Oracle® Database

New Features Guide

10*g* Release 2 (10.2)

B14214-02

November 2005

ORACLE®

Oracle Database New Features Guide, 10*g* Release 2 (10.2)

B14214-02

Copyright © 2001, 2005, Oracle. All rights reserved.

Primary Author: Rhonda Day

Contributing Author: Diana Lorentz, Antonio Romero

Contributor: Mark Townsend, Neal Brenner, Sandra Cheevers, Michael Hartstein

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

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

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software—Restricted Rights (June 1987). Oracle 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 we disclaim liability for any damages caused by such use of the Programs.

Oracle, JD Edwards, PeopleSoft, and Retek are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Contents

Send Us Your Comments xiii

Preface xv

Audience xv
Documentation Accessibility xv
Related Documents xvi
Conventions xvi

1 Oracle Database 10*g* New Features

1.1 Application Development 1-1
1.1.1 Application Program Interface (API) and Precompilers 1-1
1.1.1.1 Oracle Open Database Connectivity (ODBC) Driver Support for Linux 1-1
1.1.1.2 Oracle Call Interface (OCI) Client Version API 1-2
1.1.1.3 Start Up and Shut Down of Database OCI API 1-2
1.1.1.4 Oracle C++ Call Interface (OCCI) Transparent Application Failover Support 1-2
1.1.1.5 Instant Client Light 1-2
1.1.1.6 Instant Client Software Development Kit (SDK) 1-2
1.1.2 Globalization and Unicode Improvements 1-2
1.1.2.1 Unicode 4.0 1-3
1.1.2.2 Unicode Data Type Literal Enhancement 1-3
1.1.2.3 Globalization Development Kit (GDK) 2.0 1-3
1.1.2.4 Character Set Scanner (CSSCAN) Enhancements 1-3
1.1.2.5 Language and Character Set File Scanner (LCSSCAN) Enhancements 1-4
1.1.2.6 Consistent Linguistic Searching and Sorting 1-4
1.1.2.7 Transliteration for Japanese Kana Conversion 1-4
1.1.3 Java, Java Database Connectivity (JDBC), and Web Services 1-4
1.1.3.1 JDBC XA Connections Caching 1-4
1.1.3.2 JDBC Runtime Connection Load Balancing 1-5
1.1.3.3 JDBC Support for Data Manipulation Language (DML) in the RETURNING Clause 1-5
1.1.3.4 Support for JDBC 3.0 ResultSet Holdability 1-5
1.1.3.5 Support for JDBC 3.0 Retrieval of Auto-Generated Key 1-5
1.1.4 Oracle HTML DB Improvements 1-5
1.1.4.1 Cluster Bar Charts 1-6
1.1.4.2 Dial Charts 1-6

