interMedia* includes the following:

- Oracle's ConText text services, to deliver the powerful text retrieval capabilities fundamental to Web applications. These services let users query and analyze documents stored in document archives, online news feeds, customer call reports, and other online text information sources. Users can query for data in

the most common formats, including HTML, Word, Excel, and Acrobat/PDF formats.

- Audio, video, and image services to support integrated management of audio, video, and image information within an Oracle8i database. These services enable access to audio, video, and image data in dozens of Internet formats from a variety of sources, both within Oracle8i and from external locations such as Web URL sites or specialized servers.
- Geometric locator services to support the development of Internet applications that help users locate information, such as stores, distribution points, and events, based on their location or distance from a given address. These *interMedia* locator services enable Oracle8i to perform location queries, and support the Internet's leading online geocoding services, including Centrus from QMSoft and MapXtreme from MapInfo.

Integral to *interMedia* is the *interMedia* Clipboard and Web Agent. It provides a clipboard and Web agent working together to Web-enable *interMedia*. The clipboard can be used to:

- Capture multimedia objects from files or external sources, such as cameras, audio input devices, and scanners.
- Retrieve multimedia objects from an Oracle8i database.
- Edit objects using other tools such as Paint Shop Pro or Word.
- "Visualize" and store objects into the database.
- Drag and drop to multimedia objects from the database to Web applications and Web authoring tools such as Oracle WebDB, Symantec Visual Page, and Microsoft FrontPage.
- Construct the simple or complex URLs for accessing objects through the *interMedia* Web agent.

The Oracle8i *interMedia* Web agent decodes URLs to retrieve multimedia objects for display by Web clients, such as Web browsers. The Oracle8i *interMedia* Web Agent supports Netscape Enterprise Server, Netscape FastTrack Server, Microsoft Internet Information Server, and Oracle Application Server.

Note: For release 8.1.5, the Clipboard and Web Agent components of Oracle8i *interMedia* are not available on the Oracle8i distribution media. Instead, you can download these components from the "Free Software" download area of the Oracle Corporation Web site:

http://www.oracle.com/products/free_software/index.html

The documentation, which includes readme files and the book *Using Oracle8i interMedia with the Web*, is included in the download.

For a more detailed summary of *interMedia*, and references to specific documentation, see below.

Oracle *interMedia*, Spatial, Time Series, and Visual Image Retrieval

Oracle8i allows you to integrate all your data into its datastore. The Oracle8i extensibility framework and object relational technology have been exploited to extend the reliability, availability, and data management capabilities of the database server to multimedia data. A new multimedia option, Oracle *interMedia*, services text, document, image, audio, video, and locational data in a single integrated package. The Spatial, Time Series, and Visual Image Retrieval options, available separately, also have been designed to utilize the Oracle8i extensibility architecture.

For each of the media types, customers and partners can easily create "plug-ins" to support additional formats, new digital compression and decompression schemes (codecs), specialized indexes, custom query optimization and methods, external media data sources and even specialized data processing algorithms.

The following is a summary of functionality available with each option.

Oracle *interMedia*

Oracle *interMedia* allows multimedia data to be managed in an integral fashion with other enterprise data. Applications can access *interMedia* through both object and relational interfaces. Database applications written in Java, C++, or traditional 3GLs can interface to *interMedia* through class libraries, PL/SQL, and the Oracle Call Interface (OCI).

For a more general description of *interMedia*, refer to "[interMedia and the Web](#)" on page 3-3. This section presents specific features of the *interMedia* components and points you to supporting documentation.

Text

Oracle *interMedia* Text delivers powerful text management and retrieval capabilities. It indexes any documents or textual content to deliver fast, accurate retrieval of information from document archives, online product catalogs, news services, media asset management systems, job postings, customer call reports, and other text information sources in multiple languages. Several techniques are used to ensure fast and accurate text searching, from traditional full-text search to document theme analysis.

The base for *interMedia* Text is Oracle's ConText Cartridge. Several new features have been added, in addition to its integration into *interMedia*. Discussions of these new features follow.

For more information, see *Oracle8i interMedia Text Migration* and *Oracle8i interMedia Text Reference*.

Improved Ease of Use This is made possible through the tight integration of *interMedia* Text with SQL, PL/SQL, Oracle Enterprise Manager, SQL*Loader, and other components of the Oracle8i database engine. Some examples:

- Most queries can be performed using SQL, and text searches can be combined with regular database searches in a single SQL statement. Query performance is improved.
- SQL*Loader can be used to load data into LOBs.
- The administrative tool has been largely redesigned and is now a Java Applet integrated with the Oracle Enterprise Manager (EM).
- New out-of-box default settings make indexing easy. By default, *interMedia* Text autodetects language settings, text column datatypes, and document formats, among other things.

Extensible Index Creating an index is now much simpler because it is created and maintained using standard SQL. This is made possible through the Oracle8i extensible indexing framework as described in "[User-Defined Operators and Extensible Indexing](#)" on page 3-31.

Extensible Query Optimizer Through new functionality provided by the Oracle8i extensible optimizer, described in "[Extensible Optimizer](#)" on page 3-32, the optimizer can be used to select the best plan for executing a CONTAINS query. It does so by analyzing collected statistics of all the tables and indexes affected by a CONTAINS query.

Extensible Knowledge Base (KB) Application developers and users can import a thesaurus to augment the built-in KB. User-specific concepts, categories, words, or phrases can be added to improve theme capabilities. Information in the thesaurus takes precedence over information in the built-in KB when there is a conflict.

Filters By default during indexing, *interMedia Text* uses the Inso Corporation filtering technology, which can filter most document formats. This filtering technology automatically recognizes document formats. Thus, this filter can be used to index single or mixed column formats. A custom or other filter may also be specified to filter documents during indexing.

For document presentation, the Inso filtering technology is used to create plain text or HTML versions of formatted documents.

Hierarchical Query Feedback Given a query expression, *interMedia Text* returns related query term information (broader term, narrower term, related term) that can help refine queries.

Theme Highlighting The behavior of theme highlighting procedures has been changed. For theme queries, *interMedia Text* procedures highlight and mark up words or phrases that best represent the theme query. This behavior is different from ConText, where paragraphs are highlighted for theme queries.

Alternate Spelling Certain languages use alternate spellings or representations for characters that cannot be stored as single characters: for instance, the letter "ö" in German is alternately represented as "oe". *interMedia Text* now checks for these alternate spellings during indexing and queries.

Stoplists ConText allows stoplists to specify words that are not to be indexed (*stopwords*). The following additional stoplists are introduced in *interMedia Text*.

- *Stopthemes* specify themes that are not to be indexed
- *Stopclasses* define classes of alphanumeric characters that are not to be indexed, such as numbers.

Also, stoplists can now be changed dynamically, after an index has already been created.

Field Sections A new section type (*field section*), allows text within the field section to be indexed as a subdocument separate from the rest of the document.

Audio, Image, and Video

Oracle *interMedia* provides foundational support for audio, image, and video digital information. Services provided by *interMedia* facilitate the integral storage of information in these forms into the Oracle8*i* datastore as `BLOBs`, or as references to external `BFILEs`. These *interMedia* services also include methods for retrieval of this multimedia digital information, inserting new information, and for limited manipulation.

Summary of Audio, Image, and Video Support The following table summarizes the support for audio, image, and video data provided by Oracle *interMedia*.

Media Type	Description of Support
Audio	<ul style="list-style-type: none"> ■ Client access via Java Media Framework (JMF)--this allows any JMF player and <i>interMedia</i> to access and play audio files within an application. ■ Audio delivery through any streaming server such as Oracle Video Server or RealNetworks RealAudio Server. ■ Basic parsing of AUFF, AIFF, AIFF-C, and WAVE formats. ■ Management of audio data from a variety of sources and formats external or internal to Oracle8<i>i</i>; e.g., <code>BFILEs</code>, <code>LOBs</code>, URLs, or specialized servers. ■ Client-side access components in Java.
Image (Note that Oracle's Image Cartridge has been subsumed by Oracle <i>interMedia</i> .)	<ul style="list-style-type: none"> ■ Support for popular image formats--such as BMP, TIFF, GIF--and compression formats--such as CCITT, JPEG. ■ Support for Live Picture FlashPix format server. ■ Simplified Java and C++ access to images stored in Oracle8<i>i</i>. ■ Management of image data from a variety of sources and formats external or internal to Oracle8<i>i</i>; e.g., <code>BFILEs</code>, <code>LOBs</code>, URLs, or specialized servers. ■ Metadata extraction from image header information, conversion among image and compression formats, and support for nonstandard and custom image formats through a raw pixel format. ■ Direct access to pixel data in an image. ■ Support for basic manipulation functions including scaling and cropping.

Media Type	Description of Support
Video	<ul style="list-style-type: none"> <li data-bbox="644 267 1306 348">■ Client access via Java Media Framework (JMF)—this allows any JMF player and <i>interMedia</i> to access and play video files within an application. <li data-bbox="644 361 1279 413">■ Video delivery through any streaming server such as Oracle Video Server or RealNetworks RealVideo Server. <li data-bbox="644 425 1190 477">■ Support for popular video formats such as AVI, QuickTime, and MPEG. <li data-bbox="644 489 1310 572">■ Management of video data from a variety of sources and formats external or internal to Oracle8i; e.g., BFILES, LOBS, URLs, or specialized servers. <li data-bbox="644 585 1093 611">■ Client-side access components in Java.

For more information, see *Oracle8i interMedia Audio, Image, and Video User's Guide and Reference*.

Oracle *interMedia* Audio, Image, and Video Java Client Where a greater degree of manipulation or modification of image, audio, or video data is required, Oracle provides the *Oracle interMedia Audio, Image, and Video Java Client*.

The Oracle *interMedia* Audio, Image, and Video Java Client allows the user to use local (client side) applications to manipulate or modify multimedia data stored in a network-accessible (server side) database. It allows the user to connect to a server side multimedia object, copy that object from the server side to the client side, perform various operations on the client side object, and transfer the new multimedia object back to the server side.

For situations where the user does not have permission to modify the server side object, Oracle *interMedia* Audio, Image, and Video Java Client can retrieve the multimedia data from the server side for display purposes only.

In addition, Oracle *interMedia* Audio, Image, and Video Java Client gives the user the ability to integrate multimedia objects with various media frameworks, such as the Java Media Framework (JMF), RealPlayer, or Oracle Video Client (OVC).

For more information, see *Oracle8i interMedia Audio, Image, and Video Java Client User's Guide and Reference*.

Locator

Oracle *interMedia* Locator is the component of Oracle *interMedia* that supports online *geocoding* facilities for locator applications and proximity queries. It supports

geocoding of spatial point data of interest, stores this geocoded data locally in the Oracle8i database server, and performs within-distance queries on the geocoded data.

Geocoding represents addresses and locations of interest (postal codes, demographic regions, etc.) as geometric features (points). Geocoding services can be used to add the exact location (latitude and longitude) associated with points of interest to existing data files stored in Oracle8i. Once data is geocoded, Oracle *interMedia* Locator enables distances to be calculated and sites to be represented graphically in Web, data warehousing, customer information system, and enterprise resource planning applications.

A sample application might use *interMedia* Locator to locate restaurants and hotels within a given point-to-point distance of a specified location, such as a business or tourist information kiosk.

Locator is not designed to be an end-user application in itself. It consists of an *interMedia* locator object datatype, a geocode result object datatype, geocoding functions, a spatial locator index, and a spatial operator for performing within-distance queries. It supports only point geometry types. Web application developers can use the basic function provided here to build specialized Web-enabled *interMedia* Locator applications. Oracle *interMedia* Locator also supports server-based geocoding and data scrubbing operations for data warehouse applications.

For more information, see *Oracle8i interMedia Locator User's Guide and Reference*.

Oracle Spatial

Spatial data is any data with a location component. The location component could be the geocoded addresses of customers or suppliers, the course of a river and the outline of its floodplain, the locations of thousands of utility poles, or X and Y coordinates on a blueprint. Spatial data is not limited to the land surface but includes the subsurface, aquatic, marine, and lower atmospheric regions.

Oracle Spatial is designed for two groups of users:

- It enables traditional database customers to add useful spatial queries to their applications.
- It supports geographic information system (GIS) vendors who must store, retrieve, and manage very large spatial databases containing hundreds of gigabytes of geodata.

This section describes the enhancements to Oracle Spatial, formerly known as the Oracle8 Spatial Cartridge, for the current release. Additionally, this release completes the implementation of the OpenGIS Consortium RFP1 - Simple Features guidelines. Oracle Spatial conforms to the specifications associated with SQL92 representation of points, lines, and polygons.

For more information, see *Oracle Spatial User's Guide and Reference*.

VARRAY Storage Model

The structure of the spatial data model has changed such that ordinates are stored using the `VARRAY` storage model. A spatial object type can be created that represents single element or multi-element geometries. Where in release 8.0, the cartridge had individual columns for geometry attributes, Oracle8i combines these into a single column of type `SDO_GEOMETRY`, similar to the way Oracle *interMedia* defines specialized objects for images, text, audio, and video.

Dynamic Window Queries

For dynamic window queries, the release 8.0 cartridge required that users create and maintain a window layer that stored the definition of an area-of-interest and its corresponding spatial index entries. In Oracle8i, the spatial operators take a bind variable of type `SDO_GEOMETRY`, and the index creation and usage are handled internally.

Spatial Index Mechanism

The spatial index mechanism now uses Oracle8i extensible indexing features, including support for maintaining the index on insert, update, and delete. As a result, it is no longer necessary to explicitly call `SDO_ADMIN.POPULATE_INDEX()` or `SDO_ADMIN.UPDATE_INDEX()`.

RELATE() Operator

The `RELATE()` operator has been enhanced such that it can perform both primary and secondary filter operations. At query time, the Oracle8i extensible optimizer then determines which process to use based upon information in the query plan. The new syntax for the `RELATE()` operator greatly simplifies queries.

Extended Functionality

Two new spatial operators have been added.

- `SDO_FILTER()` implements a primary filter by determining whether the minimum bounding rectangles of the objects interact.
- `WITHIN_DISTANCE()` determines whether two geometries are within a specified Euclidean distance from each other.

Circle and *circular arc* geometric datatypes have been added, and several new spatial functions further extend the capabilities of this product.

Oracle Visual Information Retrieval

Oracle Visual Information Retrieval provides content-based retrieval for images stored in Oracle8i. Content-based retrieval--where the query takes the form "find me objects that look like this one"--reduces the time and effort required to obtain image-based information.

Users can tap into image content with self-service applications using flexible query criteria. They can locate similar images in large databases by searching visual criteria, such as color, pattern, and texture. Examples of database applications where content-based retrieval is useful include:

- Tracking images in stock photo houses, ad agencies, libraries, museums, and online networks
- Finding product information in design, manufacturing, and online shopping

Oracle Visual Information Retrieval includes foundational datatype support for images in Oracle8i, that complements and is completely compatible with Oracle *interMedia*. Oracle Visual Information Retrieval is based on the VIR™ technology from Virage, Inc., a leader in content-based retrieval software.

This section describes the enhancements to the Oracle8 Visual Information Retrieval Cartridge, now repackaged as Oracle Visual Information Retrieval. Also available with this release is the Oracle Visual Information Retrieval Java Client, similar in function to that described in "[Oracle interMedia Audio, Image, and Video Java Client](#)" on page 3-9.

For more information, see *Oracle8i Visual Information Retrieval User's Guide and Reference* and *Oracle8i Visual Information Retrieval Java Client User's Guide and Reference*.

New Extensible Indexing

Oracle Visual Information Retrieval allows users to create an index to support image objects. Oracle8i and Visual Information Retrieval cooperate to define, build, and maintain an index for image data. Once created, the index automatically

updates every time an image is inserted or removed from the database table. This feature significantly speeds up query time and improves performance.

Temporary LOB Usage

Consider the case where a thin client requests a scaled-down image from the database. Rather than storing the scaled-down image in a table, it is now possible to use temporary LOBs, described in "[Temporary LOBs](#)" on page 3-33. Methods are provided which allocate and operate on temporary LOBs. Additional performance benefits are achieved by utilizing OCI LOB buffering routines provided in Oracle8i, instead of local buffering mechanisms.

Scaling Enhancements

Scaling is now allowed to absolute dimensions. The process() method in release 8.0 only allowed scaling by a floating point number. The new scaling mechanism allows, for example, a much easier method to create a column of thumbnails for a table with images of various sizes in it.

Proprietary Image Formats

Because it is not practical to provide native support for all of the image formats in use today, Visual Information Retrieval provides a canonical raw image format gateway to ease the import of proprietary image data. This benefits the many image processing programs that can directly support input and output of uncompressed image data.

Oracle Time Series

Oracle Time Series enables time-stamped data to be stored efficiently in an Oracle database. It supports a basic set of functions--calendar, time series, and time scaling--to retrieve and process data. Its collection-based interface enables product developers and third-party providers to add functions to the base cartridge. By offering basic functions and datatypes, it allows customers and independent software vendors (ISVs) to develop data provisioning, data warehousing, and full-analysis time series applications using the power of Oracle8i.

Oracle Time Series provides the following kinds of functions:

- Calendar functions provide a convenient mechanism for defining time-related operations and ensuring the validity of time-related data. For example, arbitrary calendars can be defined over a fixed interval, such as a calendar of business days in the week or a calendar of quarterly dividend payment dates.

- Time series functions provide analysis of time series data and include support for complex aggregation (such as moving average), mathematical operations (such as cumulative sum and cell-by-cell arithmetic operations), and data verification.
- Time scaling functions allow a time series function to be transformed from one time scale to another, such as from aggregate daily data into quarterly summaries. Time scaling is flexible because the source and target scales are determined by calendars, which can be customized.

A summary of the new features of Oracle Time Series, formerly called the Oracle8 Time Series Cartridge, is presented below.

For more information, see *Oracle8i Time Series User's Guide*.

Irregular Time Series

Support is provided for irregular time series, that is, time series without an associated calendar. Irregular time series provide a means to conveniently handle high volumes of unpredictable data, as well as many applications with predictable data where timestamp validation is not needed.

Additional Calendar Frequencies

Added are *week*, *10-day*, *semimonthly*, *quarter*, and *semiannual* calendar frequencies.

Flexible Timestamp Precisions

Flexible timestamp precisions in Oracle Time Series relax the precision requirements of the previous release. For example, a monthly calendar now can be defined where all timestamps are on almost any specific day of the month (such as the 15th), whereas before they had to be on the first of the month.

Flexible Patterns

The calendar pattern has been extended to provide flexible support of a repeated pattern of included and omitted timestamps. The pattern specification extensions allow for more semantic (behavioral) options with *scaleup* and *scaledown* functions.

Time Scaling Enhancements

Time scaling enhancements in Oracle8i include the implementation of *scaledown* functions and additional overloads of scaling functions to provide an alternative to SQL null semantics.

Nested Table Support

Previously, the Time Series Cartridge supported a storage model for time series data based on a flat index-organized table (IOT). Oracle Time Series provides alternative support for a nested IOT model, enabling time series data to be stored in a nested IOT. Either method has advantages and disadvantages that must be weighed by the user to select the method providing the better performance and scalability.

Ease of Use

Administrative tools and procedures are provided to simplify the creation and use of time series schema objects. For most users, these procedures are a convenient alternative to explicitly defining the tables, views, and triggers used by Oracle Time Series.

Java

Java is becoming the standard language of the Internet. In Oracle8i, Oracle delivers an enterprise-class Java platform to develop and deploy Internet applications. Additionally, Oracle is making a significant strategic commitment to Java and is integrating it in a large part of its product offering. Oracle8i's Java offering includes:

- Oracle JServer Option--this is the Java Virtual Machine (Java VM), which runs within the Oracle8i database server address space
- Oracle JServer Accelerator Option--a native code compiler that speeds up the execution of Java code by eliminating interpreter overhead (scheduled for a later release)
- Programmatic interfaces--JDBC drivers and SQLJ
- Utilities and Development Tools--some are separate products

Each of these is summarized separately below.

For more information, see the following books:

- *Oracle8i Java Developer's Guide*
- *Oracle8i Java Stored Procedures Developer's Guide*
- *Oracle8i Enterprise JavaBeans and CORBA Developer's Guide*
- *Oracle8i JDBC Developer's Guide and Reference*
- *Oracle8i SQLJ Developer's Guide and Reference*

Oracle8i's Java Virtual Machine

Oracle's Java VM is a complete JDK 1.1.6-compliant Java execution environment. The Java VM runs in the same process space and address space as the database server, sharing its memory heaps and directly accessing its relational data. This design optimizes memory use, increases throughput, and delivers an open, highly available, secure, and manageable Java server.

The Java VM provides a runtime environment for Java objects. It fully supports Java data structures, method dispatch, exception handling, language-level threads, and all the core Java class libraries. It embeds the standard Java namespace in database server schemas allowing Java programs access to Java objects stored in Oracle databases and application servers across the enterprise.

Components of the Java VM include a library manager, bytecode compiler, interpreter, class loader, verifier, native compiler (separate option), server-side JDBC driver, and server-side SQLJ translator. Only the native compiler (JServer Accelerator), JDBC driver, and SQLJ translator components are discussed in this document.

A complete description of Oracle's Java Virtual Machine can be found in *Oracle8i Java Developer's Guide*.

JServer Accelerator

JServer Accelerator is a native code compiler that speeds up the execution of Java code by eliminating interpreter overhead. It translates standard Java binaries into C programs that are processed by a platform-dependent C compiler into native libraries that the Oracle Java VM can load dynamically. Unlike just-in-time (JIT) compilers, JServer Accelerator is portable to all OS and hardware platforms. To speed up applications, the Java VM is supplied with natively compiled versions of the core Java class libraries, ORB, and JDBC.

Note: The JServer Accelerator is not available in release 8.1.5. It is planned for a later release. However, Oracle JServer enjoys the performance advantage of having all of the JDK libraries shipped with release 8.1.5 compiled with an internal version of JServer Accelerator.

Programmatic Interfaces

Oracle offers two programmatic interfaces (Application Programmer Interfaces, or APIs) for Java Developers: JDBC (Java Database Connectivity) and SQLJ (embedded SQL in Java).

JDBC Drivers

JDBC is a standard Java interface for connecting to relational databases from Java. The JDBC standard was defined by Sun Microsystems, allowing individual providers to implement and extend the standard with their own JDBC drivers. JDBC is based on the X/Open SQL Call Level Interface, and complies with the SQL92 Entry Level standard. In addition to the standard JDBC API, Oracle drivers have extensions to properties, types, and performance.

Oracle offers three versions of its JDBC driver.

Server JDBC Driver Using low-level entry points, a specially tuned JDBC driver runs directly inside the database server. As a result, it provides the fastest access to Oracle data from Java stored procedures. The Server JDBC Driver complies fully with the Sun JDBC specification. Tightly integrated with the database server, it supports Oracle-specific datatypes, NLS character sets, and stored procedures. Also, the client-side and server-side JDBC APIs are the same, allowing for flexibility in where the application is deployed.

OCI JDBC Driver This driver is meant for developers of client/server Java applications and Java-based middle tiers. It converts JDBC invocations into OCI calls, which are sent via Net8 to the Oracle database server.

The OCI JDBC Driver supports various configurations including stateful, multi-tier configurations in which browser-based applets communicate with middle-tier servlets. However, it is not downloadable and requires the installation of OCI and Net8 libraries on each client machine or on the middle-tier Java application server.

Thin JDBC Driver This driver is meant for developers of Java applications and applets. It establishes a direct connection to Oracle over Java sockets and implements a lightweight TCP/IP version of Net8 protocol. So, unlike the OCI JDBC Driver, it works only with TCP/IP-based networks.

Written in Java, the Thin JDBC Driver is about 150 K (300 K uncompressed), so it can be downloaded with Java applets into a Web browser or Network Computer. No preinstallation of software on the client is required.

SQLJ Translator

SQLJ enables SQL statements to be embedded in Java programs. SQLJ code is much more concise and easier to write than JDBC and features static analysis and type checking.

The SQLJ translator, itself a Java program, takes as input a file of Java source code in which SQLJ clauses are embedded. Then, it translates the SQLJ clauses into Java class definitions that implement the specified SQL statements. The Java type system ensures that objects of those classes are called with the correct arguments. Oracle provides both a server-side and client-side SQLJ translator.

The server-side SQLJ translator is a highly optimized SQLJ translator that runs directly inside the database server, where it provides runtime access to Oracle data via the Server JDBC Driver. SQLJ forms can include queries, DML, DDL, transaction control statements, and calls to stored procedures. The client-side and server-side SQLJ APIs are the same, allowing for flexibility in where the application is deployed.

Development Model

Java VM supports a variety of development models, including:

- Java stored procedures--to support traditional database programmers and SQL-oriented clients.
- Enterprise JavaBeans (EJB)--Oracle8i Java VM provides a transaction server platform for distributed EJBs.
- CORBA Services--Oracle8i also allows distributed systems developers to implement CORBA services in Java. It provides a number of standard CORBA and EJB services.

These are discussed in more detail in the following sections,

Java Stored Procedures

Stored procedures are Java procedures or methods published to SQL and stored in an Oracle database for general use. With stored procedures, you can implement business logic at the server level, thereby improving application performance, scalability, and security.

Except for GUI (graphical user interface) methods, any Java method can run in the database server as a stored procedure. The runtime contexts are:

- Functions and procedures

- Database triggers
- Object-relational methods

To publish Java methods, call specifications (*call specs* for short) are written, which map Java method names, parameter types, and return types to their SQL counterparts. A call spec is *not* a wrapper. Without adding a layer of execution, a call spec simply publishes the existence of a Java method. So, when the method is called (through its call spec), the runtime system dispatches the call with minimal overhead. And, because the stored procedure is executing in the database, there is efficient access to SQL data.

Anything that can be written with PL/SQL can now be written with Java and run in the database server in the same contexts as PL/SQL. The database server allows a high degree of interoperability between Java and PL/SQL. Java applications can call PL/SQL stored procedures using an embedded JDBC driver. Conversely, PL/SQL applications can call Java stored procedures directly.

Enterprise Java Beans

Enterprise Java Beans (EJB) is a component model introduced by JavaSoft that enables developers to create custom component applications. These applications consist of server-side beans developed by yourself or by third parties. The beans provide the business logic in EJB applications.

The EJB model is flexible in that you can use different components from different vendors. For example, you can use configuration and management software from one company, bean containers from a second company, and business logic beans from a third company that specializes in providing beans for just that purpose.

CORBA

CORBA is a standards-based distributed component model proposed by the Object Management Group (OMG). It supports a development environment for building, deploying, and managing distributed object applications that are interoperable across platforms.

CORBA objects communicate using OMG's Internet Inter-ORB Protocol (IIOP), the standard for communication between and among distributed objects running on the Internet, intranets, and in enterprise computing environments. Oracle8i integrates a Java-based CORBA 2.0 compliant Object Request Broker (ORB) that provides users with the ability to call into and out of the database server using CORBA's IIOP. CORBA components written in different languages, running on different platforms, can transparently communicate and interoperate.

Utilities and Development Tools

Oracle provides utilities and tools for the development and deployment of Java applications. Some of these are discussed here.

JPublisher

JPublisher is a utility, written entirely in Java, that translates user-defined types to Java wrapper classes. It is similar to the Object Type Translator (OTT) utility used in C/C++ environments. Developers are required to have Java classes that correspond to user-defined types for developing:

- SQLJ applications that use object types, collection types (VARRAY and nested table types), or REF types
- JDBC applications that use object types, collection types, or REF types

JPublisher helps create Java language applications that use user-defined types in the Oracle server. For a Java-language application to access object data, it is necessary to represent the data in a host language format. JPublisher helps do this by creating the mapping between object types and Java classes, and between object attribute types and their corresponding Java types.

Using JPublisher to automatically generate Java wrapper classes for user-defined types is less time consuming and error prone than manually coding them.

Note: Portions of JPublisher are still considered beta, and a beta version of the JPublisher User's Guide will be made available. The full functionality of this utility is planned to be available in a later release.

JDeveloper

With the release of Oracle8i, Oracle offers a new version of the JDeveloper product. JDeveloper 2.0 provides complete support for building and deploying Java applications on Oracle8i. It provides a standard GUI based Java development environment that is well integrated with Oracle's Application Server and Database.

The following are new features included in JDeveloper 2.0:

- Advanced Enterprise JavaBeans support
- Java database programming with Oracle8i
- Servlet support for transactional Web content

- Revamped interface for true multi-tier application development

JDeveloper runs on Windows NT. General availability follows the production Oracle8i release. For more information on JDeveloper 2.0, see the product-specific documentation.

loadjava and dropjava

Provided with Java VM, the statement-line utilities loadjava and dropjava manage Java library units. Java stored procedures can be written, compiled, and partially tested and debugged on the client side in JDeveloper, or other popular Java IDE. Then, the statement-line utility loadjava can be used to upload the resulting Java sources, binaries, and resources into the database server. The companion utility dropjava can be used to drop given Java sources, binaries, and resources from a schema.

Data Warehousing and Very Large Data Bases (VLDB)

Oracle8i offers many new features and enhancements for the data warehousing and VLDB environment. The performance of queries is dramatically increased by materialized views and summary management, which allow frequently requested summaries to be stored and maintained. Parallelism is improved with more operations offered in parallel. The new database resource management facility, discussed in "[Database Resource Management](#)" on page 3-48, provides an effective way of managing CPU usage.

Moving data from one database to another is made simpler and faster with transportable tablespaces. New data mining and relational online analytical processing (ROLAP) functions are introduced. If you are using the partitioning option, then see the "[Partitioning Enhancements](#)" section on page 3-45. This discusses two new methods of partitioning, providing more partitioning options for the data warehouse environment.