1.1.4.3 Percentage Bar Charts ... 1-6
1.1.4.4 Stacked Bar Charts ... 1-6
1.1.4.5 Asynchronous Refresh for All Charts Based on Scalable Vector Graphics (SVG) ... 1-7
1.1.4.6 Create Master Detail With Choice of Layout ... 1-7
1.1.4.7 Application From a Spreadsheet Wizard ... 1-7
1.1.4.8 Create Forms and Reports Based on a Web Service Reference ... 1-7
1.1.4.9 Discover Web Services Through UDDI Repository ... 1-8
1.1.4.10 Simplified and More Consistent User Interface ... 1-8
1.1.5 PL/SQL Improvements ... 1-8
1.1.5.1 PL/SQL Conditional Compilation ... 1-8
1.1.5.2 Obfuscation of Dynamically Generated PL/SQL Source Code ... 1-9
1.1.5.3 Relaxation of Line Length and Overall Limits for the DBMS_OUTPUT PL/SQL Package ... 1-9
1.1.5.4 UTL_NLA PL/SQL Package for Matrix Math ... 1-9
1.1.6 SQL Language Improvements ... 1-10
1.1.6.1 Rules Manager ... 1-10
1.1.6.2 Regular Expression Enhancements ... 1-11
1.1.6.3 Collection Operator Performance Improvements ... 1-11
1.1.6.4 Distributed Large Objects (LOB) Support ... 1-11
1.1.6.5 Expression Filter Enhancements ... 1-11
1.1.7 XML Application Development ... 1-12
1.1.7.1 XQuery ... 1-12
1.1.7.2 Advanced XPath Rewrite ... 1-12
1.1.7.3 New InsertXML(), AppendChildXML(), InsertXMLBefore(), and DeleteXML() Functions ... 1-12
1.1.7.4 Simple Object Access Protocol (SOAP) 1.2 Implementation in C/C++ ... 1-13
1.1.7.5 SQL/XML Standard (SQL:2003 Standard Part 14) Compliance ... 1-13
1.1.7.6 SQL/XML Standard Compliance ... 1-13
1.1.7.7 Easy XML Application Development Using JAXB Customization ... 1-14
1.1.7.8 Oracle XML Developer's Kit (XDK) PL/SQL Document Object Model (DOM) API Deprecation ... 1-14
1.1.7.9 XPath 2.0 and XQuery 1.0 Functions and Operators Support in Java ... 1-14
1.1.7.10 Custom Metadata for Oracle XML DB Repository Resources ... 1-14
1.1.7.11 Secure Sockets Layer (SSL) Support ... 1-15
1.1.7.12 Transportable Tablespace Support ... 1-15
1.1.7.13 Oracle XML DB Administration in the Enterprise Manager Web Console ... 1-15
1.2 Availability ... 1-15
1.2.1 Fast-Start Failover to Standby Database ... 1-15
1.2.1.1 Fast-Start Failover ... 1-16
1.2.1.2 Automatic Conversion of the Primary Database to a Standby Database Upon Failover ... 1-16
1.2.1.3 Optimized Asynchronous Redo Transmission ... 1-16
1.2.1.4 Faster Redo Apply Failover ... 1-16
1.2.1.5 Faster SQL Apply Failover ... 1-17
1.2.2 Support for Database Transport Across Same Endian Platforms ... 1-17
1.2.2.1 Database Transport Across Same Endian Platforms ... 1-17
1.2.3 High Availability Improvements ... 1-17