Only a few of the many Oracle8i changes relating to data warehousing and the VLDB environment have been mentioned above. Read on for descriptions of these and allied new features and enhancements introduced in Oracle8i.

Summary Management Using Materialized Views

A *materialized view* is a stored summary containing precomputed results. Materialized views allow for significantly faster data warehouse query processing. With the Oracle8i summary management feature, the Oracle database server automatically rewrites queries to use the summary data, rather than retrieving data

from detail tables by doing expensive joins and aggregate operations. This *query rewrite* facility is totally transparent to the application, which is not aware of the existence of the materialized view.

The DBA's first step in creating materialized views is to define *dimensions*. These represent the hierarchies that are present in the real world; for instance, multiple months make up a quarter, multiple districts make up a region, etc. The `CREATE MATERIALIZED VIEW` statement is used to create a materialized view. This statement includes a subquery, typically a join or a data aggregation (`GROUP BY`), the results of which comprise the materialized view.

A materialized view is maintained by a refresh process. The refresh process can be done automatically when a commit is done on a detail table, or it can be controlled manually by the DBA. A refresh is specified as *complete* or *incremental*. A complete refresh truncates existing data, then repopulates the summary with new data from the detail tables. An incremental refresh updates only changed data.

To create a materialized view in a user's own schema, the user must have the `CREATE SNAPSHOT` or `CREATE MATERIALIZED VIEW`, `CREATE TABLE`, `CREATE INDEX`, and `CREATE VIEW` system privileges. To create a materialized view in another user's schema, a user must have the `CREATE ANY SNAPSHOT` or `CREATE ANY MATERIALIZED VIEW` system privilege.

To enable a materialized view for query rewrite:

- If all the master tables in the materialized view are in a user's schema, that user must have the `QUERY REWRITE` privilege.
- If any of the master tables are in another schema, a user must have the `GLOBAL QUERY REWRITE` privilege.
- If the materialized view is in another user's schema, both users must have the appropriate `QUERY REWRITE` privilege described in the preceding two items.

The schema that contains the materialized view must have sufficient quota in the target tablespace to store the materialized view's base table and index or have the `UNLIMITED TABLESPACE` system privilege.

To create and refresh a materialized view, both the creator and the materialized view owner must be able to issue the defining query of the materialized view. This capability depends directly on the database link that the materialized view's defining query uses.

For ease in managing materialized views, Oracle provides a collection of materialized view analysis and advisory functions that are callable from any

PL/SQL program. These are packaged in the `DBMS_OLAP` package, which provides recommendations and supporting data such as the following:

- Summary usage--such as number of times a query rewrite was done to use a summary, the space used by a summary, and a cost-benefit ratio for each summary.
- Summary recommendations--such as creation, retention, and dropping of summaries.

For more information, see *Oracle8i Concepts*, *Oracle8i SQL Reference*, and *Oracle8i Supplied PL/SQL Packages Reference*.

Transportable Tablespaces

This feature allows a user to move a subset of an Oracle database into another Oracle database. It is a lot like unplugging the subset from the original database and plugging it into another one. It is also possible to clone a tablespace in one database and plug it into another, thereby copying the tablespace between databases.

Moving data using the transportable tablespace feature can be orders of magnitude faster than either export/import or unload/load of the same data, because it involves only copying data files and integrating the tablespace metadata. When transporting tablespaces, index data can also be moved, eliminating the need to rebuild indexes after importing or loading the table data.

In the data warehouse environment, where huge amounts of data flow from the initial OLTP databases into the enterprise data warehouse and on to data marts, this feature presents an opportunity for a faster and more innovative means of data movement. And, its use can be extended to the archiving of data or publishing data to other databases.

For more information, see *Oracle8i Concepts* and the *Oracle8i Administrator's Guide*.

Direct Path Load API

The direct path load API is a set of OCI interfaces that provides an application access to the direct path load engine in the Oracle server. Previously, the only client of the direct path load engine has been Oracle's SQL*Loader utility. Now, it is possible for ISVs to develop applications which can take advantage of the performance benefits provided by the direct load engine, including parallel load. There is no longer a need to write SQL*Loader control and data files in order to access the direct path load engine.

For more information, see *Oracle8i Utilities* and the *Oracle Call Interface Programmer's Guide*.

ROLAP Enhancement: Extended Aggregate Operations

The `GROUP BY` clause of SQL queries can now specify additional OLAP aggregate operations. These new operations are called `ROLLUP` and `CUBE`, and are effective for computing subtotals and cross-tabulations across multiple dimensions.

While these same operations could be done using previous SQL syntax, the SQL was difficult to generate, and the execution was inefficient because the optimizer had no hints about what the user wanted to accomplish. These new operations allow client reporting and analysis tools to perform aggregate operations with less effort and greater efficiency. They are useful for creating materialized views.

For more information, see *Oracle8i Concepts* and *Oracle8i SQL Reference*.

Data Mining: SAMPLE Function

For certain types of queries, listed below, the `FROM` clause also supports a `SAMPLE` clause. The `SAMPLE` clause allows a user to specify that the results are to be based on a random (row or block) sample of a whole table. `SAMPLE` may be specified for:

- Simple select queries (no joins)
- Simple create table as select queries

This feature can be a useful for data mining tools, or in a data warehousing environment where it can be used to avoid full table scans.

For more information, see *Oracle8i Concepts* and *Oracle8i SQL Reference*.

Function-Based Indexes

Indexes can now be created on functions and expressions that involve one or more columns in the table being indexed. A *function-based index* precomputes the value of the function or expression and stores it in the index. In a data warehouse, query performance is dramatically improved when a computationally intensive expression is used to build the index, rather than using that expression in the `WHERE` clause of a query.

The function used for building the index can be:

- An arithmetic expression.

- An expression that contains a PL/SQL or Java function, package function, C callout, or SQL function.
- A method inside an object, if building an index on that object.

For more information, see *Oracle8i Concepts*.

Descending Indexes

The `DESC` keyword on the `CREATE INDEX` statement is no longer ignored. It specifies that the index should be created in descending order. Indexes on character data are created in descending order of the character values in the database character set. Neither this, nor the `ASC` keyword, may be specified for a domain index. `DESC` cannot be specified for a bitmapped index.

For more information, see *Oracle8i SQL Reference*.

Enhancements for Bitmap Indexes

Several changes enhance query performance when bitmap indexes are used. One of these changes lifts a restriction whereby the use of the `ALTER TABLE` statement could invalidate a bitmap index.

A new feature introduces two new `ALTER TABLE` statement options:

- `MINIMIZE RECORDS PER BLOCK`
- `NOMINIMIZE RECORDS PER BLOCK`

These options allow for tuning of the rowid-to-bitmap mapping. The `MINIMIZE` option can be used to optimize bitmap indexes for a query-only environment by requesting the most efficient possible mapping of bits to rowids. The `NOMINIMIZE` option disables this optimization.

For more information, see *Oracle8i SQL Reference* and *Oracle8i Designing and Tuning for Performance*.

Parallelization of Aggregate Distinct Queries

The parallelization of aggregate distinct queries is now possible even if they do not contain a `GROUP BY` clause. This is done through new hints for specifying distribution methods.

For more information, see *Oracle8i Designing and Tuning for Performance*.

Sort Improvements

The sort routines used throughout the Oracle database server have been rewritten to eliminate performance problems. Sort now performs more predictably and consistently, given the same input, system configuration, and parameter settings. Memory is utilized more efficiently, enabling more rapid sorts when additional memory is specified. The following table shows initialization parameters that have been dropped or added as a result of this rewrite:

Dropped Initialization Parameters	Added Initialization Parameter
<code>SORT_DIRECT_WRITES</code>	<code>SORT_MULTIBLOCK_READ_COUNT</code>
<code>SORT_WRITE_BUFFERS</code>	
<code>SORT_WRITE_BUFFER_SIZE</code>	
<code>SORT_READ_FAC</code>	
<code>SORT_SPACEMAP_SIZE</code>	

For more information, see *Oracle8i Designing and Tuning for Performance*.

Single-Table Hash Cluster

Using the new `SINGLE TABLE` option when creating a hash cluster specifies the creation of a hash cluster with only one table. The benefit is faster access to the rows in the table.

For more information, see *Oracle8i Concepts*, or for details about the `SINGLE TABLE HASHKEYS` clause of the `CREATE CLUSTER` statement, see *Oracle8i SQL Reference*.

Data Warehouse Statistics Gathering

To support the speed requirements of queries running in a data warehouse environment, it is critical that the optimizer selects the best plan of execution out of all those that are available for consideration. The optimizer selects the best plan based on cost, as determined by the statistics which have been gathered on the data. Therefore, it is vitally important that the optimizer has good and recent statistics on the data. Oracle8i introduces new options for the gathering and manipulation of statistics.

A new package, `DBMS_STATS`, is provided for:

- External manipulation of statistics

- Gathering statistics utilizing parallel execution
- Copying of statistics between databases

The copying of statistics between databases is useful for tuning studies. Statistics can be copied from the production database into a test database. `DBMS_STATS` can also be used to back up statistics. The backup may be restored later, if necessary, or used to study the change in data characteristics over time.

Additionally, the syntax of the `CREATE` and `ALTER INDEX` statements has been changed to include a `COMPUTE STATISTICS` clause that allows statistics to be gathered incidentally, and at minimum cost, while the index is being built.

For more information, see *Oracle8i Designing and Tuning for Performance*, *Oracle8i Concepts*, *Oracle8i SQL Reference*, and *Oracle8i Supplied PL/SQL Packages Reference*.

Automated Parallel Query Tuning

This enhancement simplifies tuning of parallel query (PQ) execution and provides better system utilization. PQ users can simply specify that a query or table is to be parallelized, and the system determines the optimal single user setting, then reduce it based on the system utilization at the time the query starts execution. There are four parts to this enhancement.

- Default settings are improved for degree and instances.
- There are improved defaults for initialization parameters when the user intends to use parallel query.
- A multi-user adaptive degree algorithm is used to select an optimal degree of parallelism based on system load.
- A load balancing algorithm evenly distributes the load across nodes in an inter-node parallel query environment.

For more information, see the *Oracle8i Parallel Server Documentation Set: Oracle8i Parallel Server Concepts; Oracle8i Parallel Server Setup and Configuration Guide; Oracle8i Parallel Server Administration, Deployment, and Performance* and *Oracle8i Designing and Tuning for Performance*.

Statistics Generation for Long-Running Operations

In large databases, operations such as complex queries, index builds, and backup operations can take a long time to complete. The dynamic `V$SESSIONS_LONGOPS` view provides a way for users and DBAs to monitor the progress of such operations and to estimate their completion times. In Oracle8i, this view is enhanced by more

internal server components collecting a greater number of statistics. Additionally, the package `DBMS_APPLICATION_INFO` contains a procedure `SET_SESSION_LONGOPS` to populate the view from an application.

For more information, see *Oracle8i Reference* and *Oracle8i Supplied PL/SQL Packages Reference*.

New Constraint Functionality

The following new features are implemented through syntax changes to the `ALTER TABLE` and `CREATE TABLE` statements.

For additional information on either of these enhancements, see *Oracle8i Administrator's Guide* and *Oracle8i SQL Reference*.

Constraint Modification

It is now possible to modify the attributes of an existing constraint through the `ALTER TABLE...MODIFY CONSTRAINT` clause. This enhancement also introduces a new constraint property, `RELY/NORELY`. This flag can be used by the DBA to indicate validity of data without enabling or validating a constraint. It is used by the query rewrite function of materialized views.

DISABLE VALIDATE Constraint State

Data warehouses may benefit from the `DISABLE VALIDATE` constraint state. This state saves space because it requires no index on a unique or primary key, yet it guarantees the validity of all existing data in the table. This is at the expense of disallowing DML and index lookups, but those may be less important considerations in some situations.

The `DISABLE VALIDATE` constraint state also enables efficient loading of data from a nonpartitioned table into a partitioned table using the `EXCHANGE PARTITION` clause of the `ALTER TABLE` statement.

To enable this feature, `DISABLE/ENABLE` and `VALIDATE/NOVALIDATE` have been made independent properties of constraints, where all four combinations are now allowed.

Index-Organized Tables

There are numerous enhancements for index-organized tables intended to make them usable for very large databases and mission-critical applications. These enhancements can be broadly classified into two areas:

1. Provide full-table functionality for index-organized tables.
2. Provide the ability to do key compression.

General Enhancements

Some of the key enhancements include:

- LOB column support provides the ability to store large objects, such as text documents and images required by various data cartridges and Web-based applications.
- Secondary index support allows for efficient access using non-primary key columns.
- Index-organized tables (IOTs) can be rebuilt with the new `MOVE` clause for `ALTER TABLE`, which also supports an `ONLINE` keyword. The `ONLINE` keyword may only be used for IOTs and allows DML operations on the IOT while the primary key index is being built.
- `CREATE TABLE... AS SELECT` enables parallel loading of an index-organized table.

For more information, see *Oracle8i Concepts*.

Logical ROWIDs

`ROWIDs` provide the fastest possible access to a given row in a given table. They contain the physical address of a row down to the specific block and allow retrieval of the row in a single block access. To expand this functionality to index-organized tables (IOTS), whose rows do not have permanent physical addresses, Oracle8i introduces primary key-based logical identifiers called logical `ROWIDs`.

Rows in index-organized tables do not have permanent physical addresses--they are stored in the index leaves and can move within the block or to a different block as a result of insertions--therefore their row identifiers cannot be based on physical addresses. Instead, Oracle provides index-organized tables with logical row identifiers, called logical `ROWIDs`, that are based on the table's primary key. Oracle uses these logical `ROWIDs` for the construction of secondary indexes on index-organized tables.

Also introduced is the `UROWID` datatype, which makes logical `ROWIDs` usable by applications in the same sense that `ROWIDs` are currently used (e.g., selecting `ROWIDs` for later update or as part of a cursor). `UROWID` can be used to store `ROWIDs` from other databases accessed via gateways, and it can be used to reference physical `ROWIDs`. PL/SQL fully supports the `UROWID` datatype.

For more information, see *Oracle8i Concepts*.

Related Features

Other features are also useful in the data warehouse environment, but are discussed elsewhere in this document.

- ["Partitioning Enhancements"](#) on page 3-45
- ["Database Resource Management"](#) on page 3-48
- ["Oracle Parallel Server"](#) on page 3-61
- ["Optimizer Plan Stability"](#) on page 3-37
- ["User-Defined Operators and Extensible Indexing"](#) on page 3-31

Database Features

With Oracle8i, the Oracle database server provides a new extensibility framework, making it easier for ISVs and other developers to extend the database to meet their needs. Many object relational enhancements have been incorporated into this release, and this functionality has been extended to all Oracle8i product configurations.

Online transaction processing (OLTP) systems are usually concerned with entering and retrieving mission-critical data from day-to-day operations. Availability, reliability, and performance are extremely important for these systems. Oracle8i delivers significant enhancements in these areas, along with improvements in scalability, serviceability, and security.

Advance Queuing has been significantly enhanced, continuing to develop the publish/subscribe messaging paradigm. National Language Support now has a new book to help users better learn how to use its new and already existing functionality.

These and other new database features are discussed in this section.

Extensibility Framework

The extensibility framework in Oracle8i extends the Oracle database server to support specialized datatypes and to provide core services for external use. Developers are able to provide their own implementations of database server services and register them with the server. The framework is safe and protects the integrity of Oracle8i.

ISVs and other developers can exploit this extensibility framework to more easily build user-defined software components, also referred to as data cartridges, that are easier to use, and that integrate seamlessly with each other and the Oracle database server. Multimedia users especially can benefit from the greater flexibility and functionality provided by this framework as it enables independent software components to interoperate in a seamless manner. For example, a query can now be based on both image and spatial data.

This section summarizes the features and enhancements that comprise the Oracle8i extensibility framework. For more information on the functionality described below, see *Oracle8i Data Cartridge Developer's Guide* and *Oracle8i Concepts*.

User-Defined Operators and Extensible Indexing

Increasingly, databases are being used to store more varied types of data, such as spatial, audio, or video data. This brings a need for indexing complex datatypes and for specialized indexing techniques. Oracle8i provides an interface that enables developers to define domain-specific operators and indexing schemes and integrate them into the Oracle database server.

Oracle provides a set of built-in operators, for use in SQL statements, which include arithmetic operators (+, -, *, /), comparison operators (=, >, <), logical operators (NOT, AND, OR), and set operators (UNION). These operators take as input one or more arguments (operands) and return a result. The extensibility framework of Oracle8i allows developers to define new operators. Their implementation is provided by the user, but the Oracle database server allows these user-defined operators to be used in SQL statements in the same manner as any of the predefined operators provided by Oracle.

The framework to develop new index types is based on the concept of cooperative indexing, where an application and the Oracle database server cooperate to build and maintain indexes for complex datatypes. The application software is responsible for defining the index structure, maintaining the index content during load and update operations, and searching the index during query processing. The index structure itself can be stored either in the Oracle database as tables or externally as files. Indexes created using these new index types are referred to as *domain indexes*.

Import/export support is provided for domain indexes. Indexes (including domain indexes) are exported by exporting the index definitions, namely the corresponding CREATE INDEX statements. Indexes are recreated by issuing the exported CREATE INDEX statements at the time of import.

Because domain index data stored in database objects, such as tables, is exported, there is a fast rebuild of domain indexes at import time.

For specific information about the extensions to SQL that support this, see *Oracle8i SQL Reference*.

Extensible Optimizer

The optimizer functionality has been extended to allow authors of user-defined functions and domain indexes to create statistics collection, selectivity, and cost functions that the optimizer uses in choosing a query plan.

Statistics Collection Functions The `ANALYZE` statement can now make a call to a user-specified statistics collection function whenever a domain index is analyzed. User-defined statistics collection functions can also be defined for individual columns of a table and for user-defined datatypes.

Selectivity Functions Users can specify user-defined selectivity functions for predicates containing user-defined operators, standalone functions, or type methods.

Cost Functions Users can define costs for domain indexes and for user-defined standalone functions, package functions, and type methods. The user-defined costs can be in the form of default costs that the optimizer simply looks up, or they can be computed by functions called by the optimizer.

Cartridge Services

In addition to the above, other basic services have been packaged for use by user-defined components. These, and the other extensibility and object enhancements in Oracle8i, have allowed the uniform integration of Oracle's former data cartridge options (ConText, Image, Visual, Time Series, Spatial) into the Oracle database server.

The services include:

- Memory management
- Context management
- Parameter management
- String and numbers manipulation
- File I/O

- Internationalization
- Error reporting
- Thread management

Additionally, the Oracle Software Packager, discussed in "[Oracle Software Packager and Oracle Universal Installer](#)" on page 3-58, allows cartridges to be packaged and installed into Oracle8i.

Object Relational Enhancements

The object relational paradigm introduced in Oracle8 gains new functionality in Oracle8i. Additionally, object relational technology is now available in all configurations of the Oracle database server. Enhancements target more flexible storage and retrieval of objects, and there is greater transparency between object views and object tables. Specifics of these changes are described below.

For more information, see *Oracle8i Concepts*, *Oracle8i Application Developer's Guide - Fundamentals*, and *PL/SQL User's Guide and Reference*. Additional references may be provided below in specific descriptions.

LOB Enhancements

Several LOB-related enhancements are included in this release. Additional references for these include *Oracle Call Interface Programmer's Guide*, *Oracle8i Application Developer's Guide - Large Objects (LOBs)*, and *Oracle8i SQL Reference*.

Temporary LOBs Oracle8 provided support for permanently storing large unstructured data by means of large object (LOB) datatypes. But many applications have a need for temporary LOBs that act like local variables. These LOBs do not exist permanently in the database, but exist mainly for the purpose of performing transformations on LOB data.

Oracle8i supports the creation, freeing, access, and update of temporary LOBs. Their default lifetime is a session, but they may be explicitly freed sooner by the application. They are ideal as transient workspaces for data manipulation, and because no logging is done or redo records generated, they have better performance than persistent LOBs.

Other Enhancements

- A new conversion function, `TO_LOB`, can be used to convert a `LONG` to a LOB.

Note: LONG column support will be discontinued in a future Oracle release. Users are advised to migrate LONG data into LOB columns.

- Usable LOB chunk size is exposed to users via an OCI interface and in DBMS_LOB.
- Open/Close/IsOpen APIs (OCI or DBMS_LOB) for internal LOBs have been added.
- Variable-width character sets are supported for CLOBs and NCLOBs.
- A DBMS_LOB package and an OCI API are provided to allow users to append data to the end of the LOB value.

User-Defined Object Identifiers

The system-generated object identifiers (OIDs) provide globally unique identifiers that are important for many applications, especially distributed processing and replication. However, they carry an overhead of increased storage requirements, and increased loading time. For applications that do not require the advantages of system-generated OIDs and which must save on overhead, Oracle8i allows users to specify their own identifier, such as a primary key, as an OID.

A related enhancement is referential integrity constraint support on REF columns.

Nested Table Enhancements

These enhancements target making the handling and storing of large nested tables more efficient. Additional references for nested table enhancements are the *Oracle Call Interface Programmer's Guide*, *Oracle8i SQL Reference*, and *Oracle8i Supplied PL/SQL Packages Reference*.

- DDL is enhanced to allow specifying that the rows of a nested table be stored in an index-organized table. Constraints may be specified on the table. Storing the nested table data in an index-organized table provides clustering of the nested table data and is an efficient storage organization, especially when the data is always retrieved based on the primary key.
- The user may specify at the time of table creation that a collection locator is to be returned when a nested table column or attribute is fetched. An OCI call (OCICollIsLocator()) and a UTL_COLL package (IS_LOCATOR function) are provided to indicate whether a collection is locator based or not. An application or PL/SQL program can now check whether a nested table attribute or variable

is locator based, before performing certain collection operations that could cause a large nested table value to be materialized in memory.

Varray Enhancements

Oracle8i provides for more efficient handling and storage of large varrays. Varray storage can now be specified. If specified, it is stored out of line in a LOB and the user can specify the storage characteristics of the LOB.

Collection Unnesting

The notion of collection datatypes (nested tables and varrays) was introduced in Oracle8. Users are allowed to create columns of collection types, and functions and procedures with parameters of collection types. When it is necessary to view these collection types in a flat (relational) form, Oracle provides a way of unnesting them.

This unnesting is implemented using the `TABLE()` operator found in the *table_collection_expression* of the SQL statement. When introduced in release 8.0, the `TABLE()` expression took an argument of a set of rows and allowed that set to be treated like a table. Oracle8i extends this functionality so that a collection column of a table in a query may be used like an independent table.

Note: `TABLE()` subsumes the `THE(subquery)` expression, thereby deprecating its usage.

Enhancement of Object Views

Object views are an extension to the basic relational view mechanism to provide an object abstraction over existing relational and object data. They allow for the retrieval, update, insertion, and deletion of relational data as if such data were stored as object types. In this release, object views are enhanced to provide more features than are available in object tables, such as, updatability of nested tables and creation of new objects in the OCI object cache. The performance of object views is improved to be comparable with that of object or relational tables.

Triggers on Nested Table View Columns

Triggers are PL/SQL procedures which are fired when various events occur in the system, such as updates to a row. In release 8.0, the trigger mechanism was defined over rows of tables or as an `INSTEAD OF` mechanism over views. Oracle8i expands the application of triggers to allow them to be defined over nested table view

columns. As a result, it is now possible to update an element of a collection synthesized using the `CAST...MULTISET` operation.

Object Cache Enhancements

The following enhancements are options that can be set for the object cache:

- The ability to detect whether an object in the cache has been changed in the server by another committed transaction
- The ability to set all of the attributes of an object to `NOT NULL` at the time of creation
- To prevent a user from blocking while trying to lock an object, `NOWAIT` can be specified when trying to lock an object
- Object cache extensibility -- allow cache clients to use their own object representations instead of C structs

For more information, see *Oracle Call Interface Programmer's Guide*.

ANALYZE Statement, Dangling REFs

The `ANALYZE` statement can now set dangling `REFs` to `NULL`. In this case, these `REFs` no longer raise an error condition and terminate execution.

For more information, see *Oracle8i SQL Reference*.

Subqueries in VALUES Clause of INSERT Statement

In release 8.0, if a relational table or an object table contained `REF` type columns, then there was no simple way of inserting a relational row or a row object into the table. First `NULLs` were required to be inserted into the columns and then an update statement was issued to set their values. This inefficiency is resolved in *Oracle8i*, by allowing a value subquery expression in the `INSERT...VALUES` list.

A value subquery expression may also be used as an argument to a function or a type constructor.

SQL*Loader Support for Objects, Collections, and LOBs

The `SQL*Loader`, used to move data from external files into tables in an Oracle database, has been extended to support the loading of objects, collections, and `LOBs`. For a summary of `SQL*Loader` enhancements, see "[SQL*Loader Enhancements](#)" on page 3-56.

For more information, see *Oracle8i Utilities*.

Other Object Relational Enhancements

Partitioning of Tables Containing Objects The partitioning of tables with columns containing objects, collections (varrays and nested tables), and `REFs` is provided in Oracle8i. The partitioning of object tables is also supported.

Parallel Query Support Oracle8i provides parallel query support on object tables.

Optimizer Plan Stability

After carefully tuning an application, a user may wish to ensure that the optimizer generates the same execution plan whenever the same SQL statements are executed. Oracle8i provides a method for stabilizing execution plans, regardless of changes to the system configuration, parameters, statistics, or even changes to the optimizer itself. This also benefits high-end OLTP sites by having SQL execute without having to invoke the cost-based optimizer at each SQL invocation.

This feature is especially useful for ISVs, who can distribute the execution plans of the SQL in their package applications so that all installed sites are executing the same execution plans.

Through a new SQL statement, `CREATE OUTLINE`, the user is able to create a *stored outline* containing a set of attributes used by the optimizer to create an execution plan. Stored outlines can also be created automatically by setting the system parameter `CREATE_STORED_OUTLINES` to `TRUE`. The system parameter `USE_STORED_OUTLINES` can be set to `TRUE`, `FALSE`, or to a category name to indicate whether to make use of existing stored outlines for queries that are being executed. The `OUTLN_PKG` package provides procedures used for managing stored outlines.

For more information, see *Oracle8i SQL Reference*, *Oracle8i Designing and Tuning for Performance*, and *Oracle8i Supplied PL/SQL Packages Reference*.

Advanced Queuing

Advanced Queuing (AQ) was introduced in release 8.0 of Oracle8 Enterprise Edition. It integrated a message queuing system with the Oracle database through a set of queue tables and queue functions, and allowed for either synchronous or asynchronous communication between programs. Oracle8i provides enhancements to the functionality of Advanced Queuing. These are discussed in this section.

For more information, see *Oracle8i Application Developer's Guide - Advanced Queuing*.

AQ-Based Publish/Subscribe

Additional features are introduced in Oracle8i Advanced Queuing to support a publish/subscribe messaging paradigm. In this model, database queues serve as a repository for the messages, rules, and control information. There is a decoupling of addressing between message senders and message recipients. It is the recipients who specify which messages they wish to receive. Sending applications anonymously publish messages and receiving applications independently and autonomously subscribe to such messages.

The components of this publish/subscribe model are:

- Publishers

Publishers are entities that publish information (as messages) to queues. They do not know or care about the interest of other applications in the messages; they only know that they are to publish the information. There are two types of publishing: database event publishing where triggers cause messages to be enqueued when certain database or system events occur, and publishing from an application. Database event publishing is new in Oracle8i and is discussed separately later.

- Subscribers

Subscribers are entities that receive information by expressing interest in certain types of messages. They do not care about the origins of the messages. They specify certain filtering rules for messages, allowing them to receive only the messages whose content satisfies the rules.

- Rules Engine

The rules engine, new in Oracle8i, provides rules-based subscriptions by evaluating incoming messages from publishers against a set of rules specified by subscribers. Only if the message satisfies a rule is it delivered to the subscriber. Subscribers can express rules using the full power of SQL. Since AQ messages are normal database objects, any SQL operation that applies to data objects can also be applied to messages.

An application (subscriber/client) may receive messages that are published in one of three ways: explicit dequeuing, dequeuing with listen, or by using OCI functions to register a callback.

For more information, see *Oracle Call Interface Programmer's Guide*.

Database Event Publication

The database event publication feature allows applications to subscribe to database events just as they subscribe to messages from other applications. Trigger syntax is extended to support system and other data events on DATABASE and SCHEMA. Trigger syntax also supports a CALL to a procedure as the trigger body.

Users can enable the publication of the following events.

- DML events (DELETE, INSERT, UPDATE)
- DDL Events (CREATE, ALTER, DROP)
- Database events (SERVERERROR, LOGON, LOGOFF, STARTUP, SHUTDOWN)

The database event publication subsystem tightly integrates with the AQ publish/subscribe engine. For a complete description of triggers for data and system events, see *Oracle8i SQL Reference*.

Enhanced Propagation

For increased scalability, the performance of basic queue operations has been enhanced and changes have been made to underlying data structures. Message propagation scheduling is improved. AQ propagation uses message streaming and eliminates the current two-phase commit mechanism. Additionally, the identity of the sender is now made part of message properties. This meets a proposed standard for messaging services and enables messaging services such as TIBCO's RVCN to distribute messages in near real time to a large number of applications.

Statistics are now maintained on the scheduled propagation, enhancing the administrator's ability to manage and tune propagation. Also, a single view is provided for the administrator that can be used to determine the propagation schedule, its status, and last error.

There is better history management for multi-consumer queues, and the history itself now contains a new message state and property for distinguishing between a message being propagated and its being processed.

Miscellaneous Enhancements

Other Advanced Queuing enhancements are as follows:

- For multiple consumer queues, a user is allowed the option of enabling or disabling the collection of statistics on the number of messages per recipient.
- System and object level security control are provided. The AQ security model is now a natural extension of the Oracle8i security model.

- Performance is improved when running in parallel server mode.
- The creation of nonpersistent queues is allowed.

DBMS_REPAIR Package

Oracle8i provides enhanced block corruption repair capability through the new DBMS_REPAIR package. It provides the DBA with a three-stage approach to addressing corruptions.

1. Detect and report corruptions.
2. Make the object usable.
3. Repair corruptions and rebuild lost data.

Depending upon the nature of the repair, data may be lost and logical inconsistencies may be introduced. The DBA must determine whether the potential loss warrants the use of this feature, and to help the DBA, this package also provides a report mechanism that aids in the determination of whether to use this feature or not.

For information on this feature, see *Oracle8i Administrator's Guide* and *Oracle8i Supplied PL/SQL Packages Reference*.

Redo Log Analysis Using LogMiner

Log files contain a wealth of useful information about the activities and history of an Oracle database, but until Oracle8i there has been no easy tool that could tap into this information. LogMiner allows online and archived redo log files to be read, analyzed, and interpreted by the user using SQL. Analysis of the log files with LogMiner can be used for the following:

- It can be used to track specific sets of changes based on transaction, user, table, time, etc. It is possible to determine who modified a database object and what the before and after data was. Understanding and being able to trace database changes back to their source and being able to undo them are valuable security and management tools.
- It can help pinpoint when a logical corruption was introduced into the database. This is important for determining how to initiate recovery for restoring the database to a consistent state prior to corruption.
- This data can also be used as a supplemental source of data for tuning and capacity planning. Performing various forms of historical analysis to determine trends and data address patterns can be done.

No additional collection overhead is incurred to obtain data for LogMiner.

For more information, see *Oracle8i Backup and Recovery Guide*.

Drop Column

A new drop column capability allows the DBA to easily remove unused columns in the database. Previously, it was necessary to use the export/import utilities. New syntax for the `ALTER TABLE` statement allows a column to be marked as unusable, without freeing up space in the table, or to be dropped from the table with the data deleted.

This feature complies with Transitional SQL92.

For more information, see *Oracle8i Administrator's Guide* and *Oracle8i SQL Reference*.

Locally Managed Tablespaces

Prior to Oracle8i, management of free and used extents within a tablespace relied heavily on data dictionary tables. Now, Oracle introduces a new mechanism for managing space within a tablespace: locally managed tablespaces. All extent information is tracked in the tablespace itself, using bitmaps.

Bitmaps manage space allocation very efficiently, and require no dictionary access to allocate or update extents for the tablespace. The result is better space management, reduced fragmentation, and increased reliability.

For more information, see *Oracle8i Concepts* and *Oracle8i SQL Reference*.

Online Index Creation, Rebuild, and Defragmentation

Creating or rebuilding an index on a table improves query performance when completed, but can be a very time-consuming operation. For large tables, it can require several hours of downtime because it forces a lock on the table and prevents concurrent DML. This is called offline index build.

Oracle8i supports the online creation or rebuilding of an index that works for partitioned or nonpartitioned B*-tree indexes, including index-organized tables. Since the table is not locked, DML operations and queries can execute on the base table while the index is being built.

New SQL syntax specifies the online index creation or rebuild. Specifically, the `ONLINE` keyword may be specified on the `CREATE INDEX` or `ALTER INDEX` statement.

Alternatively, if you do not want to rebuild the index because of space or other considerations, then the new `COALESCE` keyword may be specified to defragment the index. In either case, performance is improved and space recovered.

For more information, see *Oracle8i Administrator's Guide* and *Oracle8i SQL Reference*.

Nonpartitioned Table Reorganization

A new `MOVE` clause used in an `ALTER TABLE` statement provides a means to easily reorganize a nonpartitioned table by allowing the user to move data into a new segment while preserving all views, privileges, etc. defined on the table. The operation is performed offline.

For more information, see *Oracle8i Administrator's Guide* and *Oracle8i SQL Reference*.

Online Read-Only Tablespaces

Oracle8i improves the performance of the operation that places a tablespace in read-only mode. A tablespace in Oracle8i can be placed in read-only mode when there are no outstanding transactions in that tablespace alone, unlike previous versions of Oracle where the operation completed only when there were no outstanding transactions in the entire database.

Now, an `ALTER TABLESPACE...READ ONLY` statement makes the tablespace read only as soon as all transactions on that tablespace have completed. This capability is especially important for the transportable tablespaces feature described in "[Transportable Tablespaces](#)" on page 3-23.

For more information, see *Oracle8i Concepts*.

Temporary Tables

A temporary table is a table with session-specific or transaction-specific data. It is empty when the session or transaction begins, and discarded at the end of the session or transaction. Its definition is visible to all sessions but the data is visible to, and can be queried by, only the session that inserts the data into the table. It is created in the user's temporary tablespace. Undo (rollback to savepoint) is supported, but not redo (crash recovery). Temporary tables can be useful for saving intermediate results that can be joined back into another table. They also might facilitate migration from other databases.

For more information, see *Oracle8i Concepts*.

National Language Support (NLS)

Oracle8i National Language Support includes significant new enhancements that make it easier to deploy applications requiring local language support, or that require country-specific formatting.

General Enhancements

These are some general enhancements to NLS.

- Programming interfaces (OCI) provide cartridge developers and application developers access to international information and services.
- Function-based indexes, as described in "[Function-Based Indexes](#)" on page 3-24, allow for the creation of indexes based on linguistically sensitive NLS_SORT orders.
- Fixed-width Unicode (UCS2) character support is provided in the following client interfaces: OCI, Pro*C/C++, ODBC.
- Expanded NLS data has been included. New territories have been added for Europe (Belgium, Ireland, Luxembourg), Asia Pacific (Australia, New Zealand, Singapore), Eastern Europe (Kazakhstan, Uzbekistan), and Africa (South Africa).
- Support for the Euro currency symbol is included.

For more information, see *Oracle8i National Language Support Guide*.

Changing Character Sets with ALTER DATABASE

Oracle8i makes it possible to change both the character set the database uses to store data, and the national character set used to store data in columns specifically defined as NCHAR, NCLOB, or NVARCHAR2, *after* the database has been created. The old character set must be a strict subset of the new one.

A user with SYSDBA system privilege, may issue the following statements:

- ALTER DATABASE CHARACTER SET
- ALTER DATABASE NATIONAL CHARACTER SET

For more information, see *Oracle8i SQL Reference*.

TRIM Function

This enhancement implements the ANSI standard `TRIM` function. It combines the functionality of the existing `LTRIM` and `RTRIM` functions, allowing the user to trim leading or trailing characters, or both, from a character string.

For more information, see *Oracle8i SQL Reference*.

External Routines

An *external routine*, previously referred to as an external procedure, is a routine written in another language and stored in a dynamic link library (DLL), or `libunit` in the case of a Java class method. The routine is registered with the base language, and can be called to perform special-purpose processing.

First introduced in Oracle8, external procedures allowed the writing of C functions as PL/SQL bodies. These C functions are callable from PL/SQL and SQL (via PL/SQL). In Oracle8i, a special-purpose interface, the *call specification*, is provided that lets external routines be called from other languages.

While this service is designed for intercommunication between SQL, PL/SQL, C, and Java, it is accessible from any base language that can call these languages. For example, a routine can be written in a language other than Java or C and still be usable by SQL or PL/SQL, provided that is callable by C. Therefore, a C++ routine, would use a C++ `extern "C"` statement in that routine to make it callable by C.

Below is a brief description of the new call specification, and a discussion of other changes targeting enhancement of external routine performance.

For information about external routines, see *Oracle8i Application Developer's Guide - Fundamentals*.

The Call Specification

Until now, an external routine has been published to Oracle via an `AS EXTERNAL` clause in a PL/SQL wrapper. This wrapper defines the mapping to, and allows the calling of, external C routines. Oracle8i introduces call specifications, which include the `AS EXTERNAL` wrapper as a subset of a new `AS LANGUAGE` clause. `AS LANGUAGE` call specifications allow the publishing of external C routines, as before, but also Java class methods.

Call specifications also allow publishing with the `AS EXTERNAL` clause, introduced in Oracle8. For new applications, however, the `AS LANGUAGE` clause should be used.

External Routine DLL Caching

A DLL caching mechanism has been developed to allow reuse of already loaded libraries for subsequent external routine invocation. This avoids costly repetitive operations, thus improving the performance, usability, and scalability of external routines. DLL caching applies to both server-side external routines, invoked through the extproc agent, and to client-side external routines.

External Routine Object Support

Certain object support, previously available to trusted external routines only, is now available to untrusted and distributed routines. Specifically, all external routines now support object types as argument types, and clients are allowed to issue object-related OCI callbacks.

Database Security

For a discussion of the *virtual private database* feature, whose components are fine-grained access control and application context, see "[Securing Data in the Oracle Database Server](#)" on page 3-68. This feature secures data in the database by providing security at the row-level, across all applications, by attaching a security policy directly to a table or view.

Partitioning Enhancements

Range partitioning was introduced in Oracle8, release 8.0. Oracle8i provides enhancements for range partitioning, but more significant is the expansion of Oracle's repertoire of partitioning methods by introducing two new methods: *hash* and *composite*.

For more information about partitioning and the new features discussed in this section, see *Oracle8i Concepts* and *Oracle8i SQL Reference*.

Range Partitioning Enhancements

Merging Partitions

A new statement, `ALTER TABLE...MERGE PARTITIONS`, may be used to merge the contents of two adjacent partitions of a table partitioned using the range method. This is the inverse of a `SPLIT PARTITION` operation.

Updatable Partition Keys

A new clause, `ENABLE ROW MOVEMENT`, of the `ALTER TABLE` or `CREATE TABLE` statement allows rows to be moved between partitions. When row movement is enabled, updates affecting partitioning keys in such a way that a row no longer belongs in its current partition will proceed, resulting in a row migrating to the appropriate partition. The default behavior (`DISABLE ROW MOVEMENT`) is that row movement is disabled and such updates are disallowed, resulting in an error being returned to the user.

New Partitioning Methods

The new partitioning methods introduced in Oracle8i are hash and composite. Along with range partitioning, these present a rich set of partitioning methods that allow the DBA to choose the partitioning method that offers the best performance depending on the database workload and the application's profile. As with range partitioning, all partitioning techniques are transparent to applications and standard DML statements run against the partitioned tables.

Hash partitioning provides a very simple way to break data into evenly sized containers to be spread across multiple I/O devices, or even multiple machines in a shared-nothing cluster. Query performance is improved by spreading I/O across multiple devices, and the performance of parallel DML may also be improved. Note, however, that hash partitioning is inappropriate for rolling change windows of historical data because the data for any period of time can be spread across multiple partitions.

Composite (range/hash) partitioning provides the manageability and availability benefits of range partitioning with the data distribution advantages of hash partitioning. The user specifies ranges of values for the primary partitions of the table or index, then specifies a number of hash subpartitions. Data skew is unlikely, because the user can always add or drop subpartitions within a partition to maintain even distribution of each container. Rolling change windows of historical data are easily maintained by adding or dropping primary partitions with no effect on subpartitions in other primary partitions.

Local indexes are supported with hash and composite partitioning. All existing SQL query and DML statements work with both new partitioning methods and new or extended DDL statements are provided for maintenance.

Additional Performance Gains

Enhanced partitioning includes many performance enhancements, some of which have already been discussed. Not yet discussed are the following two.

Enhanced Partition Elimination

Partition elimination is the skipping of unnecessary index and data partitions (or subpartitions) in a query. Because Oracle8i enhanced partitioning provides the database with greater knowledge of data placement, it can perform more advanced partition elimination. Particular focus is on queries using disjunctive `OR` or `IN` predicates.

Partition-Wise Join

In Oracle8, as the first step when two tables are being joined in parallel, each table is split into some number of separate pieces, scanned, then redistributed on the join column. In Oracle8i, the redistribution step can be skipped if the data is already partitioned on the join key. This reduces memory and temporary storage requirements and increases overall performance.

Partitioned Table LOB Support

Full support for partitioned tables with LOB columns is now provided. The LOB value can be stored inline in the row or out of line in a separate segment. All LOB types are supported: BLOB, CLOB, NLOB, BFILE.

Partitioning of Index-Organized Tables

Very large databases, historical databases, and OLTP databases with special availability requirements can benefit from the partitioning of index-organized tables (IOTs). Oracle8i delivers this support for partitioning of IOTs and the secondary indexes defined on IOTs. The bulk loading of partitioned IOTs with secondary indexes via SQL*Loader direct path is supported and provides the most efficient way of loading data.

System Management

This section is concerned with the new functionality and tools that greatly enhance your ability to control resource usage, enable better recoverability and availability in the database, and make your database a lot easier to manage.

Database Resource Management

Traditionally, resource management decisions have been left to the operating system, rather than the database server. This results in the following problems.

- There is no way to partition machine resources appropriately among database tasks of varying importance.
- When the number of servers is high, there is excessive overhead due to OS context switching between database servers.
- Scheduling is inefficient--Oracle database servers are descheduled while they hold latches.

The Database Resource Manager addresses these issues by giving the database more control over the management of resources. Oracle8i's Database Resource Manager feature provides the DBA with the ability to control and limit the total amount of processing resources available to a given user or set of users.

For example, suppose that a data warehouse is shared by the marketing department and the sales department. Using the Database Resource Manager, a warehouse administrator could specify that the marketing department receives at least 60 percent of the CPU resources of the machines, while the sales department receives 40 percent of the CPU resources. The warehouse administrator could further specify limits on the total number of active sessions, and the degree of parallelism of individual queries for each department.

Specifically, using the Database Resource Manager facility, the DBA can:

- Guarantee certain users a minimum amount of processing resources regardless of the load on the system and the number of users.
- Distribute available processing resources by allocating percentages of CPU time to different users and applications. In a data warehouse, a higher priority may be given to ROLAP applications than to batch jobs.
- Limit the degree of parallelism that a set of users can use.
- Configure an instance to use a particular method of allocating resources. A DBA can dynamically change the method, for example, from a daytime setup to a nighttime setup, without having to shut down and restart the instance.

In addition, a user can select the priority of a session from a given set of priorities that the DBA has assigned to that user.

DBAs use *resource consumer groups*, *resource plans*, *resource allocation methods*, and *resource plan directives* to implement Database Resource Manager. These functions are:

Resource consumer group	They allow the administrator to group user sessions together based on their requirements for processing resources.
Resource plan	It contains resource plan directives that specify the resources that are to be given to each resource consumer group.
Resource allocation method	It determines what method (or policy) to use when allocating for any particular resource. Resource allocation methods are used by both plans and consumer groups.
Resource plan directives	There is one of these for each consumer group in the plan. They allow the administrator to assign consumer groups to particular plans and partition resources among consumer groups by specifying parameters for each resource allocation method.

Resource plans, resource consumer groups, and resource plan directives are created using the PL/SQL package `DBMS_RESOURCE_MANAGER`. Users are assigned to consumer groups using the `DBMS_RESOURCE_MANAGER_PRIVS` package. The initialization parameter `RESOURCE_MANAGER_PLAN` specifies which top plan to use for a given instance. The Database Resource Manager loads this top plan as well as all its descendants (subplans, directives, and consumer groups).

For more information, see *Oracle8i Concepts*, *Oracle8i Administrator's Guide*, and *Oracle8i Supplied PL/SQL Packages Reference*.

Recoverability and Availability in the Database

The Oracle database server's industry leading availability and reliability are further extended in Oracle8i. Significant enhancements and new features expand our customers' ability to provide highly available data center solutions and meet service level objectives.

Unless otherwise noted, information about the new recovery and availability features discussed here can be found in the *Oracle8i Backup and Recovery Guide*.

Multiple Remote Archive Destinations

Oracle8i provides the means to allow the background archive process (ARCH), or its foreground equivalent, to archive online redo log files to multiple destinations. Up to five destinations can be specified with the initialization parameter: LOG_ARCHIVE_DEST_n. A destination can be a local disk-based file or it can be a user-specified standby database that is either local or remote to the primary database.

It has been observed that two archival destinations is not enough, while more than five is excessive and provides no additional benefit.

Multiple Archive Processes

Oracle8i allows the user to specify multiple archive processes to be invoked at instance startup. This is intended to alleviate bottlenecks that can occur when the LGWR (log writer) process is able to write to the redo log file faster than ARCH can read from the log file and write to one or more destination archive logs. It is no longer necessary for users to attempt manual foreground archival operations in order to recover, as multiple archive processes automatically handle the situation.

Automated Standby Databases

Automated standby databases provide the means to create and maintain multiple remote copies of a production database for disaster recovery scenarios. This is one of the key technologies available to protect against a disaster to the primary database system or data center. The backup or automated standby database can reside in the same or different data center and take over the processing from the primary production database providing near continuous database availability. Oracle has provided the standby database feature for a number of years, but the technology is significantly extended in Oracle8i.

Automatic Archival A standby database is initially created by copying or duplicating the production database. As archived redo logs are generated on the production database, they are applied to the standby database. This allows the standby database to remain synchronized with the production database. In releases prior to Oracle8i, the redo logs were manually transported or copied to the standby database and manually applied.

With the new automated standby database feature, the archived redo logs can be automatically transferred and applied. This eliminates the need for manual procedures to copy and transmit the redo logs and the need for the operator at the backup site to manually specify which logs to apply. The standby database is placed in constant recovery mode by the user, causing it to wait for archived log files from

the primary database. When a primary archived log file arrives at the standby site, it is verified and then applied to the standby database.

Read-only Databases Another new capability is the ability to use a standby database as a read-only database. All databases, not just standby databases, benefit from this new feature, but it is especially useful for standby databases. This option can be used to make a standby database available for queries and reporting, even while archive logs are being copied from the primary database site.

Fast-Start Fault Recovery

Transparent application failover (TAF) was introduced in Oracle8 and masks system failures from applications. It is discussed in "[Transparent Application Failover \(TAF\)](#)" on page 4-8. To allow the Oracle database server to quickly recover from system faults and minimize the impact on users, Oracle8i introduces fast-start fault recovery. It consists of the following components.

Fast-Start Checkpointing As a means of controlling the time required for instance and crash recovery, this feature allows the DBA to specify an upper limit on the number of I/O operations that Oracle needs to perform during instance recovery. A new dynamic parameter, `FAST_START_IO_TARGET` is introduced to control Fast-Start checkpointing.

Faster recovery is achieved at the expense of writing out additional (meaning: more than would have been written if this feature were disabled) buffers during normal processing. The fewer the number of dirty buffers in the buffer cache at the time of the failure, the faster the recovery times. This is because recovering data blocks by applying redo records between the most recent checkpoint and the end of the log constitutes the majority of the work and time during recovery. Statistics are collected and are available to the DBA through dynamic performance views to tune the system and make tradeoffs between runtime performance and database availability.

For additional information on this feature, see *Oracle8i Designing and Tuning for Performance*.

Fast-Start On-Demand Rollback This feature enables dead transactions to be recovered, on demand, one block at a time. This improves the availability of the database for users accessing data that is locked by large dead transactions. In earlier releases, users are blocked until the entire dead transaction is recovered. With this feature, the user immediately recovers only the block under consideration and proceeds, leaving the rest of the dead transaction to be recovered in the background.

Fast-Start Parallel Rollback This Oracle8i feature is primarily targeted to solve the problem of recovering parallel transactions. Parallel transactions in release 8.0 exploited parallelism in the roll forward stage using parallel block recovery. But on a failure, the failed transactions were rolled back serially by the background process. This impacted the availability of the part of the database involved in the failed transactions.

Fast-start parallel recovery allows transactions to be recovered in parallel, resulting in significantly increased recovery throughput. Inter-transaction recovery is performed like parallel block recovery, where each slave rolls back a separate transaction. But if one of those transactions is extremely long, then intra-transaction recovery is invoked, where that transaction is divided up among the slaves.

Fast-start parallel recovery is started when SMON discovers that the amount of recovery work that is required is above a certain threshold. The threshold is the amount of work at which the time it takes to recover in parallel becomes less than the time it takes to recover serially.

Database SUSPEND/RESUME

This new enhancement provides a mechanism by which all I/O in the database can be stopped, enabling copies of the database to be made without I/O interference. Two new statements are introduced. The `ALTER SYSTEM SUSPEND` statement stops new lock and I/O activity from being initiated. The creation of backups or archived logs is not affected. The database remains suspended until an `ALTER SYSTEM RESUME` statement is issued.

Control File Character Set

This enhancement enables the Recovery Manager to correctly interpret tablespace names for recovery when the character set is other than the default. This is done by storing a character set ID in the control file. When the database is first opened, this character set ID is propagated from the data dictionary into the control file database information record. The information is then available from the control file, before database open, for any subsequent recovery attempts.

If the control file is lost, recovery is still possible. The DBA may specify the character set as an argument on the `CREATE CONTROLFILE` statement.

Recovery Manager (RMAN)

Changes to the Recovery Manager are presented below. Additionally, the *Oracle8i Backup and Recovery Guide* has been extensively rewritten. Most notably, it now

contains an exhaustive reference for RMAN statement syntax and a chapter on recovery catalog views.

Recovery Catalog No Longer Required

The use of a recovery catalog is now optional. If you choose not to use a recovery catalog, then you can still use RMAN very effectively. RMAN obtains the information it needs from the control file of the target database. When using a recovery catalog, RMAN can perform a wider variety of automated backup and recovery functions. For this reason, Oracle recommends using a recovery catalog with RMAN whenever possible.

A related enhancement is discussed above in "[Control File Character Set](#)".

Media Management API, Version 2.0

RMAN includes changes to support version 2.0 of the Oracle Media Management API, commonly referred to as System Backup to Tape API, or SBT. Support for the version 1.1 Media Management API is maintained.

In order to utilize different types of tertiary storage for Oracle database backups, RMAN requires that third party media management software be installed. This software interfaces to Oracle8i and RMAN using SBT. It is the responsibility of this vendor supplied software to load, label, and unload sequential media such as tape drives for the purpose of backing up and recovering data.

Proxy Copy Through an enhancement called *proxy copy*, the media management vendor software is able to take over the entire data movement involved in a backup or restore. RMAN merely provides a list of files requiring backup or recovery to the media manager, which in turn makes all decisions regarding how and when to move the data. The `PROXY` option of the `BACKUP` statement invokes this facility. The facility must be specifically supported by the media management software--not all media management software provides this support.

Media Pool Selection Some media management software allows backup media to be arranged into storage pools based on media type, retention period, or other criteria. The new `pool` operand of the `backup` statement provides tight integration between such products and RMAN.

Vendor Identification and Improved Error Messages When a channel is allocated, RMAN displays in its log a text message identifying the media management product that will be used to take backups on that channel. When the media manager encounters

an error, it returns a text error message explaining the error, which is displayed in the RMAN log.

Send Statement The new **send** statement allows statements to be sent directly from an RMAN session to the media management software. Typically, this statement is used to control runtime options of the media manager that can not otherwise be controlled using RMAN.

Crosscheck Catalog

RMAN can detect when backups expire from the media management software catalog or on disk, and it can update its own catalog accordingly. Two new statements, **crosscheck backup** and **delete expired backup**, allow for the synchronization of the recovery catalog with the media manager's catalog.

The new **change...crosscheck** statement can perform crosschecks on both backup sets and image copies.

List and Report Statements

The output of the **list backup** statement is significantly improved. It now prints the list of backups belonging to a backup set in a separate section of the report from the list of data files or archived logs included in the backup set.

A new statement, **report need backup redundancy**, has been implemented. When fewer than a user-specified number of backups of a data file exist, this report alerts the user that a new backup is required.

Recovery Catalog Maintenance Statements

The following new statements make it easier to create, upgrade, and drop the recovery catalog.

CREATE CATALOG This statement creates the recovery catalog. It replaces `catrman.sql` and associated scripts in the admin directory, which had to be run manually to create the recovery catalog schema.

UPGRADE CATALOG When RMAN detects that the version of the recovery catalog is too old, it issues an error message indicating that the recovery catalog requires upgrading. Previously it was necessary to run a SQL script to perform the upgrade. Now the upgrade is done with the **upgrade catalog** statement. This statement upgrades the recovery catalog from any prior version. No data is lost; the backups which were taken with the recovery catalog remain intact.

DROP CATALOG This statement removes the recovery catalog schema. The backups themselves are not removed, so if there is a duplicated or replicated catalog available, or the information is available in the control file, the backups can still be used for restore.

Database Maintenance Statements

These new statements allow for more types of recovery with a single RMAN session. It is no longer necessary to exit RMAN to shutdown, startup, mount, or open the database.

STARTUP/SHUTDOWN These statements have the same syntax as the corresponding SQL*Plus statements.

ALTER DATABASE The following two forms of the ALTER DATABASE statement are supported.

- ALTER DATABASE...MOUNT
- ALTER DATABASE...OPEN

DUPLICATE Statement

The **duplicate** statement allows creation of a replicated database using the backups of another database. Typically, a database is duplicated in order to create a test database that is separate from the production database.

Node Affinity Detection

When backing up on multiple nodes of a parallel server, it is possible that some disks have 'affinity' to certain nodes in the cluster such that access to those disks is faster from those nodes than from other nodes in the cluster. RMAN recognizes node affinity, if it exists, and attempts to schedule data file backups on channels allocated at nodes that have affinity to those files. To use this feature, channels must be allocated at more than one node in a parallel server cluster.

Backups

The following are RMAN changes affecting backups.

Duplexed Backup Sets Recovery Manager can now create up to four concurrent copies of each backup piece. This feature is particularly useful for archivelog backups. If a data file backup is bad, then an earlier backup of the same data file can usually still be used for recovery. However, if an archivelog backup is bad, then no

recovery beyond the point in time when that log was created is possible. Because archivelogs are critical to recovery, customers often require that archivelogs be backed up twice before being deleted from disk.

Naming RMAN no longer requires that backup piece names be explicitly specified using the **format** parameter. By default, RMAN chooses a unique name for each backup piece.

Is No Longer Overwritten A backup piece is no longer overwritten if an attempt is made to create a backup piece with the same name as an existing one. RMAN now issues an error.

Performing TSPITR Without A Recovery Catalog

It is now possible to perform TSPITR (Tablespace Point in Time Recovery) without a recovery catalog, however some restrictions apply.

Backup Performance Views

Two new views are available to monitor the progress and performance of Recovery Manager backups. These are `V$BACKUP_SYNC_IO` and `V$BACKUP_ASYNC_IO`. These views can be used to monitor the progress of individual files in a backup or restore, and to identify performance bottlenecks. A complete description of how to use these views to tune backups is in *Oracle8i Designing and Tuning for Performance*.

SQL*Loader Enhancements

For Oracle8i provides the following the SQL*Loader enhancements:

- SQL*Loader now includes support for the loading of objects, collections, and LOBs.
- There is no longer a 64K physical record size limit.
- A new keyword, `FILLER`, can be used to specify a filler field: *a data file mapped field which corresponds to no database column*. The filler field is assigned values from the data field to which it is mapped.
- New datatypes `VARCHARC`, `VARRAW`, `LONG VARRAW`, and `VARRAWC` are length-value pair datatypes similar to `VARCHAR`.
- The user can now specify a custom record separator in the OS-file-processing string.
- `DEFAULTIF` and `NULLIF` predicates can now be ANDed together.

- Field delimiters can now be one *or more* characters in length.

For more information, see *Oracle8i Utilities*.

Export and Import Utilities

Many of the changes for the Import and Export utilities in Oracle8i are in support of other features. However, there are a few enhancements in these utilities which stand on their own.

- The ability to specify a query for the select statements that export uses to unload the tables.
- The ability to specify multiple dump files for an export statement. (This allows users to circumvent the previous 2Gb limit for export dump files.)
- The ability to export tables containing `LOBs` and objects, even if direct path is specified on the statement line.
- The ability to export and import precalculated optimizer statistics instead of recomputing the statistics at import time. (This capability is only applicable to certain exports and tables.)
- The ability to selectively inhibit the validation of type object identifiers by import.
- Support for larger dump files and dump tapes.

For more information on any of the above, see *Oracle8i Utilities*.

Operational Simplicity

Oracle is committed to making the Oracle database server simpler to use and administer. Reducing the number and duration of administrative tasks reduces the total cost of ownership for a system. Fewer administrators are required for day-to-day operations, and they can spend their time on more challenging issues such as performance tuning and database design. For very low end systems in particular, Oracle is working toward making them as self-administering and self-tuning as possible.

The following Oracle8i components support Oracle's goal toward achieving operational simplicity.

Oracle Software Packager and Oracle Universal Installer

The Oracle Software Packager and Oracle Universal Installer work together to provide a robust method for packaging and installing Oracle8i software, whether for simple installations or those with complex interdependencies.

Oracle Software Packager The new Oracle Software Packager is a Java-based, object-oriented tool used for packaging products into components that can later be installed by the Oracle Universal Installer. Using wizards, it provides a visual development environment for generating installation packages in a format that can be interpreted by Oracle Universal Installer.

A component is the basic object to be installed by Oracle Universal Installer. The component wizard guides the install developer through the specification of the installation properties of the component, like component name, version number, operating systems and languages, the files or groups of files to be installed as part of the component, dependencies, variables, and the sequence of *dialogs* and *actions* that make up the installation plan. These properties are stored in a component definition file.