1.2.3.1 Add and Create a Transportable Tablespace to a Database ... 1-17
1.2.3.2 Additional Data Type Support in LogMiner and SQL Apply ... 1-18
1.2.3.3 Automatic Deletion of Applied Archive Logs ... 1-18
1.2.3.4 Data Guard: Improved Manageability for Redo Transport, Log Apply, and Broker ... 1-18
1.2.3.5 Dynamic RMAN Channel Allocation for Backing Up and Recovering RAC Database ... 1-18
1.2.3.6 RMAN Backup Encryption ... 1-19
1.2.3.7 Easy Conversion of a Physical Standby Database to a Reporting Database ... 1-19
1.2.3.8 Eliminate Control File Re-Creation ... 1-19
1.2.3.9 Enhanced Memory Corruption Checking ... 1-19
1.2.3.10 Enhanced RMAN Backup Job Views ... 1-20
1.2.3.11 Enhanced Server Side Shutdown ... 1-20
1.2.3.12 Enhanced Space Management Verification ... 1-20
1.2.3.13 Failover Improvements for Real Application Clusters (RAC) Distributed Transactions ... 1-20
1.2.3.14 Faster Instance Startup for Large Memory Configurations ... 1-20
1.2.3.15 Fine-Grained, Automated Tracking of SQL Apply Runtime Performance ... 1-21
1.2.3.16 Flashback Across Data Guard Switchovers ... 1-21
1.2.3.17 Flashback Database Through Resetlogs ... 1-21
1.2.3.18 Hardware Assisted Resilient Data (HARD) Compliant RMAN Backups ... 1-21
1.2.3.19 Increase Supported Real Application Clusters Instances for a Database ... 1-21
1.2.3.20 Incremental Roll Forward of Database Copy ... 1-22
1.2.3.21 Online Redefinition Supports Clustered Tables, ADT's, MV Logs, AQ Tables, and Preserves Database Statistics ... 1-22
1.2.3.22 Online Redefinition of a Single Partition ... 1-22
1.2.3.23 Optimized Creation of Logical Standby Database ... 1-22
1.2.3.24 Unused Block Compression ... 1-23
1.2.3.25 Reorganized Object Types Improves User Interface ... 1-23
1.2.3.26 Restore Point ... 1-23
1.2.3.27 Temporary Datafiles Are Re-Created on RMAN Recovery ... 1-23
1.3 Business Intelligence ... 1-24
1.3.1 SQL Support for Analytic Applications ... 1-24
1.3.1.1 Enhanced Upsert Abilities in MODEL Clause ... 1-24
1.3.1.2 Enhancements to MODEL Rules Which Use FOR Loops ... 1-24
1.3.1.3 Use of Analytic Functions in the SQL MODEL Clause ... 1-24
1.3.1.4 Intra-Partition DML in the Presence of Bitmap Indexes ... 1-24
1.3.2 Improved Very Large Database (VLDB) Support ... 1-25
1.3.2.1 Increased Maximum Number of Partitions Per Object ... 1-25
1.3.2.2 Enhanced Dynamic Partition Pruning ... 1-25
1.3.2.3 Resource Optimized DROP TABLE for Partitioned Tables ... 1-25
1.3.3 Reduced Information Cycle Time ... 1-25
1.3.3.1 Distributed Asynchronous Change Data Capture (CDC) ... 1-25
1.3.3.2 DML Error Logging Table ... 1-26
1.3.3.3 Partition Change Tracking Refresh Without Materialized View Logs ... 1-26
1.3.3.4 Query Rewrite Using Multiple Materialized Views ... 1-26
1.3.3.5 Skip Level Dimension Hierarchies ... 1-26
1.3.4 Online Analytical Processing (OLAP) Support for Analytic Applications ... 1-27

1.3.4.1 Incremental Aggregation of Compressed OLAP Cubes ... 1-27
1.3.4.2 OLAP API Supports Calculated (Virtual) Dimension Members ... 1-27
1.3.4.3 OLAP Compressed Cubes Support Semiadditive Aggregation Methods ... 1-27
1.3.5 Data Mining Support for Analytic Applications ... 1-28
1.3.5.1 Decision Tree Algorithm ... 1-28
1.3.5.2 Predictive Analytics Package ... 1-28
1.3.5.3 SQL-Level Data Mining Scoring Capability ... 1-28
1.3.5.4 Java Data Mining (JDM) Compliant Java API ... 1-28
1.4 Clustering ... 1-29
1.4.1 Real Application Clusters (RAC) Performance ... 1-29
1.4.1.1 Load Balancing Advisory ... 1-29
1.4.1.2 Oracle Cluster File System for Solaris ... 1-29
1.4.2 Real Application Clusters Robustness ... 1-29
1.4.2.1 Clusterware File Redundancy ... 1-29
1.4.2.2 Cluster Verification ... 1-30
1.4.2.3 Cluster Tracing and Diagnostic Enhancements ... 1-30
1.4.3 Application High Availability With Oracle Clusterware ... 1-30
1.4.3.1 High Availability API ... 1-30
1.5 Content Management ... 1-31
1.5.1 Oracle Text Improvements ... 1-31
1.5.1.1 CTXRULE Asian Language Support ... 1-31
1.5.1.2 Filtering ... 1-31
1.5.1.3 Key Word in Context (KWIC) ... 1-31
1.5.1.4 Online Indexing for Local Partitioned Index ... 1-31
1.5.2 *inter*Media Improvements ... 1-32
1.5.2.1 *inter*Media DICOM Medical Imaging Format Support ... 1-32
1.5.2.2 *inter*Media Support for Image Content Metadata and XML DB Integration ... 1-32
1.5.2.3 Java Advanced Imaging 1.1.2 Support ... 1-33
1.6 Database Overall ... 1-33
1.6.1 Database Utilities ... 1-33
1.6.1.1 Easier Manual Upgrades ... 1-33
1.6.1.2 Enhanced Pre-Upgrade Information Utility ... 1-33
1.6.1.3 Default DATA_PUMP_DIR Directory Object ... 1-33
1.6.1.4 Dumpfile Compression ... 1-33
1.6.1.5 Extended Job Status Display ... 1-34
1.6.1.6 FILESIZE Value for Dumpfiles ... 1-34
1.6.1.7 Monitoring of Data Pump Jobs Started Outside of Enterprise Manager ... 1-34
1.6.1.8 New DBMS_DATAPUMP API Calls ... 1-34
1.6.1.9 New PCTSPACE Transform Parameter on Data Pump Import ... 1-34
1.6.1.10 Object ID (OID) Transform Parameter ... 1-34
1.6.1.11 SQL*Plus AUTOTRACE EXPLAIN Conversion to DBMS_XPLAN ... 1-35
1.6.1.12 SQL*Plus DBMS_OUTPUT Buffer and Line Size Limit Increase ... 1-35
1.6.1.13 SQL*Plus XQuery Support ... 1-35
1.7 Database Security ... 1-35
1.7.1 Privacy and Confidentiality ... 1-35
1.7.1.1 Transparent Data Encryption ... 1-35
1.7.1.2 Key Management for Data Encryption ... 1-36