Dialogs represent user interface screens that are displayed to the installing user at run time and are used to assign values to variables. Action blocks are the state-change portions of the installation, and perform such tasks as copying files, setting environment variables, creating users, etc. The Oracle Software Packager is extensible and customizable. The install developer can take advantage of the packager's Java class libraries implementing predefined actions or dialogs, or request that the packager import custom Java classes.

Once the developer has defined the installation characteristics, Oracle Software Packager can verify the component definition and stage it. The verify wizard checks the component's integrity, while the staging wizard moves the component from the development area to a distribution medium, such as CD-ROM, Internet Server, or NFS system.

The Oracle Software Packager does not require access to a component to allow it to be defined as a dependent of the component being packaged. From simple installation definitions to very complex bundled installations, Oracle Software Packager enables packaging of Oracle software, as well as easy integration of ISV and value-added reseller (VAR) software into a common installation.

For more information, see *Oracle Software Packager User's Guide*.

Oracle Universal Installer Oracle8i is installed using the new Java-based Oracle Universal Installer, presenting the same look and feel across all platforms. Oracle

Universal Installer interprets the staging area prepared by Oracle Software Packager and performs the appropriate installation actions. Its features include:

- It has the ability to read predefined answers to installation dialogs from response files defined by the administrator using the response file editor. This allows Oracle Universal Installer to run in silent mode.
- Installs can be from a CD, or they can be Web-based.
- Multiple Oracle_Homes are allowed, enabling multiple versions (different release levels) of the same component to be installed.
- Dependencies are automatically resolved.
- Postinstallation tools, such as the Database Configuration Assistant, are automatically launched at the end of installation, if specified by the software packager.
- Logging

The user is guided through the installation process by wizards and online help.

Configuration Improvements

The Oracle Database Configuration and Data Migration Assistants have been enhanced to further simplify the creation, deletion, modification, and migration of Oracle databases.

For information on either of these products, refer to your operating system specific Oracle installation documentation. Some features may not be available on all operating systems.

Oracle Database Configuration Assistant The Oracle Database Configuration Assistant enables the creation, modification, or deletion of an Oracle database. In Oracle8i, it can be invoked as a standalone Java application from the Oracle Universal Installer, or as an applet from the Java-based Oracle Enterprise Manager.

A wizard interviews the user for information relevant to the use and environment of the database. During this process, new Oracle users can create a pre-tuned starter database by choosing to copy files (a seed database) from the installation CD-ROM to their machine. Alternatively, more experienced Oracle users can customize the type of database needed for their environment, and can choose to create the database immediately or generate a SQL script to create the database at a later time. DBCA also supports silent mode creation, which is useful for remote installations or integrated ISV installations. The creation of remote databases is supported using the Oracle Enterprise Manager's intelligent agent.

The Database Configuration Assistant provides the ability to create an Oracle OFA (Optimal Flexible Architecture) compliant instance. The OFA-like file layout provides a standardized file layout and eases management and maintenance.

If the user selects to modify a database, installed cartridges that have not previously been configured for use with the database can be configured. Also, multi-threaded server support for the database can be enabled or disabled.

Oracle 8.1.5 has improved error detection mechanisms, enhanced MTS support, and better cleanup on failure and cancellation of database creation. In addition, when a database is deleted, the entire Oracle Flexible Architecture (OFA) structure is removed for that database.

Oracle Data Migration Assistant Oracle Data Migration Assistant is a GUI tool that seamlessly transforms the data of an existing Oracle database to the current Oracle database release. It changes data file headers but leaves actual data unchanged. It does not copy data.

The Oracle Data Migration Assistant can migrate an Oracle7 database to Oracle8i. See your operating system specific Oracle documentation for information about the earliest release that the Oracle Data Migration Assistant can migrate on your operating system.

Oracle Enterprise Manager (EM)

Version 2.1 of the Oracle Enterprise Manager delivers greater scalability and ease of use through the following enhancements.

- It uses a new three-tier model, allowing flexible deployment options.
- It is Java-based, allowing it to be accessed from any operating system, not just Windows NT.
- A shared repository allows administrator access from multiple accounts and consoles.

The first tier, of Oracle Enterprise Manager's three-tier model, consists of a Java-based console and integrated applications which can be installed or run from a Web browser. Administrators can access the console from virtually anywhere, and the 100% Java framework provides heterogeneous platform support.

The middle-tier component is the Oracle Management Server (OMS). Its main function is to provide centralized intelligence and distributed control between clients and managed nodes, thereby processing and administering all system management tasks. As the number of system management tasks increases, the architecture scales through the addition of more OMSs. Failover and load balancing

are also automated within the OMSs, providing much greater reliability in notification processing.

The third tier is comprised of managed services or targets, such as databases, nodes, applications or application servers. An intelligent agent functions as the executor of jobs and events sent by the OMS. The agent works independently from the Console, the OMS, or network connections. Once an Agent is assigned a task from the OMS, it performs the work at the time scheduled, and report status back to the OMS. An example of one such task might be performing weekly database backups.

Oracle Enterprise Manager, version 2.1, is a multi-user system. Multiple administrators with separate accounts and permissions can log into their own consoles. A shared repository allows them access to the same information from their separate consoles.

Oracle Parallel Server

Numerous enhancements to Oracle Parallel Server deliver major advances in performance, scalability, and manageability. Additionally, by extending commonly used single-instance diagnostic tools and functionality to be Oracle Parallel Server-aware, DBAs, development, and support personnel are better able to diagnose, tune, and monitor systems. Summaries of these enhancements are presented below.

For more information, see the *Oracle8i Parallel Server Documentation Set: Oracle8i Parallel Server Concepts; Oracle8i Parallel Server Setup and Configuration Guide; Oracle8i Parallel Server Administration, Deployment, and Performance*.

New Architecture

In Oracle8i, new architecture dramatically decreases the need for workload partitioning and lets applications benefit from an Oracle Parallel Server implementation. Changes include the following:

- A new "diskless" ping architecture, called cache fusion, that provides copies of blocks directly from the holding instance's memory cache to the requesting instance's memory cache. This functionality greatly improves inter-instance communication. Cache fusion is particularly useful for databases where updates and queries on the same data tend to occur simultaneously and where, for whatever reason, the data and users have not been isolated to specific nodes so that all activity can take place on a single instance. With cache fusion, there is less need to concentrate on data or user partitioning by instance.

- Deeper integration of the distributed lock manager (DLM) into the database engine, increasing the efficiency of parallel cache management.
- A new, inter-process communication protocol that is highly efficient on modern interconnects.

Easier System Administration

Oracle 8i improves the diagnosis of system performance problems and the administration of the Oracle Parallel Server environment with enhanced tools.

Diagnostic Enhancements

Oradebug is a utility used by consulting and support personnel to diagnose and troubleshoot poorly behaving systems at run time. Previously, it was best suited for single-instance situations, but now that functionality is extended for the Oracle Parallel Server.

Oracle Parallel Server Management (OPSM)

Enhancements to OPSM provide for greater ease in administration of the Oracle Parallel Server. A single generic interface is presented for administering parallel servers on any platform, helping administrators to easily manage the complexities of the Oracle Parallel Server product. Here are some of the features.

- A new Enterprise Manager applet, Parallel Server Manager, has been introduced for managing parallel servers.
- Installation and creation of Oracle Parallel Server databases is handled by a separate tool called Oracle Parallel Server Assistant.
- Group Membership Services (GMS) has been moved to the vendor-specific Cluster Managers (CM) and the Oracle database engine. There is no longer a discrete GMS module visible to the user that must be explicitly started before starting instances.

For more information about enhancements to OPSM, see *Oracle Parallel Server Management Users Guide* and *Oracle8i Parallel Server Setup and Configuration Guide*.

Parallel Server Installation and Database Configuration

The Oracle Universal Installer (see "[Oracle Universal Installer](#)" on page 3-58) and Oracle Database Configuration Assistant (see "[Oracle Database Configuration Assistant](#)" on page 3-59) are both cluster aware.

In Oracle8i, only a single installer session is required to install an Oracle Parallel Server database. The installer collects node information from the user, distributes the required Oracle products to the specified nodes, and then invokes the OPS Assistant to set up the instances and create the database. When OPS Assistant is finished, the parallel server is up and available on all nodes, and the parallel server configuration information is saved so that OPSM can use it to manage the new parallel server.

For more information, see *Oracle8i Parallel Server Setup and Configuration Guide*.

Instance Affinity for Jobs

Instance affinity for jobs is the association of jobs to an instance. Through the new DBMS_JOB package, the user can indicate whether a particular instance, or any instance, can execute a user-submitted job in the Oracle Parallel Server environment.

This Oracle8i feature can be used to improve load balancing and limit block pinging. For instance, using Oracle Parallel Server and replication at the same time results in block pinging problems on the deferred transactions queue if all instances in a clustered environment decide to propagate transactions from the deferred transaction queue. By limiting the activity against tables to only one instance within the parallel server cluster, pinging can be minimized.

For more information, see *Oracle8i Supplied PL/SQL Packages Reference* and *Oracle8i Administrator's Guide*.

Recoverability and Availability Improvements

There are recoverability and availability improvements which enhance the reliability of the Oracle Parallel Server. These are covered in more detail elsewhere in this chapter, but are mentioned briefly here.

- Parallel transaction recovery is performed on each node. See "[Fast-Start Parallel Rollback](#)" on page 3-52.
- Node affinity for Recovery Manager increases backup and restore throughput. See "[Node Affinity Detection](#)" on page 3-55.
- Establishing bounds on recovery time reduces false pinging. See "[Fast-Start Checkpointing](#)" on page 3-51.

Networking Improvements

The following new Net8 enhancements, described elsewhere, improve ease of use and scalability for the parallel server.

- A database can now have more than one service associated with it. See "[Service Naming](#)" on page 3-75.
- Connection load balancing balances the number of active connections among various instances and dispatchers for the same service. See "[Connection Load Balancing](#)" on page 3-77.

Distributed Systems

Oracle8i extends the functionality of advanced replication, focusing on mass-deployment applications. Heterogeneous services, introduced in release 8.0 and which implemented an extensibility framework for accessing non-Oracle systems, has also been further enhanced.

Advanced Replication Enhancements

The Oracle8i features and enhancements described below comprise the overall effort to optimize replication performance and make snapshot environment distribution and security more effective. All are included in advanced replication.

For more information, see *Oracle8i Replication* and *Oracle8i Replication Management API Reference*.

Performance Improvements

Significant performance gains are realized by the internalization of PL/SQL replication packages and by optimizations to snapshot refresh.

Internal Apply Packages Continuing the trend started with release 8.0, more replication code has been moved into the database engine. The PL/SQL generated packages used to apply replicated transactions at a remote site have been internalized. This allows replicated transactions to be more efficiently applied at remote sites and because packages are not generated, a site can be more quickly instantiated. Internal packages are also more secure because they are tamper proof.

Faster Snapshot Refresh Snapshot refresh has been optimized to support large refresh groups. There is improved support for subquery snapshots, and for null refresh (no changes to the master tables since the last refresh). A single refresh group can now

contain 400 snapshots, and the number of roundtrips required to refresh snapshots in a refresh group has been reduced. (This feature was first added in release 8.0.5.)

Improved Mass Deployment Support

In Oracle8i, Oracle adds new functionality to better support mass deployment and front office applications.

Parameterized Snapshot Deployment Templates These facilitate the mass deployment of information to support such applications as field service and sales force automation. These templates represent a grouping together of snapshots and other database objects to be instantiated at a node. They allow a DBA to centrally package a snapshot environment for easy, custom, and secure distribution to one or multiple sites. The goal is to create the environment once, then deploy the snapshot deployment template as often as necessary. Template parameters allow data subsetting at a remote site without redefining the template, and a template may be defined as public or private. Public templates may be instantiated at any site, whereas private ones can be instantiated only at predefined, authorized sites. An Oracle Replication Manager deployment wizard guides the DBA through the selection of schema objects to add to the template, the selection of parameters, and defining authorizations.

Column Level Snapshot Subsetting Updatable snapshots can now be subsetted horizontally (selected rows) or vertically (selected columns). Previously, only horizontal subsetting was allowed. Vertical partitioning allows the deployment of the minimum amount of data needed by a remote site, thus reducing connection time. It also protects snapshot sites from changes to their associated masters. A column can be added to a master site without impacting the snapshot site, or a column can be deleted and not impact the snapshot site, if the snapshot site does not currently reference that column.

Offline Instantiation Snapshot deployment templates can be instantiated online or offline. Online instantiation allows a snapshot site to instantiate the template while connected to the target master site. The advantage here is that the data is current. However, this is at the cost of requiring a live connection, possibly of long duration and having the potential of generating extreme network traffic that could degrade other network services.

Offline instantiation allows the DBA to package the deployment group templates and required data onto some type of storage media (tape, CD-ROM, etc.) for distribution to a snapshot site. Instead of connecting to the master site, instantiation can be done by pulling the template and data from the storage media. Users can fast

refresh immediately after completing an offline instantiation; a full refresh is not required. Offline instantiation is an ideal solution for mass deployment situations where many disconnected laptops are instantiating the target template.

Improved Security

While the scripts used to instantiate a snapshot site are generated at the master site and can control access to data, it is still necessary to connect to a receiver and proxy snapshot administrator to propagate replicated transactions and to refresh snapshots. Oracle8i enhancements to the replication security model eliminate certain security deficiencies regarding the granting of privileges to untrusted sites.

- Object privileges, as required by a receiver of remote procedure calls (RPCs) at a master site, are now automatically managed. Only required object privileges are granted to untrusted sites.
- Proxy snapshot administrators now have a way of accessing objects in object groups without being granted excessive privileges.

Replication Manager

A number of improvements have been made to the Oracle Replication Manager; some have already been mentioned. A major and noticeable enhancement is that it has been rewritten to conform to the new Oracle Enterprise Manager (EM) Java interface. Oracle Replication Manager can now be run from anywhere in the network, and it is not constrained to a Windows-based machine. There is also a replication class of events in EM which, for example, can be used to monitor errors or delinquent snapshot refreshes.

Improved Oracle Lite Integration

The integration between the Oracle8i Enterprise Edition and Oracle Lite has been improved to provide better performance and increased functionality. All of the replication changes previously described are supported, and Oracle Lite users with laptops at remote sites especially benefit from reduced connection time.

Heterogeneous Services

Heterogeneous services (HS) implements an extensibility framework for accessing non-Oracle systems. Introduced in Oracle8, HS integrates the core of Oracle's gateway technology directly into the Oracle database server by extending the Oracle SQL engine to optimize and rewrite SQL for non-Oracle datastores. This closer integration has improved performance. For instance, the Transparent

Gateway for DB2 has provided a substantial performance improvement, particularly where large numbers of rows are retrieved.

Oracle8i enhances HS with the following changes. For more information, see *Oracle8i Distributed Database Systems*.

Improved Management Interface: Fixed Views

A set of fixed views is now available to provide heterogeneous services status and to monitor information. The following views are added:

V\$HS_AGENT	Identifies agent processes
GV\$HS_AGENT	Identifies agent process used by all instances of an Oracle Parallel Server configuration
V\$HS_SESSION	Identifies HS connections
GV\$HS_SESSION	Identifies HS connections used by all instances of an Oracle parallel server

The views are present at all times in the Oracle database server and can be accessed directly through SQL queries. An Oracle Enterprise Manager GUI applet also displays the information.

A full description of each of these views is available in *Oracle8i Reference*.

Miscellaneous Improvements

Agent self-registration was introduced in a previous release. It reduces or eliminates the need for DBA intervention in configuring heterogeneous services. In Oracle8i, code has been rewritten to make the self-registration process more efficient.

New in Oracle8i is an agent-specific shared library (DLL) for HS object files, other than drivers, that is substituted when linking agent executables. While the benefits from this change are platform-specific, it can improve scalability by using a single agent library for all types of agents (extproc, hsalloct, hssqlps, hsdepca, and hsots). Also, memory requirements may be reduced because agent executables become quite small.

Remote Join Enhancements

In distributed database systems where it may be necessary to join tables across systems, poor data access techniques can cause user data flow between database

servers to be a major factor in overall performance. In order to determine the best access technique for an optimum execution plan, the optimizer must first be able to determine a sufficient number of alternate paths. The remote join enhancements in Oracle8i present more options, allowing for better execution plans to be generated with a corresponding performance increase.

Networking, Security, and Oracle Advanced Security

Oracle8i security and networking enhancements support standards, provide stronger security, and lower the cost of ownership. The support of standards provides for more choice, and provides for greater ease of system integration between Oracle and non-Oracle components. When allowing remote access and extending networks to the Internet, you now have a variety of standards-based and proven methods for data encryption and authenticating users, databases, and Web servers.

The *virtual private database* feature secures data in the database by providing security at the row-level, making it easier to handle sophisticated security requirements. It also moves the definition of the security policy out of the application and into the database, thereby eliminating the potential to circumvent the security enforcement.

Centralized user management, through integrated security and directory services, lowers your cost of ownership. Fewer accounts means less administrative overhead, and single sign-on increases productivity by reducing the number of logins. It can also reduce lockouts resulting from a failure to remember all of the necessary passwords.

Securing Data in the Oracle Database Server

Oracle 8i introduces the virtual private database feature. With so many means of accessing the database available, both locally and remotely, and most notably through the Internet, server-enforced and granular access control becomes crucial. Security implemented in application code is insufficient, because to bypass the security one only needs to access the data from outside the application (through SQL*Plus, for example). The virtual private database feature addresses this problem by attaching the security policy directly to a table or view.

An important benefit of the virtual private database feature is that security is built only once, and not into every application that accesses the data. Then, no matter how a user gets to the data, the same security policy is enforced.

The virtual private database feature has two components: fine-grained access control and application context. Here is how they work together to implement a safe and consistent security policy across applications.

Fine-Grained Access Control

Fine-grained access control works using a dynamic predicate, or `WHERE` statement. This dynamic predicate allows the security rules to be acquired at SQL statement parse time when the base table or view is referenced in a select or DML statement.

Implementing fine-grained access control consists of creating a function to implement a security policy and associating the function with a table or view. The Oracle supplied `DBMS_RLS` package is used to associate security policy functions with the table or view. When a query or subquery which affects that table or view is executed, the server dynamically rewrites the query appending a `WHERE` condition generated by the function implementing the security policy.

For example, fine-grained access control can be used to enforce a security policy "customers can see only their own orders". An application might submit a SQL statement to select "all" orders. The security policy would append an additional predicate on the `WHERE` clause restricting the statement to retrieve only those orders submitted by that customer. This provides row-level, server enforced, access control, that needs to be built only once to protect your data no matter how the data is accessed.

For more information, see *Oracle8i Application Developer's Guide - Fundamentals* and *Oracle8i Supplied PL/SQL Packages Reference*.

Application Context

An application context functions as a secure data cache for storing information used to make access control decisions. It provides more flexibility and space than using `CLIENT_INFO`, and represents a new approach.

Application contexts are extensible. Each application can have its own context, with its own attributes. For example, an order entry context could have customer number and sales region attributes, while a human resources application might have employee number, position, and country code attributes. These attributes are used by fine-grain access control in enforcing a security policy.

An application context can be set and reset only by a designated package. The `CREATE CONTEXT` statement specifies the namespace of the context and the package to which it is bound. A secure and consistent API is provided to manage and access the application context.

For more information, see *Oracle8i Application Developer's Guide - Fundamentals* and *Oracle8i SQL Reference*.

Secure Network Computing Through the Oracle Advanced Security Option

With the continuing growth of distributed systems involving numerous databases and applications comes an ever increasing challenge of user authentication and user management. To meet this challenge of *enterprise user management*, Oracle8i provides methods of encryption, authentication, and authorization, along with integrated security and directory services, packaging them as Oracle Advanced Security.

The Advanced Networking Option (ANO) has been renamed to Oracle Advanced Security to better describe its focus and new functionality. Oracle Advanced Security runs on top of Net8 to deliver security solutions to the Oracle Network and beyond through the integration of industry standards for encryption, authentication, and remote access. Oracle8i expands the choice of security solutions by offering SSL (Secure Socket Layer) and RADIUS (Remote Authentication Dial-In User Service) protocol adapters.

Oracle8i also enhances enterprise user management by offering integrated security and directory services through Oracle Internet Directory (OID). User information is stored and managed in a directory and is accessed by an LDAP (Lightweight Directory Access Protocol) version 3 compliant server.

Note: The Oracle Security Server (OSS) product no longer exists in Oracle8i, but its functionality has been integrated into the standards-based products: Oracle Advanced Security and Oracle Internet Directory. Users of OSS must migrate their applications to use these Oracle8i components.

For more information, see *Oracle Advanced Security Administrator's Guide*.

SSL

Oracle Advanced Security includes implementation of a SSL protocol adapter. SSL is an industry standard protocol for securing network connections and is widely used over the Internet. It secures Net8 connections and other protocols as well, including IIOP connections used with thin clients and Enterprise JavaBeans.

SSL uses digital certificates and a public key infrastructure (PKI) to provide the major pieces of security:

- Authentication of people and machines
- Encryption techniques for privatizing
- Checksums for protecting against data modification or snooping

The importance of verifying the identity not only of users, but also of machines, becomes crucial when an organization opens its doors to the Internet.

Most organizations have chosen to protect their databases by keeping them inside of the network and placing a Web server outside of the network. Used in combination with a firewall or other security measures, the data can remain safe. The data server's ability to authenticate the Web server, therefore, is extremely important. You must verify that the Web server is what it claims to be, so that you can trust it to pass data--encrypted or unencrypted--to the outside world. SSL provides the critical authentication piece, along with the data privacy piece, of the security solution.

Oracle's implementation of SSL uses industry-standard X509 version 3 digital certificates for authentication. Briefly, the authentication works as follows:

- A client is issued an X.509v3 certificate and a private key by a *certificate authority* (CA), and stores the certificate in an *Oracle wallet*.
- When the client initiates a Net8 connection to the server, SSL uses the private key and certificate to perform a handshake (transparent to the user) between the two processes.
- If it is successful, the client is granted access.

SSL supports single sign-on in that the user only requires one password: that which opens the wallet. The certificate and private key are used to authenticate the user to multiple services. The complete package for this authorization requires components and functionality as described in "[Integrated Security and Directory Services](#)" below.

SSL does not provide authorizations (Oracle roles), that must be provided by the application. SSL provides only public key-based authentication; this is distinct from other authentication methods provided by Oracle Advanced Security, which include tokens, smart cards, and biometric devices.

Note: Use of SSL for client authentication and single sign-on was a beta feature in 8.1.5. It is production in 8.1.6.

RADIUS

The RADIUS (Remote Authentication Dial-In User Service) protocol is an industry standard for remote authentication and controlled access to networks. It is widely accepted because of its flexibility, its ability to handle many devices, and its ability to provide user authentications, authorization, and accounting between a network client and an authentication server.

RADIUS support in Oracle Advanced Security offers two major benefits.

- It readily integrates into existing systems by making the Oracle8i database server a RADIUS client and thereby capitalizing on the infrastructure and investment that organizations have already made.
- It extends the solutions for authenticating users to Oracle by enabling support for new authentication technologies such as token cards, smartcards, and challenge-response mechanisms.

Your choice of authentication mechanisms is extensively expanded, with little development effort required.

Integrated Security and Directory Services

Note: The features discussed here for integrated security and directory services are available only as a beta release. They are not an integral part of this release (8.1.5), but are planned to be available in the next production release.

Enterprise user management and security presents many challenges. Users often have too many passwords, consequently they write them down or choose the same password for all accounts. They have multiple accounts that can be cumbersome and time-consuming for an organization to manage. Another challenge is that common application information is often fragmented across the enterprise, leading to data that is redundant, inconsistent, and expensive to manage.

Oracle has recognized these challenges and has continued to evolve solutions. The Advanced Networking Option was introduced in Oracle7, providing data encryption and supporting multiple authentication mechanisms, including biometrics, DCE (Distributed Computing Environment), and Kerberos. Oracle8 introduced Oracle Security Server (OSS), supporting single sign-on using X.509 version 1 certificates and proprietary authentication protocol and components. Oracle8i enhances the security architecture of Oracle8 by providing an open, standards-based approach for enterprise user management.

Oracle8i addresses the enterprise's need for strong security and centralized management of application information by offering integrated security and directory services; specifically, by storing and managing user information in a directory. Multiple Oracle applications can rely on a common centralized definition of a user to determine which applications, services, and data servers a user may access, and with what privileges. The benefits of this approach include:

- Single sign-on to multiple services throughout the enterprise
- Single enterprise user, instead of multiple accounts per user
- Reduced total cost of ownership through single station user administration
- Well integrated, standards-based public key infrastructure (PKI)

Oracle8i's integrated security and directory services incorporates multiple components, described below.

Oracle Wallet Manager Oracle Wallet Manager is a tool that is used to manage an Oracle wallet. An Oracle wallet contains a user's credentials used for authenticating the user to multiple services, such as data servers and application servers. Specific wallet contents include:

- A user's X.509v3 certificate. This certificate is analogous to a driver's license containing a user's identifying information, and it allows a user to be uniquely identified within an organization, thus enabling strong authentication. Oracle requires this certificate be issued by an X509v3-compliant certificate authority.
- A key pair, issued by the wallet, for a user. This includes a private key, which never leaves the wallet, and a public key. The wallet also issues a PKCS#10 (Public Key Certificate Standard #10) certificate signing request. The PKCS#10 request must be fulfilled by a certificate authority (CA), who validates the identity of the user, typically in an out-of-band mechanism. After the CA has issued a certificate, the user must load the certificate back into the wallet using a cut-and-paste mechanism.
- Certificate trustpoints. This is a list of certificate authorities that the user trusts.

The Oracle Wallet Manager gives users (and administrators of data servers) complete control over the contents of their wallets. The contents of the wallet are encrypted with a key based on a user-specified password. The wallet is stored, encrypted with the user's password, on the operating system.

The user needs only to remember one password, which opens the wallet, to access his or her credentials, which can then be used to authenticate the user to multiple

databases. Those databases need no longer store and manage local passwords for users.

Note: Oracle Wallet Manager is required to enable SSL. You must have a server-side wallet to set up an SSL connection.

Certificate Authority Users may obtain certificates from any X509v3 certificate authority, such as VeriSign, Inc. This is an offline and email process, and may take from a few days to weeks. Having a certificate issued in this manner guarantees its authenticity. Certificates can be loaded into Oracle wallets using Oracle Wallet Manager.

Note: While OSS provided a means to create user certificates, Oracle currently has no such functionality. As mentioned above, users may obtain certificates from any X509v3 certificate authority

Oracle Internet Directory The Oracle Internet Directory is an LDAPv3 compliant directory. It combines a native implementation of the Internet Engineering Task Force's (IETF) LDAP v3 standard with an Oracle8 back-end datastore. This design provides the Oracle LDAP Directory Server with a high degree of scalability, reliability, and compatibility with standard interfaces. It is specifically designed to meet the needs of managing user and system configuration data, including security attributes and privileges, in an Internet environment.

The management of Oracle wallets and their contents has already been discussed, but Oracle8i also extends the benefits of *enterprise roles* by storing them in and retrieving them from the Oracle Internet Directory. Oracle8 introduced enterprise roles, but they were stored in a proprietary repository.

Enterprise roles enable centralized authorization of users. For example, a user may be granted the enterprise role "HR clerk", which contains the global role "HR user" on the human resources database, and the "employee" global role on the corporate information database. If a user changes jobs, then his enterprise role assignment can be changed, altering his privileges in multiple databases throughout the enterprise. Also, an administrator can add capabilities to enterprise roles (granted to multiple users) without having to update the authorizations of each user independently.

Oracle Enterprise Security Manager Directory-enabled Security Manager, which is named Oracle Enterprise Security Manager in release 8.1.6, is an administrative tool that provides single station administration of privileges throughout the enterprise.

Security administrators can manage users and authorizations across multiple Oracle8i databases and Oracle Internet Directory from a single console.

Networking Improvements

The following new Net8 enhancements, described elsewhere, improve ease of use and scalability for the parallel server.

- A database can now have more than one service associated with it. See "[Service Naming](#)" on page 3-75.
- Connection load balancing balances the number of active connections among various instances and dispatchers for the same service. See "[Connection Load Balancing](#)" on page 3-77.
- If more than one listener supports a service, then clients can use client load balancing and connect-time failover. See "[Client Load Balancing](#)" on page 3-76 and "[Connect-Time Failover](#)" on page 3-76.

It is also now easier to set up, configure, and administer the Oracle network, using the new Net8 Easy Config and enhanced Net8 Assistant tools. These tools are not discussed in this book, but they are described in Net8 documentation.

For more information, see *Net8 Administrator's Guide*.

Service Naming

An Oracle database can span multiple instances over multiple computers, making the database service a distributed service.

Prior to Oracle8i, an Oracle database service was identified by an Oracle System Identifier (SID). The SID was also used internally by the database as pointer to the System Global Area (SGA). Clients connected to a database instance by specifying the SID in the connect descriptor. This naming scheme did not distinguish services from instances.

In Oracle8i, a new naming schema has been implemented. Because an Oracle database can span over multiple computers, both the service as a whole and each of its instances are specified.

A database is now identified by its service name with the `SERVICE_NAMES` parameter in the initialization file. `SERVICE_NAMES` specifies the name of the highest-level view of Oracle database service, which may span instances and/or nodes. `SERVICE_NAMES` is defaulted to the global database name, a name comprised of the database name (`DB_NAME`) and domain name (`DB_DOMAIN`).

Database instances are identified by an instance name with the `INSTANCE_NAME` parameter in the initialization file. `INSTANCE_NAME` corresponds to the `SID` of the instance.

Connect descriptors used by clients should be configured with the `SERVICE_NAME` (*singular*, not plural) parameter to connect to an Oracle8i database.

Optionally, the connect descriptor can also be configured with the `INSTANCE_NAME` parameter to connect to a particular instance of the database. This may be useful if you have an Oracle Parallel Server with multiple instances.

Service Registration

The PMON process—an instance background process—registers information with the listener. Prior to this release, information about the instance was required to be manually configured in the `listener.ora` file. Service registration provides the listener with information about:

- Service name(s) for each running instance of the database
- Instance name(s) of the database
- Service handlers (dispatchers and dedicated servers) available for instance
This allows the listener to direct a client's request appropriately.
- Dispatcher, instance, and node load information
This allows the listener to determine which dispatcher can best handle a client connection's request. If all dispatchers are blocked, then the listener can spawn a dedicated server for the connection.

This feature minimizes the system configuration, because there is no need to manually enter instance names in `listener.ora` files. Database instance registration enables dynamic connection load and connect-time failover.

Client Load Balancing

When more than one listener supports a service, a client can randomize requests to the various listeners. The randomization serves to distribute the load, so as not to overburden a single listener.

Connect-Time Failover

When more than one listener supports a service, a client can be configured to failover the client request to a different listener if the first listener fails.

Connect-time failover continues until the client successfully connects to a listener.

Connection Load Balancing

Connection load balancing improves connection performance by balancing the number of active connections among multiple dispatchers. In an Oracle Parallel Server environment, connection load balancing also has the capability to balance the number of active connection among multiple instances.

Because PMON can register with *remote* listeners, a listener can always be aware of all instances and dispatchers regardless of their location. This way, a listener can send an incoming client connection request to the least loaded instance and least loaded dispatcher that may be running on a different node.

If a database service has multiple instances on multiple nodes, the listener chooses the least loaded instance on the least loaded node. Based on the selected instance, the least loaded dispatcher is then chosen.

Net8 Configuration Assistant

The Net8 Configuration Assistant is a post-installation tool that lets you configure basic network components. It runs automatically after software installation. It can be used on either the client or server.

It may be also be run in stand-alone mode to configure naming methods usage, the listener, net service names in the `tnsnames.ora` file, and directory server access.

Native Authentication on Windows NT

The Net8 native authentication adapter on NT has been improved in the following ways:

- Authentication is faster. The server uses the same authentication mechanism that the client used to log into the NT system.
- The authentication adapter uses Security Support Provider Interface (SSPI) to authenticate.
- The client and server mutually authenticate each other. Previously, only the server authenticated the client

Backward compatibility with previous versions of Oracle is maintained, while integration with new security mechanisms in the future is easier.

For more information about native authentication on Windows NT, see the Windows NT documentation "Getting Started" guides for Oracle8i and Net8.

Multi-Tier Authentication and Authorization

Oracle8i also offers security to multi-tier architectures. In some systems, the middle tier is super-privileged to perform any action on behalf of any user, and the identity of the real client is not preserved through the middle tier. Middle tiers, especially Web servers or application servers, often sit on or outside a firewall, so limiting their access and auditing their actions is especially important.

Oracle8i provides the ability to preserve the real client identity through the middle tier and limit the users on whose behalf a middle tier can connect. The server can also audit actions taken by the middle tier on behalf of a particular user.

To support this feature, syntax changes have been made to the `ALTER USER` and `AUDIT` statement. For more information, see *Oracle8i SQL Reference*.

Program Interfaces

This section covers the Oracle8i new features and enhancements for the various program interfaces used to access data in the Oracle database. PL/SQL, the Pro*C/C++ and Pro*COBOL precompilers, and the Oracle Call Interface (OCI) are included. The new features supporting Windows NT integration are also discussed.

PL/SQL

For more information on any of the PL/SQL features or enhancement described here, see the *PL/SQL User's Guide and Reference*. Additional references may be noted in the descriptions.

Autonomous PL/SQL Blocks

Autonomous PL/SQL blocks are PL/SQL blocks that have a transaction scope independent of the transaction scope of the calling PL/SQL block. They can perform operations, commit, and rollback independent of the transactions in the calling block, before returning to the calling block. Transactions within an autonomous PL/SQL block are referred to as autonomous transactions. They are not nested transactions in that they do not share resources with the calling transaction and committed changes are immediately visible to other transactions regardless of whether the calling PL/SQL block commits or rolls back. A PL/SQL stored procedure/function, local procedure/function, a package/type method, or a top-level anonymous block can be declared to be an autonomous block using the pragma specification:

```
pragma AUTONOMOUS_TRANSACTION
```

The above specification can appear anywhere in the declaration section of a PL/SQL block, but at most once.

The Oracle8i autonomous transactions feature provides a means of constructing reusable application subcomponents that perform specialized operations.

Invoker Rights

The invoker-rights model, introduced in this release, allows programs to be executed with the privileges of the calling user. Prior releases used a *definer-rights* model wherein programs execute with the privileges of the creating user. The syntax of the CREATE FUNCTION, CREATE PROCEDURE, CREATE PACKAGE, and CREATE TYPE statements has been altered to allow an *invoker_rights_clause*. If AUTHID CURRENT_USER is specified, then the program becomes an invoker-rights program. If it is not specified, then the default is AUTHID DEFINER and there is no change from past releases.

The invoker-rights model is another step toward the creation of reusable code--a single code base that can be used by multiple users to manage data stored in their schemas.

For more information, see *Oracle8i Concepts*.

PL/SQL Bulk Binds

For PL/SQL, Oracle8i provides functionality to send a collection of bind variables (bulk bind) to the SQL engine with a single SQL statement, rather than requiring the same statement to be sent multiple times, differently bound each time. This functionality is similar to that already provided with OCI and Pro*C/C++. It reduces communication overhead and can significantly improve performance.

Dynamic SQL in PL/SQL

Dynamic SQL adds power and flexibility to host languages by allowing an application to generate and submit SQL statements for execution at run time. Prior releases provided this capability through programmatic interfaces in the DBMS_SQL package. Oracle8i provides dynamic SQL processing support in PL/SQL that makes dynamic SQL processing extremely simple, efficient, and easy to use. It is seamlessly integrated with SQL, much like static SQL in PL/SQL, and its performance is comparable to that of static SQL.

Parameter Passing by Reference

In Oracle8i, PL/SQL supports three parameter passing modes: IN, IN OUT, and OUT. IN parameters are passed by reference; IN OUT parameters support copy-in and copy-out semantics; OUT parameters support copy-out semantics. Through a new syntax, using NOCOPY mode, Oracle8i allows all parameters to be passed efficiently by reference. This achieves a significant performance gain for any application, such as a data cartridge, that passes large data structures as IN OUT and OUT parameters.

PL/SQL Procedures for Supporting REF-based Operations

Oracle 8i provides the UTL_REF package, containing procedures to support reference-based operations. Specifically, given a REF, there are procedures for selecting, locking, updating, and deleting an object.

Unlike SQL, the UTL_REF procedures enable the writing of generic type methods without knowing the object table name.

For more information, see *Oracle8i Supplied PL/SQL Packages Reference*.

Monitoring and Analysis of Program Execution

Oracle8i provides three new APIs, in the form of packages, that facilitate the tracing, debugging, or profiling (including code coverage) of PL/SQL applications.

For more information about any of these packages, see *Oracle8i Supplied PL/SQL Packages Reference*.

DBMS_TRACE This package provides a means for tracing the execution of PL/SQL programs on the server. It contains procedures to start and stop the tracing of calls to PL/SQL functions or procedures, and exceptions. Data is written to the database trace log file.

DBMS_DEBUG Third-party tool vendors can leverage this low-level API to provide debuggers to PL/SQL developers.

Debugging requires two database sessions. One session runs the target code, which invokes DBMS_DEBUG to cause the PL/SQL server to generate debug events. The other session, the supervising session or debugger, uses DBMS_DEBUG to communicate with and read events posted by the target session. Several debuggers are already available using this API, such as Oracle Procedure Builder.

DBMS_PROFILER This profiling API provides PL/SQL programs with services for collecting and persistently storing data, in database tables, that can help identify

and isolate performance problems or provide code coverage information. It is possible to generate profiling data for all named library units in a single session. For this release, the data includes the total number of times each line of code is executed, the total amount of time spent executing the line, and minimum and maximum times spent on a particular execution of the line.

A simple sample suite of analysis and reporting tools is provided, and the use of database tables allows ad hoc querying. Third-party tool vendors can now provide more sophisticated analysis and reporting tools for profiling and code coverage.

For more information on these packages, see *Oracle8i Supplied PL/SQL Packages Reference*.

Increased Package Body Size

Oracle8i relaxes the limit on the size of PL/SQL package bodies.

Purity Rules Relaxed

In earlier releases, users were expected to qualify functions and procedures with the `RESTRICT_REFERENCES` pragma, which has various options (`WNDS`, `RNDS`, `WNPS`, `RNPS`). This specifies the purity of a function or procedure if the user wishes to use the function for certain operations. This purity information, combined with additional computed information, is then used to determine whether it is safe to perform the following operations:

- Calling a function from a query or DML statement
- Parallelizing a query or DML statement which calls a function
- Creating a snapshot with fast refresh

In Oracle8i, these rules have been relaxed to be more user-friendly.

More details about this, refer to the *Oracle8i Application Developer's Guide - Fundamentals*.

Precompiler Enhancements

The Pro*C/C++ and Pro*COBOL compilers, which allow the embedding of SQL statements in source code, have been significantly enhanced for Oracle8i.

Pro*C/C++

The following table lists the Oracle8i enhancements for the Pro*C/C++ Compiler.

Enhancement	Description
CALL statement	The CALL embedded SQL statement invokes a stored procedure. It can be used instead of an embedded PL/SQL block in new applications.
LOB support	A new embedded SQL interface allows the manipulation of LOBs (large objects).
Collections	A new embedded SQL interface to allows the manipulation of varray and nested table data.
ANSI dynamic SQL	This is the implementation of ANSI dynamic SQL Method 4.
PREFETCH option	This precompiler option speeds up database access by "prefetching" values, thus cutting down the number of network round-trips.
External routines	External routines written in C can be called from PL/SQL blocks. The REGISTER CONNECT embedded SQL statement is used by the procedures.
HEADER option	The precompiler option, HEADER, specifies that precompiled header files are to be created and used to reduce the time and computer resources needed for developing large projects.
Calling Java from PL/SQL	Stored procedures written in Java can be called from an application.
DML returning clause	This clause is allowed in INSERT, DELETE, and UPDATE statements.
Universal ROWID	The universal ROWID datatype is supported. Index-organized tables use this concept.
SYSDBA/SYSOPER privileges in CONNECT statements	These privileges can now be set in the CONNECT statement.
CLOSE_ON_COMMIT precompiler option	The CLOSE_ON_COMMIT micro precompiler option allows a choice of whether or not to close all cursors when a COMMIT is executed and the macro option MODE=ANSI.

For more information, see the *Pro*C/C++ Precompiler Programmer's Guide*.

Pro*COBOL

The Pro*COBOL precompiler was largely rewritten for release 8.0, introducing many new features, but with an emphasis on greater compatibility with DB2. Oracle8i continues Oracle's effort to enhance the functionality of the Pro*COBOL

precompiler, providing support for new Oracle8i database server features. The following table summarizes the Oracle8i enhancements.

Enhancement	Description
CALL statement	The CALL embedded SQL statement invokes a stored procedure. It can be used instead of embedded PL/SQL block in new applications.
LOB support	An embedded SQL statement interface allows LOBs (large objects) to be used in precompiler applications.
ANSI dynamic SQL	This is the implementation of ANSI dynamic SQL Method 4.
PREFETCH option	This precompiler option speeds up database access by "prefetching" values, thus cutting down the number of network round-trips.
Calling Java from PL/SQL	Stored procedures written in Java can be called from a Pro*COBOL application.
DML returning clause	This clause, which can save round-trips to the database server, is now allowed in INSERT, DELETE, and UPDATE statements.
Universal ROWIDs	The support for universal ROWID datatype is presented. Index-organized tables use this concept.
User-specified runtime contexts	COBOL subprograms can now be called by C programs. The context can be passed by a Pro*COBOL program to a C subprogram
SYSDBA/SYSOPER privileges in CONNECT statements	These privileges may now be set using the CONNECT statement.
Tables of group items	Tables of group items are now allowed as host variables in Pro*COBOL.
WHENEVER Do Call Branch	The WHENEVER directive has a DO CALL action: a subprogram is called.
Decimal-Point is Comma	The DECIMAL-POINT IS COMMA clause is supported. This permits commas to be used instead of decimal points in numeric literals.
Optional divisions	The following divisions and their contents are now optional: IDENTIFICATION, ENVIRONMENT, DATA.
NESTED Option	When set to NO, the NESTED precompiler option prevents generation of the GLOBAL clause for non-nested programs.

For more information, see *Pro*COBOL Precompiler Programmer's Guide*.

Java Interfaces

The Java interfaces, JDBC and SQLJ, are discussed in the Java section of this document. See:

- "JDBC Drivers" on page 3-17
- "SQLJ Translator" on page 3-18

OCI Enhancements

For information about these OCI enhancements, see *Oracle Call Interface Programmer's Guide*.

DML Returning Enhancements

Enhancements have been made to make the DML RETURNING clause more optimized for common usage patterns. Data transfer is minimized and the interface is simplified for these common cases. Specifically, OCI provides selective data transfer extensions to the use of the RETURNING clause which returns data only when it has been changed in the server. Also, there is additional OCI functionality for single row and array DML operations in which each iteration returns more than one row.

Enhanced Array DMLs

OCI now provides an error-batching mode for an array DML statement execution. In prior releases, OCI aborted if an error was signalled during an array DML operation. Oracle8i can allow the DML to complete, but trap any erroneous rows. The erroneous rows are returned at one time after inserting/deleting all correct rows. The user can then make necessary corrections, and do another update operation. Consequently, the number of round-trips a user might have to perform has been reduced to a maximum of two.

OCI Callback Registrations

OCI callbacks now provide a means of tracing OCI calls for debugging and performance measurements. Two types of callbacks can be registered, entry and exit, wherein additional pre or post processing can be done. They can also be used to substitute the body of the function with proprietary code to execute on a foreign data source. For licensed ISVs and internal Oracle groups, a dynamic callback registration mechanism is provided. This mechanism allows a core-compliant

dynamically linked module to register callbacks at OCI environment initialization time. This is transparent to the OCI application and does not require any changes in a running application.

OCIDescribeAny Enhancements

`OCIDescribeAny` was introduced in Oracle8 as a simple and uniform OCI interface to describe schema-level database objects: tables, views, sequences, synonyms, procedures, functions, packages, and types. It has been expanded to describe attributes of sub-schema objects like columns of a table or fields in a type, and to include more attributes of currently described database objects.

Asynchronous Event Notification

The publish/subscribe and system event notification capabilities of Advanced Queuing were discussed in "[AQ-Based Publish/Subscribe](#)" on page 3-38. OCI provides a programmatic API through which an application or user can subscribe to be notified of a database or system event, and be notified asynchronously through a callback.

Non-Blocking OCI

Oracle8i provides non-blocking capability with OCI. Essentially, it allows a user to specify or set a non-blocking polling mode for OCI call execution. This mode can either be specified for a whole application or at the individual call level. OCI non-blocking mode allows applications to process other events while an OCI call is in progress.

Windows NT Integration

Microsoft Windows NT® is an increasingly popular environment for database development and deployment. Oracle8i contains several features to make it easier to develop Oracle-based applications with Microsoft products. The Oracle database server on Windows NT provides a highly available, scalable, and secure platform for application deployment. This section discusses the features supporting Windows NT integration offered in Oracle8i.

Oracle Application Generator for Microsoft Visual Studio

The Oracle Application Wizard (AppWizard) provides developers with a GUI tool with which to quickly and seamlessly create a Visual C++, Visual InterDev, or Visual Basic application that provides connectivity and data access to an Oracle

database. A developer can use this tool to generate an Oracle database enabled application that compiles, links, and runs without writing a single line of code.

AppWizard is thoroughly integrated with the Visual Studio IDE, which is the most popular IDE for developing C++ and Visual Basic applications on the Windows NT/95 platform. It allows ISVs, VARs, and other users to easily build applications that leverage Oracle database technology in the Visual Studio IDE.

The wizard can be invoked whenever a developer decides to create a new project in Visual Studio. AppWizard guides the developer through a two-step process. First, the wizard prompts the user with questions about the database connection that is to be used: connection string, username, and password to the Oracle database. Then, the user is allowed to select the specific tables and columns from which the application retrieves data.

Second, based on the developer's responses, AppWizard generates a Visual Studio project and source code that provide the developer with a custom tailored application framework from which to start developing immediately. The generated code framework consists of a mixture of Microsoft Foundation Classes (MFC) and Oracle Objects for OLE classes. The MFC code provides the basic GUI application code and the OO4O classes provide the connectivity and data access to Oracle databases.

Oracle Objects for OLE

Oracle Objects for OLE (OO4O) is a COM-based database connectivity tool that provides seamless and optimized access to Oracle databases. OO4O can be used in environments ranging from the typical two-tier client/server applications to application servers deployed in n-tier environments and Web servers such as Microsoft IIS or MTS. It can be used from virtually any programming or scripting language that supports the Microsoft COM Automation technology, such as Visual Basic, Visual C++, VBA in Excel, VBScript and JavaScript in IIS Active Server Pages. OO4O consists of an *in-process COM Automation Server*, a *C++ class library*, and the *Oracle Data Control*.

The Oracle8i version of OO4O enables developers who utilize COM/DCOM based development tools to have seamless access to all Oracle specific features that are presently inaccessible from other ODBC or OLE DB-based components such as ADO. Significant enhancements in this release include:

- Efficient management of database connections and user sessions (connection pooling and multiplexing)
- Full support for Microsoft Transaction Server (MTS) coordinated transactions

- Seamless access to instances of new datatypes introduced in Oracle8 including:
 - Object References (REFs)
 - Object Instances (Objects)
 - Nested Tables
 - Varrays
 - BLOBs, CLOBs, NCLOBs and BFILEs
- Easy-to-use interface for describing schema objects
- Full support for accessing Oracle Advanced Queuing capabilities

Extensive online help is available with Oracle Objects for OLE, including a *Quick Tour* for getting started.

Microsoft Transaction Server Integration

Microsoft Transaction Server (MTS) is a middle-tier application server for managing distributed applications. It offers an ActiveX/DCOM based programming model to develop distributed applications and a runtime environment to deploy these applications. Oracle8i provides full, native integration with MTS. It allows a user to create distributed transactions coordinated by MTS with Oracle databases on any platform as participants.

Prior to Oracle8i, customers used Microsoft's XA solution for MTS to support transactions involving Oracle data. The transactions are coordinated by MTS but are mapped to the XA protocol. But Oracle8i now provides much better integration by supporting the OLE transaction interfaces published by Microsoft. This solution performs much better and with less configuration.

COM Cartridge

This cartridge allows PL/SQL developers to programmatically manipulate COM/DCOM objects using the OLE automation interface (IDispatch). By allowing COM manipulation in PL/SQL stored procedures, functions, and triggers, developers have access to the thousands of pre-built and reusable COM objects that exist today. This feature enables developers to leverage COM objects that they build or buy from third-party vendors. Furthermore, it allows customers to extend the functionality of PL/SQL to suit their requirements.

Release 8.0 New Features and Options

This chapter contains the following sections:

- [Partitioned Tables and Indexes](#)
- [Improved Data Warehouse Performance](#)
- [Improvements for OLTP Applications](#)
- [Replication](#)
- [Object-Relational Technology](#)
- [Migration and Interoperability](#)
- [Other Enhancements](#)

Partitioned Tables and Indexes

Very large databases present significant challenges for administrators and application developers. Large amounts of data complicate administrative tasks and affect the availability of the database.

To improve availability, ease administration, and enhance query and DML performance, the Oracle8 Enterprise Edition allows tables and indexes to be partitioned, or broken up, into smaller parts based on a range of key values. Because partitions operate independently of each other, data in a partitioned table are available even if one or more partitions are unavailable. Partitions also make large tables easier to manage by breaking up administrative operations into smaller tasks, which in turn can be performed in parallel. Finally, partitioning a table or index can improve performance of operations on the data by eliminating unneeded partitions from the execution plan of the operation.

Partitioning is transparent to both applications and users, so standard SQL statements in existing applications can run against partitioned tables. The Oracle optimizer is partition-aware, and partitions which do not contain any data required by a query are eliminated from the search, often resulting in a substantial performance increase.

What is Partitioning?

A table or index can be partitioned or divided into smaller pieces. You define a table or index partitioning strategy when creating the structure. Pick a column or set of columns to act as a partition key, and this key determines which data is placed into each partition. Data from insert operations is automatically placed into the appropriate partitions, so applications do not need to be rewritten to take advantage of partitioning.

All operations you perform on regular tables can be performed on individual partitions of a partitioned table. For example, you can export or back up single partitions of a table to avoid exporting or backing up the entire table in one operation. You can also perform the `ANALYZE` statement on each partition concurrently to compute statistics needed for cost-based optimization more quickly.

Each partition of a partitioned table or index operates independently. Therefore, operations on one partition are not affected by the availability of other partitions. If one partition becomes unavailable because of a disk crash or administrative operations, then both query and DML operations on data in other partitions can still continue.

Greater Control Over Data

Administrators can specify storage attributes for each partition and the placement of the partition within the host file system, increasing the granularity of control for very large databases. Partitions can be individually taken off-line or brought on-line, backed up, recovered, exported and imported, and loaded—thereby limiting the time required for management operations. An individual index partition can be built for one table partition, reducing the time required for index maintenance operations. Partition operations can be performed in parallel. Partitions increase availability by isolating media and application failures—applications not requiring data from an inaccessible partition continue to run without impact.

Easy-To-Use Administration

The Oracle8 Enterprise Edition provides a set of easy-to-use management statements for partitions. For example, you can implement a table containing a rolling time window of data with a partitioned table so that only one year's worth of data is ever contained in the table. To accomplish this, just add new partitions to the end of the table and drop partitions containing data more than a year old. You can also split and merge partitions easily to balance data among partitions. Also, stand-alone tables that use Oracle7 release 7.3 UNION ALL views can be converted into partitions of a single table. This allows users with very large Oracle7 databases to quickly benefit from Oracle8 Enterprise Edition partitioning without having to rebuild large tables.

Improved Data Warehouse Performance

The Oracle8 Enterprise Edition introduces new features that improve data warehousing performance:

- [Enhanced Star-Query Processing](#)
- [New Parallel Operations](#)
- [Increased Database Size](#)

Enhanced Star-Query Processing

The Oracle8 Enterprise Edition introduces performance improvements to the processing of star queries, which are common in data warehouse applications. A star query, or star schema, occurs when one or more very large tables, often called fact tables, have relationships to multiple smaller tables called dimension tables.

Oracle7 introduced the functionality of star query optimization, which provides performance improvements for these types of queries. In the Oracle8 Enterprise Edition, however, star-query processing has been improved to provide even better optimization for star queries.

In the Oracle8 Enterprise Edition, a new method for executing star queries has been introduced. Using a more efficient algorithm, and utilizing bitmapped indexes, the new star-query processing provides a significant performance boost to data warehouse applications.

The Oracle8 Enterprise Edition has superior performance with several types of star queries, including star schemas with "sparse" fact tables where the criteria eliminate a great number of the fact table rows. Also, when a schema has multiple fact tables, the optimizer efficiently processes the query. Finally, the Oracle8 Enterprise Edition can efficiently process star queries with large or many dimension tables, unconstrained dimension tables, and dimension tables that have a "snowflake" schema design.

The Oracle8 Enterprise Edition's star-query optimization algorithm, unlike that of Oracle7, does not produce any Cartesian-product joins. Star queries are now processed in two basic phases. First, the Oracle8 Enterprise Edition retrieves exactly the necessary rows from the fact table. This retrieval is done via bit mapped indexes and is very efficient. The second phase joins this result set from the fact table to the relevant dimension tables. This allows for better optimizations of more complex star queries, such as those with multiple fact tables. The new algorithm uses bit-mapped indexes, which offer significant storage savings over previous methods that required concatenated column B-tree indexes. The new algorithm is also completely parallelized, including parallel index scans on both partitioned and non-partitioned tables.

For more information, see *Oracle8i Concepts*.

New Parallel Operations

Insert, update, and delete operations can now be run in parallel in the Oracle8 Enterprise Edition. These operations, known as parallel DML, are executed in parallel across multiple processes. By having these operations execute in parallel, the statement is completed much faster than if the same statement were executed in a serial fashion. Parallel DML complements parallel query by providing parallel transaction execution as well as queries. Parallel DML is useful in a decision support (DSS) or data warehouse environment where bulk DML operations are common. However, parallel DML operations can also speed up batch jobs running in an OLTP database.

The Oracle8 Enterprise Edition supports parallel inserts, updates, and deletes into partitioned tables. It also supports parallel inserts into non-partitioned tables. The parallel insert operation on a non-partitioned table is similar to the direct path load operation that is available in Oracle7. It improves performance by formatting and writing disk blocks directly into the datafiles, bypassing the buffer cache and space management bottlenecks. In this case, each parallel insert process inserts data into a segment above the high watermark of the table. After the transaction commits, the high watermark is moved beyond the new segments.

To use parallel DML, it must be enabled prior to execution of the insert, update, or delete operation. Normally, parallel DML operations are done in batch programs or within an application that executes a bulk insert, update, or delete. New hints are available to specify the parallelism of DML statements.

For more information, see the *Oracle8i Parallel Server Documentation Set: Oracle8i Parallel Server Concepts; Oracle8i Parallel Server Setup and Configuration Guide; Oracle8i Parallel Server Administration, Deployment, and Performance*.

Increased Database Size

The Oracle8 Enterprise Edition can manage databases of hundreds of terabytes in size because of partitioning, administrative improvements, and internal enhancements. Many size limitations in earlier versions of Oracle have been raised, such as the number of columns per table, the maximum database size, and the number of files per database.

Improvements for OLTP Applications

Demanding OLTP applications benefit from a number of new features which improve scalability, performance, and manageability.

Extended Backup/Recovery Subsystem

Oracle8 and the Oracle8 Enterprise Edition's server-managed backup and recovery provides a better backup and recovery functionality integrated within the Oracle database server. Detailed information is maintained on when backups are performed, exactly which parts of the database are backed up, and where the files are stored. Should a recovery be necessary, Recovery Manager analyzes the state of the database and determines the operations necessary to repair the database. It then automatically performs those operations, greatly simplifying the recovery for the administrator and reducing the possibility of human error. A simple GUI interface within Oracle Enterprise Manager controls backup and recovery. An API is also

available for third parties who may wish to provide an alternative interface. Media management layer interfaces to popular third-party tape management products are available. Legato Storage Manager is provided free with both Oracle8 and the Oracle8 Enterprise Edition.

The Oracle8 Enterprise Edition provides multilevel, incremental backups that greatly reduce the size of the backups, because only the changed blocks are backed up. This can substantially reduce the time required to back up a datafile.

Tablespace point-in-time recovery allows one or more tablespaces to be recovered to an earlier time, while the remainder of the database is up and running. This allows many types of user errors to be easily corrected. For example, if a user runs a batch job that incorrectly updates many records in a table, then the table can be restored to a time previous to the batch job. Also, if a table is accidentally dropped or truncated, then it can be restored to a time before this operation.

For more information, see *Oracle8i Backup and Recovery Guide*.

Very Large User Populations

Numerous enhancements throughout the Oracle8 and Oracle8 Enterprise Edition database servers and Net8 increase the utilization of operating system and networking resources. Connection pooling temporarily drops the physical connection for idle users (and transparently re-establishes the connection when needed), thus increasing the number of users that can be supported. Oracle Connection Manager can be used to configure a middle tier that manages the connections of very large user populations. By configuring multiple connection managers, the Oracle8 Enterprise Edition can support tens of thousands of concurrent users. Shared database links multiplex many users into a database server with a single connection, reducing resource requirements, especially for multitier application architectures.

Advanced Queuing

The Oracle8 Enterprise Edition Advanced Queuing feature supports database messaging, or queuing, through a set of queue tables and queue functions.

Advanced Queuing adds direct support in the database for deferring transaction execution to a later time and executing transactions in a particular order. This capability allows you to decouple distributed applications and eliminates dependency on external systems for applications requiring high scalability. Enqueue and dequeue operations can be used to shift processing from within a transaction to a background process, thereby improving transaction response time.

Also, queuing can be used to implement work flow applications that move data to a system as the state of the data changes. An example of this is moving orders from an order-entry application to a shipping application, and then to a billing system during the life of an order. Advanced Queuing can also work in conjunction with popular TP monitor queuing systems.

Messages can be enqueued and dequeued by applications or other queues. The propagation feature enables applications to communicate with each other without having to be connected to the same database or to the same queue. Using the familiar database links and Net8, messages can be propagated from one queue to another, irrespective of whether these are local or remote.

For more information, see *Oracle8i Application Developer's Guide - Advanced Queuing*.

Parallel Server Improvements

Numerous enhancements in Oracle Parallel Server improve performance, scalability, memory usage, and availability. A common and integrated distributed lock manager replaces the lock managers provided by the different operating system vendors, improving performance and portability on most platforms. Several performance enhancements make the Oracle Parallel Server significantly faster. Also, new global VS tables improve manageability.

Oracle Parallel Server now uses its own integrated distributed lock manager (DLM) for processing inter-node requests for resources. Previously, Oracle relied on operating system vendors to supply DLM capability for the Parallel Server to run on a given platform. Integrating the DLM within the Oracle8 Enterprise Edition allows the Oracle Parallel Server to be available on platforms that previously did not support cross-node sharing of resources.