1.7.2 Improved Manageability 1-36
1.7.2.1 Enhanced Proxy Authentication 1-36
1.7.2.2 Audit Trail in Standard XML Format 1-36
1.7.3 Secure By Default 1-37
1.7.3.1 Connect Role Privilege Reduction 1-37
1.7.4 Improved Integration and Interoperability 1-37
1.7.4.1 Secure External Password Store 1-37
1.7.4.2 Enhanced Centralized Password Policies for the Database 1-38
1.7.4.3 Oracle Label Security Support for Real Application Clusters 1-38
1.8 Grid Computing 1-38
1.8.1 Data Provisioning 1-38
1.8.1.1 File Groups Repository 1-38
1.8.1.2 Transportable Tablespace from Backup 1-38
1.8.2 Oracle Scheduler 1-39
1.8.2.1 Event Scheduling 1-39
1.8.2.2 Job Chaining 1-39
1.8.2.3 Scheduler Enhancements 1-39
1.9 Grid Management 1-39
1.9.1 Configuration Plug-ins and Assistants 1-40
1.9.1.1 Enterprise Manager Configuration Plug-in 1-40
1.9.1.2 Migrate Non-ASM Databases to ASM 1-40
1.9.1.3 Standalone Automatic Storage Management (ASM) Configuration 1-40
1.9.1.4 Standalone Disk Group Management 1-40
1.9.1.5 Upgrade 10.1 Database to 10.2 1-41
1.9.2 Database Management 1-41
1.9.2.1 Ability to Discover and Update Host Cluster Targets in Enterprise Manager 1-41
1.9.2.2 Access Advisor Enhancements in Enterprise Manager 1-41
1.9.2.3 Backup and Recovery Configuration Enhancements in Enterprise Manager .. 1-41
1.9.2.4 Clone External Objects as Part of the Cloning Operation 1-41
1.9.2.5 Database Administration Enhancements in Enterprise Manager 1-42
1.9.2.6 Database Auditing Administration in Enterprise Manager 1-42
1.9.2.7 Enable or Disable Undo Retention Guarantee 1-42
1.9.2.8 Enhanced Backup Management 1-42
1.9.2.9 Enhanced Diagnostics for Cache Coherency Statistics 1-43
1.9.2.10 Integrate Segment Advisor with Segment Findings 1-43
1.9.2.11 Lock and Unlock Statistics 1-43
1.9.2.12 Management of Directory Objects in the Enterprise Manager Web Console ... 1-43
1.9.2.13 Miscellaneous Undo Advisor User Interface Enhancements 1-43
1.9.2.14 New Storage and Configuration Policies 1-44
1.9.2.15 Oracle Spatial Administration in the Enterprise Manager Web Console 1-44
1.9.2.16 Performance Monitoring for Large Clusters 1-44
1.9.2.17 Restore Statistics From Historic Snapshots 1-45
1.9.2.18 Restore and Recovery Enhancements in Enterprise Manager 1-45
1.9.2.19 Segment Advisor to Work on the Whole Database 1-45
1.9.2.20 Show Segment Details of Tablespace 1-45
1.9.2.21 Statistics Collection for Fixed and Dictionary Objects 1-45
1.9.2.22 Tabular Display of Tablespace Segments 1-46