The Oracle8 Enterprise Edition introduces numerous improvements to Oracle Parallel Server performance. System change number (SCN) generation is now optimized for providing SCNs across instances more efficiently. This improvement alone can amount to a 10 to 15 percent improvement in Oracle Parallel Server performance. Also, the DLM now caches locks to avoid cross-node communication of lock information. Pinging, or contention for blocks across nodes, is also reduced by improving the algorithm for controlling access to contended blocks. Reverse-key indexes reduce "hot spots" in indexes, especially primary key indexes, by reversing the bytes of the leaf blocks and thus eliminating the contention for leaf blocks across instances. Partitions also help Parallel Server performance by allowing database administrators to map partitions to specific nodes, which can dramatically reduce pinging. Finally, the Oracle8 Enterprise Edition reduces the memory overhead associated with fine-grain locking in an Oracle Parallel Server.

The Oracle8 Enterprise Edition allows designation of groups of instances for parallel query or parallel DML processing. You can specify to which group an instance belongs and then use the group to process statements from certain applications. This is especially useful for separating OLTP from data warehouse processing among your clustered servers. Using separate instances in an Oracle Parallel Server ensures that data warehouse queries do not affect the performance of OLTP applications.

The Oracle8 Enterprise Edition introduces global fixed views (GV\$) for the Oracle Parallel Server. This allows administrators to log into one instance of a Parallel Server and view global views that contain data from all the V\$ views in the cluster. This makes administration operations significantly simpler and improves the productivity of database administrators.

Transparent Application Failover (TAF)

Should a node in the Oracle Parallel Server fail, transparent application failover migrates your connections and automatically re-establish their sessions on another node. Your applications continue to run, and you may be unaware of the failure. This provides continuous availability in the event of scheduled and unscheduled outages. Even if you are not using the Oracle Parallel Server option, TAF automatically reconnects and reestablishes your session.

Depending on the amount of overhead you are willing to incur on the client and backup machines, the instance failure could result in a completely transparent migration of user connections to the failover node. You can specify that you want all queries to be cached on the client so that they can be reinstantiated on the failover node. Also, you can have a session pre-connected to the failover node, eliminating the time required to reconnect to a failover instance.

Transparent application failover is useful not only for availability, but also for manual load-balancing or orderly shutdown of the system. If too many users connect to an instance, then you can terminate their sessions and have them transparently migrate to another node. You can also shut down a node and have users transparently migrate to a failover node after their current transaction completes. Shutdowns no longer need to interrupt users work.

Note: To take advantage of this functionality, applications must be written specifically using new Oracle8 OCI calls.

For more information, see *Oracle8i Administrator's Guide*.

Improved TP Monitor Support

Dynamic XA support improves performance for multitier applications with industry-standard XA-compliant Transaction Processing (TP) monitors. Oracle8 and the Oracle8 Enterprise Edition improve support of the XA interface with support for dynamic registration and loosely-coupled transaction branches. They also offer better performance and recovery when used with Oracle Parallel Server.

For more information, see *Oracle8i Application Developer's Guide - Fundamentals* or the *Oracle8i Parallel Server Documentation Set: Oracle8i Parallel Server Concepts; Oracle8i Parallel Server Setup and Configuration Guide; Oracle8i Parallel Server Administration, Deployment, and Performance*.

Better Security Administration

Oracle8 and the Oracle8 Enterprise Edition include a Security Server which provides a single sign-on environment for centrally administering users and roles. Oracle Security Server is compliant with the X.509 certificate-based security standard for public/private key authentication. Mutual authentication between client and server is also supported for protection against "rogue" databases designed to capture client communication. Also, a digital signatures tool kit is provided for creating applications designed to identify unauthorized tampering with data.

Oracle8 and the Oracle8 Enterprise Edition also contain improved password maintenance and administration. You can now define a password profile to enforce a security scheme. Passwords can expire after a certain time, or be checked for complexity (e.g., minimum length). You can define your own policy or use the standard stored function for checking the length, content, or reuse of new passwords. Also, you can create user accounts so that the user must immediately change the password upon the first access to the system.

Oracle8 and the Oracle8 Enterprise Edition also add support for privileged database links, which make it unnecessary to embed a password in a database link, and data encryption services callable from OCI and PL/SQL.

Replication

Oracle8 and Oracle8 Enterprise Edition replication provides improved support for mass-deployment distributed systems, higher throughput failover configurations, and specialized data warehouse systems.

Increased replication performance is now possible. Changes to replica sites can be propagated and applied in parallel, effectively removing the throughput limit on replication, while maintaining transactional integrity. Also, much of the replication functionality has been rewritten and moved from PL/SQL triggers into C code inside the Oracle engine, providing a significant boost in performance. Finally, the amount of data sent over the network for each replicated transaction has been reduced, and network round-trips are kept to a minimum.

Salesforce automation and other mass deployment systems are now easier to administer and design with the subsetting capabilities of both Oracle8 and the Oracle8 Enterprise Edition. Complex subquery snapshots can now include a fast refresh clause which allows you to build a snapshot of select rows based on a query to another table. For example, you can build a snapshot for each salesman based on the `ORDERS` table. That snapshot can contain only the orders by selecting the orders based on criteria from an `ASSIGNMENTS` table. Each salesman sees only his orders and only has to interact with the master table when uploading new orders.

Oracle8 Enterprise Edition replication also provides numerous improvements in manageability and ease of use. Significant enhancements to Oracle Replication Manager, including wizards, make replicated environments easier to set up and maintain. Additional improvements to replication manageability include snapshot registration, which provides valuable information about which sites have associated snapshots; primary key snapshots, which allow faster snapshot refresh after a table reorganization; a new security algorithm, which is easier to administer; and fine-grain quiesce, which allows an administrator to make schema-level changes to one replication group while others continue to operate.

For more information, see *Oracle8i Replication*.

Object-Relational Technology

The Oracle8 Enterprise Edition makes a major leap in data management technology with the introduction of an object-relational paradigm. Database schemas and applications today are becoming increasingly complex. Often, several separate applications with similar data, such as customer information, billing, and shipping, exist in different database schemas and an MIS department must manage the interoperation. Corporate management of the information becomes a difficult task of integrating different relational objects and different applications, possibly from different vendors, into a more coherent end-user data model. By enhancing the relational database with object extensions, Oracle addresses the need to simplify data modeling and extend the database with new datatypes.

The new, object-relational features include the following:

- [Object Types and Views](#)
- [Calling External Procedures From Within the Database](#)
- [Client-Side Support for Objects](#)
- [Evolution of Relational Environments](#)
- [Development Tools for Object Modeling](#)
- [Multimedia Data](#)
- [Java](#)
- [Extensibility](#)

Object Types and Views

Object types provide a way to extend Oracle's relational datatype system. Relational databases support three datatypes: characters, numbers, and dates. Object types allow you to define new datatypes and use them as you would regular relational datatypes. For example, you can create a new type called Address. This object type can have data, called attributes, such as Street, City, and Postal Code. The object type can also have methods or stored procedures, such as Distance, for computing the distance between addresses. These methods can be written in either PL/SQL or C. An address can then be used anywhere a regular datatype could, whether in column definitions, in PL/SQL variables, or even as the definition for an object table.

Oracle's object types can use powerful object modeling techniques for complex objects. For example, you can represent collections of similar objects in array structures or nested tables. You can also store references to objects for fast traversal without joining tables.

Object types allow application developers to code application logic in the database or the middle-tier application server, as opposed to using client-side code. All applications can then share the logic of the new datatypes so developers do not need to rewrite the code. This feature provides the advantages of creating re-usable code components and transparent application partitioning so that the code can reside and execute on the tier that yields the best performance: client, application server, or database server.

The Oracle8 Enterprise Edition follows the emerging SQL3 standard for object type definition and object modeling techniques. SQL3 defines syntax for creating and modifying object types, generating and storing object identifiers (OIDs), creating references or pointers to objects, and modeling collections of similar objects.

For more information, see *Oracle8i Concepts* and *Oracle8i Application Developer's Guide - Fundamentals*.

Calling External Procedures From Within the Database

Oracle8 and the Oracle8 Enterprise Edition provide a safe, fast way for the database to make a call to an external program. The call also can be made through open protocols like HTTP or IIOP (a CORBA standard). External procedures allow you to use existing application code, or write highly-optimized code for specific purposes, such as a computationally complex algorithm like Fast Fourier Transform (FFT). Also, you can use external procedures to interface with other applications or with specialized devices like embedded systems.

For more information, see *PL/SQL User's Guide and Reference* and *Oracle8i Application Developer's Guide - Fundamentals*.

Client-Side Support for Objects

The client-side object cache allows user applications to retrieve a complex hierarchy of objects into an application cache. The application can then traverse the objects without performing additional network retrievals. This provides a convenient and fast way to use objects in a client application and write code that is more like the native object-oriented code.

A new utility, the Object Type Translator, has been introduced as a quick way of generating header and implementation files for applications running against object schemas.

Evolution of Relational Environments

The Oracle8 Enterprise Edition is designed to allow users to easily evolve into using the new object-oriented functionality, as all existing applications are upwardly compatible. The new object-relational extensions are built on the same foundation as the relational functionality, which means that users do not have to discard or rewrite their existing relational applications before migrating to the Oracle8 Enterprise Edition. Unlike other object-relational databases, this design allows the older relational applications, which still read and write rows and columns, to coexist with new object-oriented applications, which read and write objects. The Oracle8 Enterprise Edition provides object views to retrieve relational data and represent the data to a client as if it were an object and vice-versa.

For example, an existing relational order-entry system might need a new front-end for the World Wide Web. The existing applications accessing the relational schema

can remain in operation, and a new set of object views can be developed as an object representation for the client. New and old applications can be based on the same data, but each has its own representation.

For more information, see *Oracle8i Concepts*.

Development Tools for Object Modeling

The Object Database Designer offers object support in a design tool. It helps you design, create, and access object-based Oracle8 Enterprise Edition systems.

For more information, see *Object Database Designer*.

Multimedia Data

Large Objects (LOBs) handle unstructured data such as images, sounds, video, and text, and have much richer functionality than their predecessors, LONG and LONG RAW. Character LOBs (CLOB or NCLOB), Binary LOBs, and BFILES (externally stored LOBs), can be replicated and can be an attribute of an object. You can also have more than one LOB per table. LOBs also have a greater maximum size than LONGs and have different mechanisms for maintaining read-consistency and random access.

LOB data is indexed for fast access starting at a specified byte. For example, you can read/write at specific byte-offsets. You can also read/write LOBs through the Oracle8 buffer cache or access them directly from disk.

LOB functionality is available with both Oracle8 and the Oracle8 Enterprise Edition and does not require the Objects option to be licensed and installed.

For more information, see *Oracle8i Concepts*.

Java

Oracle has a long history of supporting cross-platform environments. Oracle8 and the Oracle8 Enterprise Edition build upon that with strong support for Java.

Oracle currently has two methods for accessing Oracle data from Java programs: an Oracle-provided JDBC driver integrated with Oracle's object types, and JSQL for embedding SQL statements into Java code. Oracle provides its own JDBC drivers for better performance. JSQL allows you to include SQL statements in a Java application. The JSQL precompiler then converts the SQL into JDBC calls, which allows you to use existing SQL code in new Java applications.

For more information, see *Oracle8i JDBC Developer's Guide and Reference*.

Extensibility

Extensibility allows you to define your own datatypes. You can extend the capabilities of your current database by creating new datatypes for your specific applications. These new datatypes can be used in the same manner as the ones you have now; the same operations can be performed. Image, Spatial, Time Series, Visual, and ConText cartridges are now available as extensions to the Oracle database server.

Migration and Interoperability

A simple and fast migration utility rebuilds the data dictionary and converts the control files, log files, and data blocks. The migration utility converts any Oracle 7.1, 7.2, or 7.3 database into an Oracle8 or Oracle8 Enterprise Edition database. Oracle7 applications run unchanged against either of the Oracle8 products. Distributed statements from either of the Oracle8 products run against Oracle7, and vice-versa.

For more information, see *Oracle8i Migration*.

Other Enhancements

Index-Organized Tables

Index-organized tables store the data columns of a table within the leaf nodes of the Oracle8 B-tree index structure. This reduces overall storage requirements when most columns are indexed by storing the columns only once, rather than in both an index and a separate table. Index-organized tables also reduce access time by retrieving all columns from one location instead of two.

For more information, see *Oracle8i Concepts*.

Reverse Key Indexes

Reverse key indexes reduce the "hot spots" in indexes, especially ascending indexes. Unbalanced indexes can cause the index to become increasingly deep as the base table grows. Reverse key indexes reverse the bytes of leaf-block entries, therefore preventing "sliding indexes".

For more information, see *Oracle8i Concepts*.

Improved Constraint Processing

Constraint processing has been dramatically improved. You can now use non-unique indexes to enforce `UNIQUE` and `PRIMARY KEY` constraints. This eliminates redundant indexes and permits indexes to remain valid while constraints are disabled. Deferred-constraint checking shifts integrity-constraint checking from the end of statement execution to when a commit is issued. This simplifies the coding of certain operations involving integrity constraints. Also, all constraints can be enabled concurrently and in parallel with other constraints, and permit concurrent DML while the enabling continues.

Two Character Sets in One Database

In both Oracle8 and the Oracle8 Enterprise Edition, a new `NCHAR` datatype allows a second character set in one database. This improves performance and storage predictability for some Asian language, multibyte character set databases. Also, fixed-length native Unicode 2.0 and Chinese GBK support is new.

Miscellaneous Improvements

A non-updatable view (i.e., a joined view) can be updatable through the use of a new type of trigger, the `INSTEAD OF` trigger. This trigger allows you to replace `INSERT`, `UPDATE`, and `DELETE` operations on views with your own logic so that even views based on complex joins can be modified. Also, you can place subqueries in the select list of other queries to return result sets from an arbitrary list of detail tables without performing join operations. By placing a cursor expression in a select list, you can fetch rows from detail tables in a 3GL program without specifying complex join conditions in the `FROM` clause.

Oracle8i Feature and Option Availability

The Oracle8i database is available as Oracle8i, Oracle8i Enterprise Edition, and Oracle8i Personal Edition. Not all features and options are available with all editions of the Oracle8i database.

Oracle8i Personal Edition is available on Windows NT and Windows98 and includes, at no extra cost, all features and options that are available with Oracle8i Enterprise Edition, such as Oracle *interMedia* and Oracle Partitioning. Oracle8i Personal Edition on Windows98 includes the database server and most related products and features, depending on whether the required functionality is supported on the Windows98 platform.

For more details on Oracle8i feature and option availability and descriptions, see the "Oracle8i: A Family of Database Products" document at www.oracle.com/database/availability.

Note: Unlike prior releases, all features and options applicable to your configuration are included on your product CD-ROM and are shipped enabled. *Options must be licensed for use.*

This chapter contains the following sections:

- [Options](#)
- [Features](#)
- [The V\\$OPTION Table](#)

Options

Table 5–1 lists the options available with this release. Options which are new, or whose factoring status has changed, are noted in *bold italic* print.

Table 5–1 Oracle8i Options (Page 1 of 2)

Option	Oracle8i Standard Edition	Oracle8i Enterprise Edition	Oracle8i Personal Edition	Notes
<i>Oracle Advanced Security</i>	N	Y	Y	<i>Oracle Advanced Security provides a single source of integration with network encryption, single sign-on services, smartcard, token and biometric user authentication.</i>
<i>Oracle Change Management Pack</i>	N	Y	Y	<i>Oracle Change Management Pack eliminates errors and loss of data when upgrading databases to support new applications. The pack analyzes the impact and complex dependencies associated with application change and automatically performs database upgrades.</i>
<i>Oracle Diagnostics Pack</i>	N	Y	Y	<i>This is an advanced set of tools used to monitor the state of databases and systems and pinpoint, analyze, and repair any weaknesses.</i>
<i>Oracle interMedia</i>	Y	Y	Y	<i>interMedia enables Oracle8i to manage text, documents, images, audio, and video in an integrated fashion with other enterprise information.</i>
<i>Oracle JServer</i>	Y	N	N	<i>Includes Java stored procedures, methods, and triggers; EJB, CORBA, IIOP, and HTTP support.</i>
<i>Oracle JServer Enterprise Edition</i>	N	Y	Y	<i>Includes Java stored procedures, methods, and triggers; EJB, CORBA, IIOP, and HTTP support.</i>
<i>Oracle Parallel Server</i>	N	Y	N	<i>Oracle Parallel Server allows database workload to be distributed across a cluster, resulting in high scalability. High availability is assured, because even if one node fails, the remaining nodes still can access the database.</i>
<i>Oracle Partitioning</i>	N	Y	Y	<i>Partitioning allows large tables to be broken into smaller pieces that improve manageability, availability, and scalability.</i>

Table 5–1 Oracle8i Options (Page 2 of 2)

Option	Oracle8i Standard Edition	Oracle8i Enterprise Edition	Oracle8i Personal Edition	Notes
Oracle Programmer	Y	Y	Y	It is a family of products consisting of three SQL-style interfaces: precompilers, SQL*Module, and SQLJ; three call level interfaces: Oracle Call Interface (OCI), ODBC, and JDBC; Oracle Objects for OLE; and Object Type Translator and JPub.
Oracle Spatial	N	Y	Y	Oracle Spatial is an integrated set of functions and procedures that enables spatial data to be stored, accessed, and analyzed quickly and efficiently in an Oracle database.
<i>Oracle Standard Management Pack</i>	Y	N	N	<i>A complete set of database tuning, diagnostic, change management, and other manageability technologies for Oracle8i. This is an extra cost option.</i>
<i>Oracle Time Series</i>	N	Y	Y	<i>Manages time-based data and provides functionality to do time-based analysis on the data.</i>
<i>Oracle Tuning Pack</i>	N	Y	Y	<i>Oracle Tuning Pack provides database administrators with expert performance management for the Oracle environment, including SQL tuning and storage optimization.</i>
Oracle Visual Information Retrieval	N	Y	Y	Visual Information Retrieval (VIR) functionality from Virage for advanced image searching and management.
<i>Oracle WebDB</i>	Y	Y	Y	<i>The Oracle WebDB is a complete and cost-effective solution for building, deploying, and proactively monitoring Web database applications and content-driven Web sites.</i>

Features

[Table 5–2](#) lists the new features available with release 8.1.6. Features which were new to Oracle8i, or whose factoring status changed in Oracle8i, are noted in *bold italic* print.

Table 5–2 Oracle8i Release 8.1.5 Feature and Option Availability (Page 1 of 10)

Area	Feature	Oracle8i Standard Edition	Oracle8i Enterprise Edition	Oracle8i Personal Edition	Notes
Data Warehousing and VLDB	<i>Automated parallel query degree</i>	N	Y	Y	<i>Automated parallel query optimization based on system utilization.</i>
Data Warehousing and VLDB	Bitmap indexes	N	Y	Y	Index type commonly used in a data warehouse for columns with small number of distinct values, such as 'Y' or 'N', resulting in dramatic performance gains in a data warehouse application.
Data Warehousing and VLDB	<i>CUBE and ROLLUP</i>	Y	Y	Y	<i>OLAP operators CUBE and ROLLUP to produce sub-totals and crosstab reports easily and efficiently using a single SQL statement.</i>
Data Warehousing and VLDB	<i>Analytic functions</i>	Y	Y	Y	<i>These include rank, moving average, and ratio-to-report.</i>
Data Warehousing and VLDB	<i>Descending indexes</i>	Y	Y	Y	<i>Provides better performance when indexed data needs to sorted in descending order.</i>
Data Warehousing and VLDB	<i>Direct Path Load API</i>	Y	Y	Y	<i>Direct Load API allows complete access to all load functionality via OCI API for creation of high performance load programs.</i>
Data Warehousing and VLDB	<i>Export transportable tablespaces</i>	N	Y	Y	<i>Functionality to quickly move or create a copy of tablespace data. A transportable tablespace can be "plugged in" to all editions of the Oracle8i database, but only the Enterprise Edition can create a transportable tablespace.</i>

Table 5–2 Oracle8i Release 8.1.5 Feature and Option Availability (Page 2 of 10)

Area	Feature	Oracle8i Standard Edition	Oracle8i Enterprise Edition	Oracle8i Personal Edition	Notes
Data Warehousing and VLDB	Function-based indexes	N	Y	Y	Allows indexes to be created on expressions or functions.
Data Warehousing and VLDB	Import transportable tablespaces	Y	Y	Y	Functionality to quickly "plug in" a transportable tablespace. Oracle8i Enterprise Edition must be used to create (export) a transportable tablespace.
Data Warehousing and VLDB	Long operations monitor	Y	Y	Y	Progress of long running database and application operations can be monitored and information can be displayed, such as percent complete.
Data Warehousing and VLDB	Materialized views	N	Y	Y	Includes summary management functionality, join indexes, and summary advisor. Powerful functionality that automatically redirects queries to stored summary tables, improving query performance dramatically.
Data Warehousing and VLDB	Optimizer statistics management	Y	Y	Y	Schema object statistics, used by the cost based optimizer, can be copied and moved to another database.
Data Warehousing and VLDB	Parallel analyze	N	Y	Y	The ANALYZE statement, used to gather statistics on tables, can be run using parallel processing. This feature requires the Partitioning Option.
Data Warehousing and VLDB	Parallel bitmap star query optimization	N	Y	Y	This algorithm utilizes single-table bitmap indexes and an advanced star query join method, resulting in excellent performance while efficiently utilizing space.
Data Warehousing and VLDB	Parallel DML (insert/update/ delete)	N	Y	Y	DML operations are transparently divided across multiple processes. This feature requires the Partitioning Option.

Table 5–2 Oracle8i Release 8.1.5 Feature and Option Availability (Page 3 of 10)

Area	Feature	Oracle8i Standard Edition	Oracle8i Enterprise Edition	Oracle8i Personal Edition	Notes
Data Warehousing and VLDB	Parallel index build	N	Y	Y	Indexes can be created using parallel processes, significantly improving performance of the operation.
Data Warehousing and VLDB	Parallel index scans	N	Y	Y	Queries that use of an index can scan the index in parallel with multiple processes. This requires the Partitioning Option
<i>Data Warehousing and VLDB</i>	<i>Parallel load</i>	Y	Y	Y	<i>Fast load of data accomplished by use of multiple, parallel processes.</i>
Data Warehousing and VLDB	Parallel query	N	Y	Y	Oracle8i transparently distributes query execution across multiple processes, resulting in excellent performance.
<i>Data Warehousing and VLDB</i>	<i>Sample scan</i>	N	Y	Y	<i>The SAMPLE keyword is used to in SQL statements to easily select a percentage of random data within a table; this is useful in data mining applications.</i>
Data Warehousing and VLDB	Star query optimization	Y	Y	Y	This star-query optimization algorithm utilizes B-tree indexes (not bitmap indexes)
Database Features	Advanced Queuing	N	Y	Y	Rules-based, publish and subscribe queuing system that can be used to develop large-scale, message-oriented distributed applications.
<i>Database Features</i>	<i>Database event triggers</i>	Y	Y	Y	<i>Database triggers, such as database startup or shutdown triggers, DDL statement triggers, and logon/logoff triggers are fired when the event occurs.</i>
<i>Database Features</i>	<i>DBMS_REPAIR package</i>	Y	Y	Y	<i>Provides early detection and correction of software and hardware corruption.</i>
<i>Database Features</i>	<i>Drop column</i>	Y	Y	Y	<i>Functionality to remove an unwanted column from a table.</i>

Table 5–2 Oracle8i Release 8.1.5 Feature and Option Availability (Page 4 of 10)

Area	Feature	Oracle8i Standard Edition	Oracle8i Enterprise Edition	Oracle8i Personal Edition	Notes
Database Features	<i>Fine-grained access control</i>	N	Y	Y	<i>Includes functionality to create security policies that implement row-level security.</i>
Database Features	<i>Index coalesce</i>	N	Y	Y	<i>Free space in index leaf blocks can be defragmented while table is online.</i>
Database Features	Index-organized tables	Y	Y	Y	Table where data is stored in a B*-tree index structure for better performance and reduced storage.
Database Features	<i>Indexes on NLS collating sequences</i>	N	Y	Y	<i>Efficient linguistic sorting using NLS sorting through use of function-based indexes (available in Enterprise Edition and Personal Oracle8i).</i>
Database Features	Instead-of triggers	Y	Y	Y	Triggers that execute instead of the DML transaction execution.
Database Features	LOB (large object) support	Y	Y	Y	Datatypes and functionality for storing and manipulating large objects, such as images or other unstructured data.
Database Features	<i>Locally-managed tablespaces</i>	Y	Y	Y	<i>Efficient storage management functionality that improves reliability and reduces fragmentation.</i>
Database Features	<i>LogMiner</i>	Y	Y	Y	<i>LogMiner provides the functionality to analyze transaction log information online to enable auditing of transactions and the ability to "back out" individual transactions.</i>
Database Features	National language support (NLS)	Y	Y	Y	Native language support
Database Features	<i>Objects and extensibility</i>	Y	Y	Y	<i>Object relational functionality, including object types and methods. Also, database extensibility via APIs to database services and functionality.</i>

Table 5–2 Oracle8i Release 8.1.5 Feature and Option Availability (Page 5 of 10)

Area	Feature	Oracle8i Standard Edition	Oracle8i Enterprise Edition	Oracle8i Personal Edition	Notes
Database Features	<i>Online index build</i>	N	Y	Y	<i>Indexes can be built and re-built without locking the table during the operation.</i>
Database Features	Password management	Y	Y	Y	Provides password expiration, complexity, and security policy management.
Database Features	PL/SQL stored procedures, triggers	Y	Y	Y	Provides mechanism to deploy logic directly in the database.
Database Features	<i>PL/SQL Server Pages</i>	Y	Y	Y	<i>Server-side Web pages (in HTML or XML) with embedded PL/SQL scripts, enabling rapid development of dynamic Web pages.</i>
Database Features	<i>Plan Stability</i>	N	Y	Y	<i>This allows execution plans for SQL to be stored so that the plan remains consistent throughout schema changes, database reorganizations, and data volume changes.</i>
Database Features	Reverse key indexes	Y	Y	Y	Indexed columns can be indexed on reversed column values for better index performance.
Database Features	<i>Temporary tables</i>	Y	Y	Y	<i>Temporary tables to allow users to manipulate data for the duration of a transaction or session.</i>
Distributed	Advanced Replication	N	Y	Y	Advanced replication includes basic replication functionality plus multi-master replication, Replication Manager, and parallel propagation to maximize throughput.
<i>Distributed</i>	<i>Basic replication</i>	Y	Y	Y	<i>Oracle8i server fully supports bi-directional replication with automated conflict detection and resolution. Supported configurations include a single updatable master site with multiple updatable or read-only snapshot sites.</i>
Distributed	Distributed queries	Y	Y	Y	Queries can include tables from multiple databases.

Table 5–2 Oracle8i Release 8.1.5 Feature and Option Availability (Page 6 of 10)

Area	Feature	Oracle8i Standard Edition	Oracle8i Enterprise Edition	Oracle8i Personal Edition	Notes
Distributed	Distributed transactions	Y	Y	Y	Transactions can include tables from multiple databases. Includes transparent two phase commit functionality and XA support.
Distributed	Heterogeneous services	Y	Y	Y	Queries and transactions can involve Oracle and non-Oracle databases transparently. This functionality is used by Oracle Transparent Gateways.
Networking	<i>N-tier authentication/authorization</i>	<i>N</i>	<i>Y</i>	<i>Y</i>	<i>N-tier authentication functionality that preserves client identity through all tiers.</i>
Networking	Network access control	N	Y	Y	Provides filtering capabilities between clients and servers and acts as relay for thin JDBC connections that have no access privileges to the database host
Networking	Connection pooling	Y	Y	Y	Connection pooling enables the server to time-out idle sessions and use that connection to service an active session, resulting in more efficient network usage.
Networking	Multi-protocol connectivity	N	Y	Y	Bridges Net8 communities of users that use different network protocols such as SPX/IPX, TCP/IP or LU6.2
Networking	Multiplexing	N	Y	Y	Multiplexing enables multiple network sessions to coexist over a single physical transport, reducing the number of physical connections a server must maintain to support a population of clients.
Networking	Net8	Y	Y	Y	Net8 enables heterogeneous, distributed computing across machines regardless of vendor, operating system, or hardware architecture.

Table 5–2 Oracle8i Release 8.1.5 Feature and Option Availability (Page 7 of 10)

Area	Feature	Oracle8i Standard Edition	Oracle8i Enterprise Edition	Oracle8i Personal Edition	Notes
Networking	Oracle Connection Manager	N	Y	Y	Multiplexing, Multi-protocol connectivity, Network Access Control.
Networking	Oracle Names	Y	Y	Y	Provides fully integrated, multi-protocol network naming service that allows administrators to centrally define service addresses, inter-database links, aliases and client configuration profiles.
Program Interfaces	AppWizard for Visual Studio (NT only)	Y	Y	Y	AppWizard to ease development of Oracle based applications using Visual Studio.
Program Interfaces	Autonomous transactions	Y	Y	Y	Blocks of PL/SQL can commit independently of each other to ease complex transaction programming.
Program Interfaces	COM cartridge (NT only)	Y	Y	Y	COM cartridge allows PL/SQL developers to programmatically manipulate COM objects through the OLE Automation interface. Entry points are exposed for access to the methods defined by the OLE Automation interface for easy application integration.
Program Interfaces	JDBC drivers	Y	Y	Y	JDBC access to Oracle8i: Java call-level interface to SQL that is standards compliant and supports Oracle-specific features.
Program Interfaces	Microsoft Transaction Server Integration (NT only)	Y	Y	Y	Full, native integration with Microsoft Transaction Server (MTS) to allow developers to deploy COM-based applications using MTS against Oracle.
Program Interfaces	Objects for OLE	Y	Y	Y	Oracle Objects for OLE is a COM-based database connectivity tool that provides seamless and optimized access to Oracle databases.

Table 5–2 Oracle8i Release 8.1.5 Feature and Option Availability (Page 8 of 10)

Area	Feature	Oracle8i Standard Edition	Oracle8i Enterprise Edition	Oracle8i Personal Edition	Notes
Program Interfaces	ODBC driver	Y	Y	Y	ODBC driver for access to the Oracle database.
Program Interfaces	Oracle Call Interface (OCI)	Y	Y	Y	Lowest level API interface to the Oracle database.
Program Interfaces	Pro*C	Y	Y	Y	Allows C programs to easily use SQL to access data. Requires Oracle Programmer
Program Interfaces	SQLJ	Y	Y	Y	Allows embedded SQL statements in Java for concise and easy access to Oracle data. Requires Oracle Programmer.
System Management	Automated standby database	N	Y	Y	Multiple remote standby databases can easily be configured and automatically managed by the server, including copying and applying the transaction log files at the standby sites.
System Management	Readable standby database	N	Y	Y	Databases, including a standby database, can be opened as "read-only", where no DML operations are allowed.
System Management	Database resource management	N	Y	Y	Resources, such as CPU and degree of parallelism, can be allocated to groups of users such that a prioritization of tasks can be managed easily.
System Management	Duplexed backup sets	N	Y	Y	Backup sets can be written to multiple devices in parallel.
System Management	Oracle DBA Management Pack	Y	Y	Y	Set of management tools bundled with Oracle Enterprise Manager.
System Management	Fast-start fault recovery	N	Y	Y	Fast-start fault recovery provides fast and predictable recovery from system failures.

Table 5–2 Oracle8i Release 8.1.5 Feature and Option Availability (Page 9 of 10)

Area	Feature	Oracle8i Standard Edition	Oracle8i Enterprise Edition	Oracle8i Personal Edition	Notes
System Management	Incremental backup and recovery	N	Y	Y	This functionality allows a backup process to be set up such that only changed blocks are written to the backup file.
System Management	Legato Storage Manager	Y	Y	Y	Integrated functionality from Legato to allow backups directly to tapes.
System Management	Online backup and recovery	Y	Y	Y	Backups can be run while the database is online. Recovery operations can also be run while the database is running.
<i>System Management</i>	<i>Oracle DBA Management Pack</i>	Y	Y	Y	<i>This Oracle Enterprise Manager management pack includes Schema Manager and other tools.</i>
System Management	Oracle Enterprise Manager	Y	Y	Y	Single, integrated console that allows easy management and monitoring of the Oracle and non-Oracle database environment.
System Management	Oracle Fail Safe for Oracle8i on NT	Y	Y	Y	Oracle Fail Safe makes it easy to deploy highly available single-instance Oracle8i database solutions on Microsoft Windows NT clusters.
System Management	Parallel backup and recovery	N	Y	Y	Both backup and recovery can be run using parallel processing.
System Management	Point-in-time tablespace recovery	N	Y	Y	Allows a tablespace to be recovered up to a specified point-in-time after a failure or inadvertent transaction execution.
System Management	Recovery Manager	Y	Y	Y	Easy to use, wizards-based functionality to setup and manage the entire backup and recovery process.

Table 5–2 Oracle8i Release 8.1.5 Feature and Option Availability (Page 10 of 10)

Area	Feature	Oracle8i Standard Edition	Oracle8i Enterprise Edition	Oracle8i Personal Edition	Notes
System Management	Server managed backup and recovery	Y	Y	Y	The database server manages the backup process, such as scheduling of backups, as well as the recovery process, such as applying the correct backup file when recovery is needed.
System Management	Transparent Application Failover	N	Y	Y	Oracle8i Transparent Application Failover (TAF) can completely mask many failures from end-users, by preserving the state of their application and resuming any work that had been in progress at the time of failure.

The V\$OPTION Table

To check the values for different options on your database, query the V\$OPTION table:

```
SQL> SELECT * FROM V$OPTION;
```

The response, of course, varies depending upon the specific installation. A sample result might be:

PARAMETER	VALUE
Partitioning	TRUE
Objects	TRUE
Parallel Server	FALSE
Advanced replication	TRUE
Bit-mapped indexes	TRUE
Connection multiplexing	TRUE
Connection pooling	TRUE
Database queuing	TRUE
Incremental backup and recovery	TRUE

PARAMETER	VALUE
Instead-of triggers	TRUE
Parallel backup and recovery	TRUE
Parallel execution	TRUE
Parallel load	TRUE
Point-in-time tablespace recovery	TRUE
Fine-grained access control	TRUE
N-Tier authentication/authorization	TRUE
Function-based indexes	TRUE
Plan Stability	TRUE
Online index segment creation	TRUE
Index segment coalesce	TRUE
Managed Standby	TRUE
Materialized view rewrite	TRUE
Materialized view warehouse refresh	TRUE
Database Resource Manager	TRUE
Spatial	TRUE
Visual Information Retrieval	TRUE
Export transportable tablespaces	TRUE
Transparent Application Failover	TRUE
Fast Start recovery	TRUE

Oracle8*i* Documentation

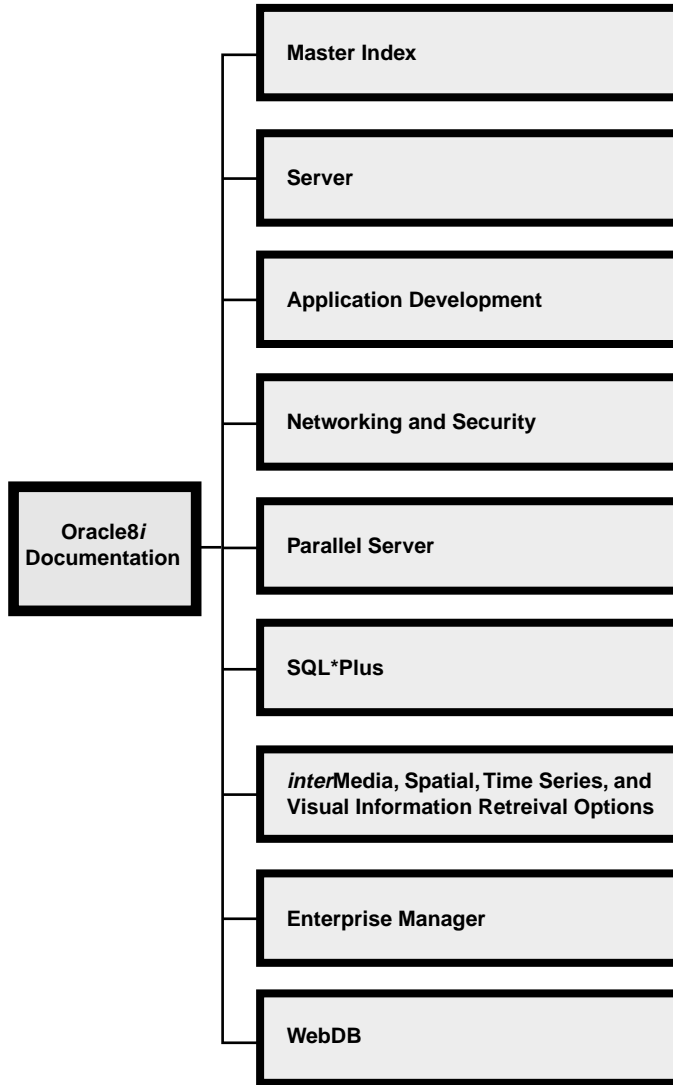
This chapter lists the generic (non-operating system-specific) books that document Oracle8*i*, release 8.1.6. Oracle8*i* documentation is available online, and some books are available in printed format.

This chapter contains the following sections:

- [Online Documentation Structure](#)
- [Documentation Titles](#)

Online Documentation Structure

All books are available in electronic format in the *Oracle8i, Release 8.1.6, Online Generic Documentation* CD-ROM. The following figure shows how the books are organized on this CD.



Documentation Titles

The following tables list the book titles in the generic documentation areas.

Note: This is the list of books and part numbers that were available at release time. The documentation CD may be slightly different than what is listed below. You can also view Oracle documentation at the Oracle Technology Network Web site: <http://technet.oracle.com/>.

Oracle8i Master Index

The Master Index includes indexed entries for all the books included in the [Oracle8i Server](#), [Application Development](#), [Networking and Security](#), and [Parallel Server](#) subsections listed below.

Table 6–1 Oracle8i Master Index Documentation

Oracle8i Master Index	Part Number
<i>Oracle8i Generic Documentation Master Index</i>	A76991-01

Oracle8i Server

Table 6–2 Oracle8i Server Documentation

Oracle8i Server Documentation	Part Number
<i>Getting to Know Oracle8i</i>	A76962-01
<i>Oracle8i Concepts</i>	A76965-01
Volume 1	A76963-01
Volume 2	A76964-01
<i>Oracle8i Error Messages</i>	A76999-01
Volume 1	A76998-01
Volume 2	A76996-01
Volume 3	A76997-01
<i>Oracle8i National Language Support Guide</i>	A76966-01
<i>Oracle8i Replication</i>	A76959-01
<i>Oracle8i Replication Management API Reference</i>	A76958-01

Table 6–2 Oracle8i Server Documentation

Oracle8i Server Documentation	Part Number
<i>Oracle8i Reference</i>	A76961-01
<i>Oracle8i SQL Reference</i>	A76989-01
Volume 1	A76967-01
Volume 2	A76988-01
<i>Legato Storage Manager Administrator's Guide</i>	TBD
<i>Oracle8i Administrator's Guide</i>	A76956-01
<i>Oracle8i Backup and Recovery Guide</i>	A76993-01
<i>Oracle8i Recovery Manager User's Guide and Reference</i>	A76990-01
<i>Oracle8i Standby Database Concepts and Administration</i>	A76995-01
<i>Oracle8i Distributed Database Systems</i>	A76960-01
<i>Oracle8i Migration</i>	A76957-01
<i>Oracle8i Designing and Tuning for Performance</i>	A76992-01
<i>Oracle8i Utilities</i>	A76955-01
<i>Oracle8i Data Warehousing Guide</i>	A76994-01

Application Development

Table 6–3 Application Development Documentation

Application Development Documentation	Part Number
<i>PL/SQL User's Guide and Reference</i>	A77069-01
<i>Oracle Call Interface Programmer's Guide</i>	A76975-01
Volume 1	A76977-01
Volume 2	A76978-01
<i>Oracle8i Application Developer's Guide - Fundamentals</i>	A76939-01
<i>Oracle8i Application Developer's Guide - Large Objects (LOBs)</i>	A76940-01
Volume 1	A76949-01
Volume 2	A76950-01
<i>Oracle8i Application Developer's Guide - Advanced Queuing</i>	A76938-01
Volume 1	A76947-01
Volume 2	A76948-01
<i>Oracle8i Application Developer's Guide - Object-Relational Features</i>	A76976-01
<i>Oracle8i Supplied PL/SQL Packages Reference</i>	A76936-01
Volume 1	A76946-01
Volume 2	A76945-01
<i>Oracle8i Supplied Java Packages Reference</i>	A76935-01
<i>Pro*C/C++ Precompiler Programmer's Guide</i>	A76942-01
Volume 1	A76943-01
Volume 2	A76944-01
<i>Pro*COBOL Precompiler Programmer's Guide</i>	A76951-01
<i>Oracle8i Data Cartridge Developer's Guide</i>	A76937-01
<i>* Programmer's Guide to the Oracle Precompilers (release 7.3.4)</i>	A42525-01
<i>* Pro*Fortran Supplement to the Oracle Precompilers Guide (release 7.3.4)</i>	A42523-01
<i>* SQL*Module for Ada Programmer's Guide (release 8.0.4)</i>	A58231-01

*** Note:** The last three books have not been updated. They are provided for software which is in maintenance mode (bug fixes only).

Online-Only Books

Oracle8i Application Developer's Guide - Large Objects (LOBs) Using PL/SQL

Oracle8i Application Developer's Guide - Large Objects (LOBs) Using OCI

*Oracle8i Application Developer's Guide - Large Objects (LOBs) Using Pro*C*

*Oracle8i Application Developer's Guide - Large Objects (LOBs) Using Pro*COBOL*

Oracle8i Application Developer's Guide - Large Objects (LOBs) Using VB (OO4O)

Oracle8i Application Developer's Guide - Large Objects (LOBs) Using Java

Oracle8i Application Developer's Guide - Advanced Queuing Using PL/SQL

Oracle8i Application Developer's Guide - Advanced Queuing Using VB (OO4O)

Oracle8i Application Developer's Guide - Advanced Queuing Using Java

Networking and Security

Table 6–4 Networking and Security Documentation

Networking and Security Documentation	Part Number
<i>Net8 Administrator's Guide</i>	A76933-01
<i>Oracle Internet Directory Administrator's Guide</i>	A76931-01
<i>Oracle Advanced Security Administrator's Guide</i>	A76932-01

Parallel Server

Table 6–5 Parallel Server Documentation

Parallel Server Documentation	Part Number
<i>Oracle8i Parallel Server Documentation Online Roadmap</i>	online only
<i>Oracle8i Parallel Server Administration, Deployment, and Performance</i>	A76970-01
<i>Oracle8i Parallel Server Setup and Configuration Guide</i>	A76934-01

Table 6–5 Parallel Server Documentation

Parallel Server Documentation	Part Number
<i>Oracle8i Parallel Server Concepts</i>	A76968-01

SQL*Plus

Table 6–6 SQL*Plus Documentation

SQL*Plus Documentation	Part Number
<i>SQL*Plus User's Guide and Reference</i>	A75664-01
<i>SQL*Plus Quick Reference</i>	A75665-01

interMedia, Spatial, Time Series, and Visual Image Retrieval

Table 6–7 *interMedia*, Spatial, Time Series, and Visual Image Retrieval

<i>interMedia</i>, Spatial, Time Series, and Visual Image Retrieval	Part Number
<i>Oracle8i interMedia Text Migration</i>	A77061-01
<i>Oracle8i interMedia Text Reference</i>	A77063-01
<i>Oracle8i interMedia Audio, Image, and Video User's Guide and Reference</i> (release 8.1.5)	A67299-01
<i>Oracle8i interMedia Audio, Image, and Video Java Client User's Guide and Reference</i> (release 8.1.5)	A67296-01
<i>Oracle8i interMedia Locator User's Guide and Reference</i> (release 8.1.5)	A67298-01
<i>Oracle Spatial User's Guide and Reference</i>	A77132-01
<i>Oracle8i Time Series User's Guide</i> (release 8.1.5)	A67294 -01
<i>Oracle8i Visual Information Retrieval User's Guide and Reference</i> (release 8.1.5)	A67293-01
<i>Oracle8i Visual Information Retrieval Java Client User's Guide and Reference</i> (release 8.1.5)	A67300-01

Oracle Enterprise Manager

Table 6–8 Oracle Enterprise Manager

Oracle Enterprise Manager	Part Number
<i>Oracle Enterprise Manager Installation</i>	A75681-01
<i>Oracle Enterprise Manager Readme</i>	A75682-01
<i>Oracle Enterprise Manager Administrator's Guide</i>	A75683-01
<i>Oracle Enterprise Manager Messages Manual</i>	A75690-01
<i>Oracle SNMP Support Reference Guide</i>	A75687-01
<i>Oracle Enterprise Manager Configuration Guide</i>	A75685-01
<i>Oracle Enterprise Manager Concepts Guide</i>	A75684-01
<i>Oracle Intelligent Agent User's Guide</i>	A75686-01

WebDB

Online-Only Books

Release Notes

Oracle WebDB Installation

Oracle WebDB Tutorial

Oracle WebDB Creating and Managing Components - Task Help

Oracle WebDB Creating and Managing Components - Field-Level Help

Oracle WebDB Creating and Managing Sites - Task Help

Oracle WebDB Creating and Managing Sites - Field-Level Help

Deprecated and Desupported Features

This and prior releases of the Oracle database server have made several features or certain functionality no longer advisable for use. Some of this functionality has been, or will be, desupported.

This chapter contains the following sections:

- [Strings of Zero Length Are Not Equivalent To a NULL](#)
- [The SELECT Privilege](#)
- [Date Format Strings Are Stricter](#)
- [SERIALIZABLE=TRUE Is No Longer Supported](#)
- [Non-Deferred Linking](#)
- [Single-Task Linking](#)
- [CONNECT INTERNAL](#)
- [Partition Views](#)
- [V6 Compatibility Behavior](#)
- [Use of "THE\(subquery\)" Expression](#)
- [Server Manager Desupport](#)
- [The SGADEF File](#)
- [LONG Column Support](#)
- [The Oracle Security Server and Cryptographic Toolkit](#)
- [Dynamic Views Used for Monitoring Parallel Execution Performance](#)

Strings of Zero Length Are Not Equivalent To a NULL

A string of zero length ('') is not equivalent to a NULL.

According to the ANSI SQL 1992 Transitional standard, a zero-length or empty string is not the same as NULL. The Oracle database server *may comply fully with this aspect of the standard in the future*. Therefore, it is recommended that applications ensure that empty strings values and NULL are not treated equivalently.

The SELECT Privilege

The SELECT privilege *may be required* on tables that users update. Always grant the SELECT privilege to a user or role if you grant the UPDATE or DELETE privileges on the table.

In conformance with SQL92, Oracle releases starting with release 7.1, allow the security administrator to require that users have SELECT privilege on a table when executing an UPDATE or DELETE statement that references table column values in a WHERE or SET clause.

For more information, refer to the discussion of the SQL92_SECURITY initialization parameter in *Oracle8i Reference*.

Date Format Strings Are Stricter

In Oracle7, a space or punctuation character in the format string caused the corresponding character in the date string to be discarded. This caused incorrect dates to be entered into the database since alphanumeric characters were thrown out. *Starting in Oracle8, an error occurs* if an alphanumeric character is found in the date string when a punctuation character or space is found in the format string.

Example:

```
TO_CHAR(TO_DATE('0297', 'MM/YY'), 'MM/YY')
```

Oracle7 result: 02/07

Oracle8 result: ORA-1861

SERIALIZABLE=TRUE Is No Longer Supported

The initialization parameter SERIALIZABLE=TRUE is *no longer supported in Oracle8 and beyond*. The default behavior henceforth is as if SERIALIZABLE was set to

FALSE. Use the `SET TRANSACTION ISOLATION LEVEL SERIALIZABLE` statement to achieve similar transaction isolation behavior.

Non-Deferred Linking

Application developers are cautioned that Oracle *plans to desupport non-deferred mode linking beginning with the release of Oracle9* (it continues to be supported with all the releases of Oracle8). Recognizing these plans, application developers should no longer use non-deferred mode linking in developing new applications. Currently Oracle supports two linking modes:

1. Non-deferred linking

The Oracle6 OCI (client) only supported non-deferred linking which meant that for each SQL statement, a parse, a bind and a define call were each executed separately with individual round trips between the client and the server. This significantly increased network traffic between the client and the server and reduced both the performance and scalability of OCI applications.

2. Deferred linking

Starting with the Oracle7 OCI, both non-deferred linking and deferred linking are supported. Deferred mode linking essentially defers the bind and define steps until the statement executes, so it automatically bundles and defers the bind and define calls to execution time. Further, when the application is linked with deferred mode and a special parsing call is used (the `OPARSE` call with the `DEFFLG` set to a non-zero value), even the parse call can be deferred to execution time. Note that deferred mode linking does not depend on the specific OCI calls that the application uses, only on the link option that is selected.

Deferred mode linking therefore significantly reduces the number of round trips between the client and the server and as a result improves the performance and scalability of OCI applications. The default behavior of Oracle7 OCI connected to the Oracle7 database server is deferred mode linking. However, Oracle7 OCI also supports non-deferred linking by setting specific link time options.

Further, starting with Oracle8 OCI, there are two types of calls. First, all the Oracle7 OCI calls are supported with Oracle8 OCI; i.e., they work with an Oracle8 OCI client by relinking the Oracle8 OCI libraries. Second, there are additional Oracle8-specific OCI calls. The default mode with the first type of calls continues to be deferred mode linking; however, non-deferred mode linking is supported for these calls through all releases of Oracle8 by setting link time options. However,

Oracle8-specific calls use a different paradigm and as a result non-deferred mode linking is not necessary.

The various combinations of client side libraries and server with which non-deferred linking is currently supported are summarized below:

Table 7-1 Non-deferred Linking: Server and Client-side Libraries

Server	Client-side Libraries				
	OCI release 6	OCI release 7	OCI release8 (release 7 calls)	OCI release 8 (release 8 calls)	OCI release 9
Oracle9	Not supported	Default: deferred Non-deferred supported	Default: deferred Non-deferred supported	Not supported	Not supported
Oracle8	Not supported	Default: deferred Non-deferred supported	Default: deferred Non-deferred supported	Not supported	Not supported
Oracle7	Non-deferred mode only	Default: deferred Non-deferred supported	Default: deferred Non-deferred supported	Not supported	Not supported
Oracle6	Non-deferred mode only	Default: deferred Non-deferred supported	Default: deferred Non-deferred supported	Not supported	Not supported

Oracle continues to support deferred-mode linking with all the releases of Oracle8 (all 8.* releases). This has implications as discussed in the following.

Applications Using Oracle6 OCI Libraries

Since the Oracle6 OCI library is not supported against the Oracle8 database, applications using the Oracle6 library cannot be run against an Oracle8 database.

Applications Using Oracle7 OCI Libraries

Applications using Oracle7 OCI libraries can run in two configurations against an Oracle8 database:

1. They can be run with Oracle7 OCI libraries against an Oracle8 database in non-deferred mode provided link time options are set appropriately.
2. They can also be relinked with the Oracle8 OCI libraries and run in non-deferred mode provided link time options are set appropriately.

Oracle supports the first configuration through all the releases of Oracle8. However, the second configuration will not be supported in Oracle9. Therefore, applications that require non-deferred linking will not be able to upgrade to Oracle9 client-side libraries.

Applications Using Oracle8 OCI Libraries

Applications using Oracle8 specific OCI calls, such as those used to access Oracle8's object types, do not need to use non-deferred mode linking since the Oracle8 OCI uses a different paradigm. Applications using only Oracle7 OCI calls can use non-deferred mode linking, but only through Oracle Release 8.1 (Oracle8i).

Single-Task Linking

Single-task linking is a feature used by a limited number of Oracle's customers primarily on the OpenVMS platform. Application developers are cautioned that Oracle continues to support single-task linking with all the releases of Oracle8 (all 8.* releases) but *will desupport it beginning with the first release after Oracle8i*.

With single-task linking, Oracle supports two configurations to link Oracle products and user-written applications against the Oracle database:

1. **Single-task linking:** In this case, applications are directly linked against the Oracle shareable image making single-task connection to Oracle.
2. **Two-task linking:** In this case, applications linked in a standalone configuration can only connect to Oracle using SQL*Net's two task drivers such as SQL*Net DECnet or SQL*Net VMS Mailbox on the OpenVMS platform. This is the typical configuration used in the large majority of client-server applications. With two task linking, applications and tools connect with the Oracle7 database through a programmatic interface that creates a shadow process for each user process. This shadow process runs a copy of the Oracle shareable image on behalf of the user process using SQL*Net protocols to communicate between the user and shadow processes. Therefore, with this interface, user routines that invoke the Oracle7 database server functions run as one process or task, and the Oracle7 routines that execute these functions run as the second task.

Application developers who would like to use single-task linking to run their applications will not be able to do so against the first server release after Oracle8.

CONNECT INTERNAL

CONNECT INTERNAL is currently supported for backwards compatibility only.

CONNECT INTERNAL *will be completely desupported in Oracle8i, release 8.2*. If you have not done so already, then you should plan to migrate your applications to use other connection syntax. See *Oracle8i Administrator's Guide* for more details.

Partition Views

In Oracle8i, *partitioned tables are strongly recommended in preference to partition views*. If you must use partition views, then see *Oracle8i Concepts*.

V6 Compatibility Behavior

With Oracle7, Oracle offered a Version 6 [V6] compatibility flag that allowed application developers developing Oracle7 applications to emulate Oracle6 behavior. With the release of Oracle8i (but starting in Oracle8, release 8.0), users are again cautioned that the *Version 6 compatibility flag has been desupported in all of the Oracle8 products* including PL/SQL8, all the Oracle Precompilers, the Oracle8 Oracle Call Interface, SQL*Module, and SQL*Plus. The desupport of the V6 compatibility flag is consistent with Oracle's policy of supporting backwards compatibility and behavior from one version release upgrade to another; i.e., from Oracle6 to Oracle7 but not for more than one version release upgrade.

Users who do not absolutely need to maintain V6 behavior are encouraged to upgrade their Oracle7 clients to Oracle8i. Users who absolutely need to continue to emulate V6 behavior for certain applications need to maintain one \$ORACLE_HOME with an Oracle7 client for those applications. They can create a separate \$ORACLE_HOME with an Oracle8i client for those applications for which they do not need V6 behavior.

Specifically, the V6 compatibility flag emulated the following aspects of OracleV6 behavior with Oracle7:

- String literals are fixed length in Oracle7 but are treated as variable length with the V6 flag.
- PL/SQL local CHAR variables are fixed length in Oracle7 but are treated as variable length with the V6 flag.
- Return value of SQL functions (e.g., USER) are fixed length characters in Oracle7 but are treated as variable length characters with the V6 flag.

- Select/Fetch of a NULL with no indicator raises an ORA-1405 error with Oracle7, but it returns no error with the V6 flag.
- SQL group function is called at FETCH time with Oracle7, but it is called at query execution time with the V6 flag.
- Describe of a fixed length string returns Type=96 with Oracle7 but returns Type=1 with the V6 flag.

All of the above are desupported with the desupport of the V6 Compatibility Flag in all releases above Oracle8, release 8.0.

Use of "THE(subquery)" Expression

The current syntax for the *table_expression_clause* of a SELECT, INSERT, UPDATE, or DELETE statement includes a *table_collection_expression*. This expression is recognized by the TABLE() keyword (i.e., TABLE(*collection_expression*)), and it informs Oracle that the collection value expression should be treated as a table. Further, *collection_expression* may be a subquery that selects a nested table column from a table or view.

Prior to Oracle8, release 8.0, *table_collection_expression* was expressed as "THE (subquery)". *That usage is now deprecated.*

For more information, see *Oracle8i SQL Reference*.

Server Manager Desupport

The Server Manager *will be desupported in a future release* of Oracle. You should begin to migrate your Server Manager SQL scripts to SQL*Plus scripts as soon as possible.

The SGADEF File

The contents of this file have been deleted, but the file remains because its existence is required by Oracle. *Do not delete this file.* It will be obsoleted in a future release.

LONG Column Support

LONG column support will be discontinued *in a future Oracle release*. You are advised to migrate LONG data into LOB columns.

The Oracle Security Server and Cryptographic Toolkit

The Oracle Security Server and Cryptographic Toolkit products do not exist in Oracle8i, but their functionality has been integrated into the standards based products: Oracle Advanced Security option and Oracle Internet Directory. *Users must migrate* their applications to use these Oracle8i components.

Dynamic Views Used for Monitoring Parallel Execution Performance

The following dynamic views, used for monitoring parallel execution performance, *will be removed, renamed, or replaced in the next major Oracle release (Release 9.0).*

View	Description	Action
V\$PQ_SESSSTAT	Provides parallel execution statistics for each session	Removed
V\$PQ_SYSSTAT	Provides parallel execution statistics for the system	Replaced by new view: V\$PX_PROCESS_SYSSTAT
V\$PQ_SLAVE	This view tallies the current and total CPU time and number of messages sent and received per parallel server process	Replaced by new view: V\$PX_PROCESS
V\$PQ_TQSTAT	Provides a detailed report of message traffic at the table queue level	Renamed to V\$PX_TQSTAT

Glossary

application context

A secure data cache for storing information used by fine-grained access control to make access control decisions. Used in conjunction with fine-grained access control.

autonomous transactions

Transactions within an autonomous PL/SQL block, which is a block that has a transaction scope independent of the transaction scope of the calling PL/SQL block. Autonomous transactions are not nested transactions, since they do not share resources with the calling transaction and committed changes are immediately visible to other transactions regardless of whether the calling PL/SQL block commits or rolls back.

cache fusion

A new "diskless" ping architecture, used in the Oracle Parallel Server, that provides copies of blocks directly from the holding instance's memory cache to the requesting instance's memory cache.

composite partitioning

Partitions data using the range method and, within each partition, subpartitions it using the hash method (see hash partitioning). It provides the manageability and availability benefits of range partitioning with the data distribution advantages of hash partitioning.

CUBE

One of the Oracle8i operators that can be used in a GROUP BY clause for dimensional analysis. Similar to the ROLLUP operator (see ROLLUP), but also

includes cross-tabulations--i.e., it uses *every* possible combination of the columns or expressions in the `GROUP BY` clause to produce super-aggregate rows.

data sampling

A process used by the `ANALYZE` statement or the `DBMS_STATS` package to *estimate* statistics for a database object (table, index, or cluster). It involves analyzing a specified (or default) percent (or number) of random rows or blocks to estimate statistics.

Database Configuration Assistant

A Java-based utility that enables the creation, modification, or deletion of an Oracle database. It can be invoked as a standalone Java application from the Oracle Universal Installer or as an applet from the Java-based Oracle Enterprise Manager.

DBA Studio

DBA Studio is a new component in the DBA Management pack which combines the functionality of Schema, Security, Storage, and Instance Managers into one tool. This improves the workflow between the various day-to-day administrative tasks performed by the DBA, and allows more than one database to be managed in the Navigator tree when not connected through the OMS. This provides DBAs with an additional tool that offers scalability from the smallest to the largest Oracle site.

dimension

Used in structuring materialized views (see materialized view), a dimension defines hierarchical (parent/child) relationships between pairs of columns or column sets. A hierarchical relationship is a functional dependency from one level of a hierarchy to a more abstract level in the hierarchy. Each value at the child level is associated with one and only one parent-level value.

DISABLE VALIDATE constraint state

A database state that saves space by requiring no index on a unique or primary key, and guarantees the validity of existing data in the table. This state also enables efficient loading of data from a non-partitioned table into a partitioned table.

domain index

A user-defined index for a complex datatype. The application software is responsible for defining the index structure, maintaining the index content during load and update operations, and searching the index during query processing. The index structure itself can be stored either in the Oracle database as tables or externally as files.

enterprise roles

Stored in the Oracle Internet Directory, they enable centralized authorization of users by defining a users privileges in the enterprise. If a user changes jobs, his enterprise role assignment can be changed, altering his privileges in multiple databases throughout the enterprise. Also, an administrator can add capabilities to enterprise roles (granted to multiple users) without having to update the authorizations of each user independently.

event triggers

Refers to triggers that can now be specified on database and schema events. These are:

- DML events (DELETE, INSERT, UPDATE)
- DDL Events (CREATE, ALTER, DROP)
- Database events (SERVERERROR, LOGON, LOGOFF, STARTUP, SHUTDOWN)

extensibility

Capability of being extended. For example, the Oracle8i extensibility framework allows partners to easily provide their own implementations of RDBMS services and register them with the server.

Fast-Start Fault Recovery

A cluster of features that provide high availability by enabling the database to recover quickly from instance failures. Fast-Start Parallel Rollback rolls back transactions in parallel. Fast-Start On-Demand Rollback recovers individual blocks in large transactions that require user access. Fast-Start Checkpointing allows users to set an upper limit for the number of blocks requiring recovery. Replaces incremental checkpoint feature of Oracle8.

fine-grained access control

A feature that supports user-defined security policy for row-level access control. Rather than embedding security rules in views, the server dynamically appends a WHERE condition to a query that allows security rules to be acquired when a table or view is referenced in a DML statement. This enforces security 100%, across all applications, since it is in the database that the security is enforced.

function-based index

The value of a function or expression is precomputed and stored in the index. Performance is dramatically improved when a computationally intensive expression is used to build the index. When an application uses such an expression

in its where clause, much faster access is provided to the data. This greatly enhances the use of parallel I/O.

hash partitioning

Partitions data according to a hash function. Hash partitioning provides a very simple way to break data into evenly sized containers to be spread across multiple I/O devices, or even multiple machines in a shared-nothing cluster. This greatly enhances the use of parallel I/O.

interMedia

Oracle8i *interMedia*. This is Oracle's offering for managing multimedia data in an integrated fashion with other enterprise data. It is specifically aimed to support the data requirements of the Internet. See multimedia data.

Internet Computing

Using the infrastructure of the Internet to bring business applications to the enterprise. These applications are deployed on scalable, professionally managed servers, and are accessible across the organization in a low cost computing environment. Internet computing makes use of such components as browsers, Web servers, and internet standards such as SSL and LDAP.

Internet File System (iFS)

A Java application that runs within Oracle8i and provides universal access to any data in the database. The *iFS* appears as if it were just another volume on the network. Relational data can appear as files; so too can hybrid documents that combine relational and non-relational data. Whether the user accesses the contents of *iFS* through Windows Explorer, a Web browser, an FTP client, or an e-mail client, the files appear the same.

Java VM

A complete, JDK 1.1-compliant Java execution environment. The Java VM runs in the same process space and address space as the database kernel, sharing its memory heaps and directly accessing its relational data. This design optimizes memory use, increases throughput, and delivers an open, highly available, secure, and manageable Java server.

JPublisher

A utility that translates user-defined types to Java wrapper classes. It is similar to the Object Type Translator (OTT) utility used in C/C++ environments.

JServer

Oracle's Java Virtual Machine (Java VM). It includes Java stored procedures, methods, and triggers. It also provides EJB, CORBA, and HTTP support.

JServer Accelerator

A native code compiler that speeds up the execution of Java code by eliminating interpreter overhead. It translates standard Java binaries into C programs that are processed by a platform-dependent C compiler into native libraries that the Oracle Java VM can load dynamically.

linguistic-sorted index

An NLS-supported, function-based index that uses a linguistic sort order for optimized performance. For example, you can create an NLS index on a table that returns an extended Spanish collation order.

locally managed tablespaces

A mechanism that manages space within a tablespace by using bitmaps to track all extent information in the tablespace itself. This optimizes space usage and consumption in the tablespace, minimizing fragmentation so that data does not need to be re-organized. Bitmaps require no dictionary access to allocate or update extents for the tablespace.

logical ROWID

Primary-key-based logical identifiers that expand ROWID functionality to index-organized tables (IOTS), whose rows do not have permanent physical addresses. Oracle uses these logical ROWIDs for the construction of secondary indexes on index-organized tables.

LogMiner

A utility that allows you to analyze online or archived redo logs according to user-specified criteria. For example, you can track changes to the database made by a specified user.

materialized view

A stored summary or snapshot (remote materialized view used for replication) containing precomputed results. A CREATE MATERIALIZED VIEW statement containing a subquery, typically a join or a data aggregation, is used to create the materialized view. This provides much faster access to summary data. Local materialized views can be used for query rewrite.

multimedia data

Includes text, audio data, video data, images, and graphic objects. Oracle8i *interMedia* enables Oracle8i to manage this data in an integrated fashion with other enterprise data.

Oracle Advanced Security

A renaming and expansion of the Oracle Advanced Networking Option.

Oracle Internet Directory (OID)

An LDAPv3-compliant, hierarchical data repository for storing enterprise user information, including enterprise roles and Oracle Wallets. It combines a native implementation of the Internet Engineering Task Force's (IETF) LDAP v3 standard with an Oracle8i back-end datastore.

Oracle Universal Installer

A Java-based installation utility that has the same interface across all platforms. Oracle Universal Installer interprets the staging area prepared by Oracle Software Packager and performs the appropriate installation actions.

Oracle wallet

An encrypted data structure that contains a user's credentials: a private key, a user certificate, and a list of trustpoints (certificate authorities that the user trusts). Users can store wallets in a directory and thus access their credentials from anywhere in the organization.

plan stability

A method for stabilizing execution plans used by the optimizer that is independent of changes to the system configuration, parameters, statistics, or even the optimizer itself. You can use initialization parameters or SQL statements to create stored outlines containing a set of attributes used by the optimizer. Useful for optimizing applications for multiple environments.

query rewrite

With the Oracle8i summary management feature, this is an automatic rewrite of a query by the Oracle database server to use the summary data, rather than retrieving data from detail tables by doing expensive joins and aggregate operations. This rewrite facility is totally transparent to the application, which is not aware of the existence of the materialized view.

read-only standby database

A standby database that is opened in read-only mode, making it available for queries. You can move a standby database from recovery mode to read-only mode and back again.

ROLLUP

One of the Oracle8i operators that can be used in a `GROUP BY` clause for dimensional analysis. A commonly performed operation in a data warehousing environment, and in producing materialized views, is to show totals at a coarse granularity and then to show subtotals at successively finer levels. `ROLLUP` produces these results.

rule-based publish/subscribe

An Advanced Queuing model in which the sending applications anonymously publish messages and the receiving applications independently and autonomously subscribe to these messages (using SQL-based rules). This feature allows messages to be filtered by content.

secondary index

An index on an index-organized table that enables performing operations based on a column that isn't part of the primary key. Oracle constructs secondary indexes on index-organized tables using logical row identifiers (logical `ROWIDs`) that are based on the table's primary key.

security policy

Defines who is allowed access and to which data. Various means are available within Oracle8i for enforcing security policies including, but not limited to, privileges, roles, the virtual private database feature, and network authentication and authorization.

snapshot refresh group templates

A feature that allows a single refresh group to contain up to 400 snapshots (or remote materialized views) and also reduces the number of round trips required to refresh snapshots in a group.

SQLJ

Oracle's precompiler for Java source code. It enables a user to embed SQL statements in the Java source. The precompiler takes the source code and produces as output new Java source code, where all of the SQL statements are translated into Java class definitions that implement these statements.

summary advisor

A `DBMS_OLAP` package to help users design and evaluate materialized views for query rewrites.

temporary LOBs

`LOBs` (Large object datatypes) that do not exist permanently in the database, but exist mainly for the purpose of performing transformations on `LOB` data. Many applications have a need for temporary `LOBs`, which act like local variables.

temporary tables

A table with session-specific or transaction-specific data that is empty when a session or transaction begins and discarded when it ends. Temporary tables are useful for saving intermediate results that can be joined back into another table.

transportable tablespaces

A feature that allows a user to move a subset of an Oracle database into another Oracle database. It is also possible to clone a tablespace in one database and plug it into another, thereby copying the tablespace between databases.

user-defined operator

A SQL operator whose definition and implementation is provided by the user. Part of the extensibility architecture, the Oracle database server allows these user-defined operators to be used in SQL statements in the same manner as any of the predefined operators provided by Oracle.

vertical partitioning

Selecting only the relevant *columns* for replication to a remote site. In addition to minimizing the amount of data sent to a site, it provides protection from changes at the master site. When adding a column to a master table (or dropping a column not in use by a dependent snapshot), the snapshot can continue to perform fast refreshes, without being affected by the DDL change.

virtual private database

A feature consisting of the fine-grained access control and application context.

WebDB

An HTML-based development tool for building HTML Web pages with content based on data stored in the Oracle database. It also provides an ease of use GUI interface to Oracle data, so a user is not required to know SQL to access their data.

Index

A

Advanced Networking Option, 3-70

Advanced Queuing, 3-37 to 3-40, 4-6

- database event publication, 3-39
- multiple consumer queues, 3-39
- nonpersistent queues, 3-40
- object level and system level privileges, 3-39
- object type support, 2-5
- parallel server mode, 3-40
- propagation, enhanced, 3-39
 - addition of a sender's ID column, 3-39
 - history management information, 3-39
 - single view, 3-39
 - statistics, 3-39
- rules engine, 3-38
- see also* publish/subscribe

advanced replication, 3-64 to 3-66

- column level snapshot subsetting, 3-65
- internal apply packages, 3-64
- mass deployment support, 3-65
- offline snapshot instantiation, 3-65
- Oracle Lite integration, improved, 3-66
- Replication Manager, 3-66
- security, improved, 3-66
- snapshot deployment templates,
 - parameterized, 3-65
- snapshot refresh, 3-64

Advanced Security option. *See* Oracle Advanced Security.

aggregate distinct queries, parallelization of, 3-25

ALTER INDEX statement

- COMPUTE STATISTICS clause, 3-27
- online rebuild, 3-41

ALTER TABLE statement

- drop column support, 3-41
- ENABLE ROW MOVEMENT, 3-46
- merging partitions, 3-45
- MODIFY CONSTRAINT clause, 3-28
- MOVE clause for IOT or non-partitioned table, 3-29
- no longer invalidates bitmap indexes, 3-25

ANALYZE statement

- dangling REFs, 3-36
- parallel, 4-2

ANO. *See* Advanced Networking Option.

application context

- CREATE CONTEXT statement, 3-69

application development, 3-69

Application Generator for Microsoft Visual Studio, 3-85

application schema, 2-9

automated parallel query tuning, 3-27

automated standby databases, 3-50

- automatic archival, 3-50
- read-only databases, 3-51

autonomous PL/SQL blocks, 3-78

B

backup and recovery, incremental, 4-6

bitmap indexes, 4-4

- ALTER TABLE no longer invalidates, 3-25
- new ALTER TABLE clauses
 - MINIMIZE RECORDS PER BLOCK, 3-25
 - NOMINIMIZE RECORDS PER BLOCK, 3-25

C

- cache fusion, 3-61
- CACHE READS option, 2-13
- call specs, 3-19, 3-44
- cartridge services, 3-32
- CASE expressions, 2-14
- certificate authority, 3-74
- collections
 - collection unnesting, 3-35
 - SQL*Loader support, 3-36
- COM Cartridge, 3-87
- compatibility flag, 7-6
- composite partitioning, 3-46
- COMPUTE STATISTICS clause of CREATE and ALTER INDEX statements, 3-27
- CONNECT INTERNAL desupported, 7-6
- Connection Manager, Oracle, 4-6
- connection multiplexing, 4-6
- connection pooling, 4-6
- constraint processing, 4-15
- constraints
 - constraint modification with ALTER TABLE, 3-28
 - DISABLE VALIDATE constraint state, 3-28
- control files, 3-52
- CORBA, 3-19
- CREATE CONTEXT statement, 3-69
- CREATE MATERIALIZED VIEW statement, 3-22
- CREATE OUTLINE statement, 3-37
- CREATE_STORED_OUTLINES session parameter, 3-37
- Cryptographic Toolkit, 7-8
- CUBE operation, GROUP BY clause, 3-24

D

- Data Migration Assistant, 2-4, 3-60
- data mining, 3-24
- data warehousing, 3-21 to 3-30
 - Database Resource Manager, 3-48
 - DISABLE VALIDATE state, 3-28
 - function-based indexes, 3-24
 - materialized views, 3-21
 - related features, 3-30
 - SAMPLE function, 3-24

- statistics gathering
 - COMPUTE STATISTICS clause of CREATE and ALTER INDEX statements, 3-27
 - DBMS_STATS package, 3-26
 - transportable tablespaces, 3-23
- Database Configuration Assistant, 2-4, 3-59
- database instance registration
 - client load balancing, 3-76
- Database Resource Manager, 3-48
 - DBMS_RESOURCE_MANAGER package, 3-49
 - DBMS_RESOURCE_MANAGER_PRIVS package, 3-49
 - resource allocation methods, 3-49
 - resource consumer groups, 3-49
 - resource plan directives, 3-49
 - resource plans, 3-49
- database SUSPEND/RESUME, 3-52
- date format strings, 7-2
- DBA Management Pack, 2-3
- direct path load API, 3-23
- distributed lock manager (DLM), 3-62
- distributed systems, 3-64 to 3-68
- documentation
 - online, 6-2
- drop column, 3-41
- dynamic (connection) load balancing, Net8, 3-77

E

- Enterprise Java Beans, 3-19
- enterprise roles, 2-9
- Export/Import utilities, 3-57
- extended aggregate operators, 3-24
- extensibility, 4-14
- extensibility framework, 3-30 to 3-33
 - cartridge services, 3-32
 - extensible indexing, 3-31
 - extensible optimizer, 3-32
 - heterogeneous services, 3-66
 - user-defined operators, 3-31
- extensible indexing, 3-31
- extensible optimizer, 3-32
- external procedures (routines)
 - call specification, 3-44
 - DLL Caching, 3-45

- object support, 3-45
- Pro*C/C++, 3-82
- external procedures. See external routines
- external routines, 4-12
- external routines. See external procedures (external routines)

F

- factoring matrix
 - features, 5-4
 - options, 5-2
- failover
 - with Oracle Enterprise Manager, 3-76
- Fast-Start fault recovery, 3-51
 - see also transaction recovery
- feature factoring matrix, 5-4
- fine-grained access control, 3-69
 - DBMS_RLS package, 3-69
- free lists
 - reducing contention, 2-13, 2-17
- function-based indexes, 3-24

G

- global roles, 2-9
- Group Membership Services (GMS), 3-62

H

- hash clusters
 - single-table-only, 3-26
- hash partitioning, 3-46
- heterogeneous services
 - agent self-registration, 3-67
 - agent shared library, 3-67
 - fixed views, 3-67
- high availability, 2-2

I

- iFS. See Internet File System
- IIOp, 3-70
- indexes
 - building online, 3-41
 - defragment using COALESCE, 3-42

- descending, 3-25
- extensible indexing, 3-31
- function-based, 3-24
- secondary for index-organized tables, 3-29
 - see also bitmap indexes
- index-organized tables, 4-14
 - full table functionality, 3-29
 - key compression, 3-29
 - LOB column support, 3-29
 - logical ROWIDs, 3-29
 - nested table data, 3-34
 - parallel loading of, 3-29
 - partitioning, 3-47
 - rebuilding with ALTER TABLE MOVE, 3-29
 - secondary indexes, 3-29
- instance affinity for jobs, 3-63
 - DBMS_JOB package, 3-63
- INSTEAD OF trigger, 4-15
- interMedia, 2-6
 - Audio, Image, and Video, 3-8
 - Audio, Image, and Video Java Client, 3-9
 - Locator, 3-9
 - Text, 3-6
 - Web based tools, 3-3
- Internet File System, 3-3
- invoker rights, 3-79
- IOTs. See index-organized tables

J

- Java, 4-13
 - JDBC drivers, 3-17
 - SQLJ translator, 3-18
- Java development tools/utilities
 - JDeveloper, 3-20
 - JPublisher, 3-20
 - loadjava and dropjava, 3-21
- Java Message Service, 2-5
- Java VM, 2-5, 3-16
 - development models
 - CORBA, 3-18
 - EJBs, 3-18
 - Java stored procedures, 3-18
 - JServer Accelerator, 3-18
- JavaASO, 2-8

JDBC, 4-13
JDeveloper, 2-5, 3-20
JPublisher, 3-20
JServer Accelerator, 3-16
JServer. *See* Java VM
JSQL, 4-13
JVM. *See* Java VM

L

LDAP, 2-10, 3-74
Legato Storage Manager, 4-6
Lightweight Directory Access Protocol. *See* LDAP.
listener
 configuring for Oracle Enterprise Manager, 3-76
loadjava and dropjava, 3-21
LOBs
 convert LONG to LOB, 3-33
 DBMS_LOB API to append data to LOBs, 3-34
 in index-organized tables, 3-29
 LOB chunk size, 3-34
 Open/Close/IsOpen APIs, 3-34
 partitioning of tables with, 3-47
 SQL*Loader support, 3-36
 temporary, 3-33
 variable width character sets for CLOBs and
 NCLOBs, 3-34
locally managed tablespaces, 3-41
LOG_ARCHIVE_DEST_n initialization
 parameter, 3-50
logical ROWIDs, 3-29
LogMiner, redo log analysis, 3-40
long running operations, statistics generation
 for, 3-27

M

materialized views, 2-11
materialized views. *See* summary management
Media Management API
 version 2.0, 3-53
Microsoft Transaction Server integration, 3-87
migration utility, 4-14
multimedia data, 4-13
multiplexing, connection, 4-6

multi-tier authentication and authorization, 3-78

N

national language support, 3-43
native authentication on Windows NT, 3-77
NCHAR datatype, 4-15
nested tables
 collection locators, 3-34
 data stored in indexed-organized table, 3-34
 OCICollIsLocator(), 3-34
 triggers on view columns, 3-35
 user-specified storage clauses for, 3-34
 UTL_COLL package, 3-34
Net8, 3-64, 3-70, 3-75
 dynamic (connection) load balancing, 3-77
 native authentication on Windows NT, 3-77
 Net8 Assistant, 3-75
 Net8 Easy Config, 3-75
 service naming, 3-75
Net8 Configuration Assistant, 2-4, 3-77
NOCOPY parameter passing mode, 3-80
non-deferred linking, 7-4

O

object cache, 3-35, 3-36, 4-12
Object Database Designer, 4-13
object identifiers, 7-8
object relational enhancements, 2-12, 3-33 to 3-37
 ANALYZE statement, dangling REFs, 3-36
 collection unnesting, 3-35
 object cache enhancements, 3-36
 object identifiers, user defined, 3-34
 object views, 3-35
 parallel query support, 3-37
 partitioning of tables containing user-defined
 types, 3-37
 referential integrity constraint on REF
 column, 3-34
 SQL*Loader support, 3-36
 value subqueries, 3-36
 see also LOBs
 see also nested tables
 see also varrays

- Object Type Translator (OTT), 4-12
- object types, 4-11
- object views, 3-35, 4-11, 4-12
- object-relational technology, 4-10
- Objects for OLE, 3-86
- OCI enhancements, 3-84
 - API to append data to LOBs, 3-34
 - AQ support, 3-38
 - array DMLs, 3-84
 - asynchronous event notification, 3-85
 - collection locator, 3-34
 - direct path load, 3-23
 - DML returning, 3-84
 - LOBs, 3-34
 - NLS interface, 3-43
 - non-blocking OCI, 3-85
 - object cache, 3-36
 - OCI callback registrations, 3-84
 - OCIDescribeAny, 3-85
 - Open/Close/IsOpen API for LOBS, 3-34
 - UCS2 (fixed-width Unicode) support, 3-43
- OCI non-deferred linking, 7-4
- OCI single-task linking, 7-5
- OCI support for client notification, 3-38
 - see also* OCI enhancements, asynchronous event notification
- OiD. *See* Oracle Internet Directory.
- online index builds, 3-41
- online read-only tablespaces, 3-42
- optimizer
 - and remote join enhancement, 3-68
 - and ROLAP enhancement, 3-24
 - extensible, 3-32
 - plan stability, 3-37
- option factoring matrix, 5-2
- Oracle Advanced Security, 3-70
 - integrated security and directory services, 3-72
 - RADIUS, 3-72
 - SSL, 3-70
- Oracle Advanced Security option, 7-8
- Oracle Connection Manager, 4-6
- Oracle Database Configuration Assistant, 3-59, 3-62
- Oracle Enterprise Login Assistant, 2-9
- Oracle Enterprise Manager, 2-3, 3-60

- connect-time failover, 3-76
- static service information in
 - LISTENER.ORA, 3-76
- Oracle Enterprise Security Manager, 2-9
- Oracle Internet Directory, 3-70, 3-74
- Oracle Management Server (OMS), 3-60
- Oracle Parallel Server, 3-61 to 3-63
 - cache fusion, 3-61
 - database configuration, 3-62
 - distributed lock manager (DLM), 3-62
 - Group Membership Services (GMS), 3-62
 - installation, 3-62
 - instance affinity for jobs, 3-63
 - inter-process communication protocol, 3-62
 - Oracle Parallel Server Assistant, 3-63
 - Oradebug, 3-62
 - see also*, OPSM
- Oracle Parallel Server Assistant, 3-62
- Oracle Parallel Server Management (OPSM), 3-62
 - Group Membership Services (GMS), 3-62
 - Oracle Parallel Server Assistant, 3-62
 - Parallel Server Manager, Oracle, 3-62
- Oracle Security Server, 3-70, 3-72, 7-8
- Oracle Software Packager, 3-58
- Oracle Spatial, 2-7, 3-10
- Oracle Time Series, 3-13
- Oracle Universal Installer, 2-4, 3-58, 3-62
- Oracle Visual Information Retrieval, 3-12
- Oracle Wallet Manager, 2-8, 3-73
- Oradebug, 3-62
- OSS. *See* Oracle Security Server.

P

- packages
 - DBMS_APPLICATION_INFO, 3-28
 - DBMS_DEBUG, 3-80
 - DBMS_JOB, 3-63
 - DBMS_LOB, 3-34
 - DBMS_OLAP, 3-22
 - DBMS_PROFILER, 3-80
 - DBMS_REPAIR, 3-40
 - DBMS_RESOURCE_MANAGER, 3-49
 - DBMS_RESOURCE_MANAGER_PRIVS, 3-49
 - DBMS_RLS, 3-69

- DBMS_SQL, 3-79
- DBMS_STATS, 3-26
- DBMS_TRACE, 3-80
- OUTLN_PKG, 3-37
- UTL_COLL, 3-34
- UTL_REF, 3-80
- parallel ANALYZE, 4-2
- parallel DML, 3-46, 4-4
- parallel index scans, 4-4
- parallel query, 4-4
 - automated tuning, 3-27
- Parallel Server Manager, Oracle, 3-62
- parallelization of aggregate distinct queries, 3-25
- parameters
 - passing by reference, 3-80
- partition views, 7-6
- partitioning, 2-11, 4-2
 - "rolling window", 2-11
 - new partitioning methods
 - composite, 3-46
 - hash, 3-46
 - local indexes, 3-46
 - queries, DML, and DDL for, 3-46
 - of index-organized tables, 3-47
 - of object tables, 3-37
 - of tables containing user defined types, 3-37
 - of tables with LOBs, 3-47
 - range partitioning enhancements
 - merging partitions, 3-45
 - updatable partition keys, 3-46
- partitions
 - elimination, 3-47
- plan stability for optimizer, 3-37
- PL/SQL, 3-78
 - autonomous blocks, 3-78
 - bulk binds, 3-79
 - dynamic SQL
 - DBMS_SQL package, 3-79
 - increased package body size, 3-81
 - invoker rights, 3-79
 - monitoring and analysis of program execution
 - DBMS_DEBUG, 3-80
 - DBMS_PROFILER, 3-80
 - DBMS_TRACE, 3-80
 - NOCOPY parameter passing mode, 3-80

- parameter passing by reference, 3-80
- procedures for supporting REF-based operations
 - UTL_REF package, 3-80
- purity rules relaxed, 3-81
 - see also* external procedures (routines)
- PL/SQL Server Pages, 2-6
- precompilers
 - enhancements, 3-81
- Pro*C/C++, 3-81
- Pro*COBOL, 3-82
- program interfaces, 3-78 to 3-87
- publish/subscribe
 - client notification, 3-38
 - database event publication, 3-39
 - publishers, 3-38
 - rules-based subscriptions, 3-38
 - subscribers, 3-38
- publish/subscribe, AQ based, 3-38

Q

- query rewrite, 2-12

R

- RADIUS, 3-72
- range partitioning, 3-45
- read-only databases, 3-51
- recoverability and availability enhancements
 - multiple ARCH processes, 3-50
 - multiple, remote archive destinations, 3-50
 - see also* automated standby databases
- Remote Authentication Dial-In User Service. *See* RADIUS
- remote join enhancements, 3-67
- Reorganizing non-partitioned table with ALTER TABLE MOVE clause, 3-42
- replication, 4-9
 - Replication Manager, 3-66
 - see also* advanced replication
- Replication Manager, 3-66
- reverse key indexes, 4-14
- ROLAP enhancements, 3-24
- ROLLUP operation, GROUP BY clause, 3-24
- ROWIDs

logical, 3-29
universal, 3-29

S

SAMPLE function, 3-24
Secure Socket Layer. See SSL
security enhancements, 3-68 to 3-77
 application development, 3-69
 fine-grained access control, 3-69
 integrated security and directory services, 3-72
 multi-tier authentication and authorization, 3-78
 native authentication on Windows NT, 3-77
SERIALIZABLE=TRUE, 7-2
Server Manager, 7-7
service naming, Net8, 3-75
SGADEF file, 7-7
single sign-on, 4-9
single-table-only hash clusters, 3-26
single-task linking, 7-5
Software Packager, Oracle, 3-58
sort improvements, 3-26
spatial data, 3-10
SQL, 2-14
 analytic functions, 2-10
SQL*Loader, 3-36, 3-47, 3-56
SQL92_SECURITY initialization parameter, 7-2
SSL, 2-8, 3-70
 Public Key Infrastructure, 3-73
 X.509v3 certificates, 3-71
star-query optimization, 4-3
statistics generation for long running operations, 3-27
stored procedures, Java, 3-18
summary management, 3-21
 CREATE MATERIALIZED VIEW statement, 3-22
 dimensions (dimension schema object), 3-22
 materialized views, 3-21
 query rewrite, 3-22
 refreshing materialized views, 3-22
 summary advisory functions package (DBMS_OLAP), 3-22

T

tables
 drop column support, 3-41
 temporary, 3-42
tablespace point-in-time recovery, 4-6
tablespaces
 locally managed, 3-41
 online read-only, 3-42
 point-in-time recovery, 4-6
 transportable, 3-23
temporary tables, 3-42
THE keyword for subqueries, 7-7
Time Series, 3-13
transaction recovery
 fast-start checkpointing, 3-51
 Fast-Start on-demand rollback, 3-51
 Fast-Start parallel rollback, 3-52
 inter-transaction, 3-52
 intra-transaction, 3-52
transparent application failover, 4-8
transportable tablespaces, 3-23
triggers
 on DATABASE and SCHEMA, 3-39
 on nested table view columns, 3-35
 with CALL to a procedure as trigger body, 3-39
TRIM function, 3-44

U

Universal Installer, 3-58, 3-62
universal ROWIDs (UROWIDs), 3-29
UROWID datatype, 3-29
USE_STORED_OUTLINES session parameter, 3-37
user defined object identifiers, 3-34
user-defined operators, 3-31

V

V\$OPTION table, 5-13
V\$SESSIONS_LONGOPS view, 3-27
Version 6, 7-6
value subqueries
 argument to a type constructor or function, 3-36
 in VALUES clause of INSERT statement, 3-36
varrays

storage parameters for storing as LOBs, 3-35
virtual private database, 3-45, 3-68
Visual Information Retrieval, 3-12

W

Wallet Manager, Oracle, 2-8, 3-73
WebDB, 2-6, 3-2, 3-5

X

X.509 certificates, 3-72
X.509v3 certificates, 3-71, 3-73
XA support, 4-9
XML Class Generator, 2-6
XML Parser for Java, 2-6

Z

zero-length strings, 7-2