1.9.2.23 Track Interconnect Information Statistics ... 1-46
1.9.2.24 Use RMAN Compression During Cloning ... 1-46
1.9.2.25 Wizard-Based Loading of Flat Files in Enterprise Manager ... 1-46
1.9.3 Patching ... 1-46
1.9.3.1 OPatch the Patching Tool ... 1-46
1.9.4 Software Install and Cloning ... 1-47
1.9.4.1 Improve Automatic Storage Management (ASM) Installation Process ... 1-47
1.9.4.2 Real Application Clusters Cloning ... 1-47
1.9.4.3 Real Application Clusters Install Framework ... 1-47
1.9.4.4 Silent Install for Real Application Clusters ... 1-47
1.10 Information Integration ... 1-47
1.10.1 Oracle Streams Performance and Usability Improvements ... 1-48
1.10.1.1 Streams Performance ... 1-48
1.10.1.2 Support for Hot Mining With Downstream Capture ... 1-48
1.10.1.3 Buffered Messaging ... 1-48
1.10.1.4 Commit-Time Ordering of Messages ... 1-48
1.10.1.5 Apply Handler Dependency Declaration ... 1-49
1.10.1.6 Simplified Large Object (LOB) Handling in Streams Apply Process ... 1-49
1.10.1.7 Declarative LCR Transformations ... 1-49
1.10.1.8 Notification Enhancements ... 1-49
1.10.1.9 Oracle Java Message Service (OJMS) Nonpersistent Messaging ... 1-49
1.10.1.10 Simplified Configuration APIs ... 1-50
1.10.1.11 Oracle Streams Administration in the Enterprise Manager Web Console ... 1-50
1.10.2 Improved Client/Server and Server/Server Communications ... 1-50
1.10.2.1 Local Naming Support for JavaNet ... 1-50
1.10.2.2 Oracle Streams AQ Administration in the Enterprise Manager Web Console 1-50
1.10.2.3 Transparent Application Failover (TAF) Notification ... 1-50
1.11 Location Services ... 1-51
1.11.1 Extend Business Geographics Features to Support Europe ... 1-51
1.11.1.1 Oracle Spatial Routing Engine Support for Western Europe ... 1-51
1.11.2 Extend Coordinate System in Core Spatial ... 1-51
1.11.2.1 Oracle Spatial Support for EPSG Coordinate Systems ... 1-51
1.11.3 GeoRaster Compression Standards ... 1-51
1.11.3.1 Oracle Spatial GeoRaster Compression ... 1-51
1.11.4 Improved Performance of Network Data Model ... 1-52
1.11.4.1 Graph Partitioning for Very Large Networks in Oracle Spatial ... 1-52
1.11.5 Improved Topology Operations ... 1-52
1.11.5.1 Whole Feature Inserts and Updates for the Oracle Spatial Topology Data Model ... 1-52
1.12 Performance and Scalability ... 1-52
1.12.1 Improved Large Object (LOB) Performance ... 1-52
1.12.1.1 LOB Access Statistics in Dynamic Performance Views ... 1-53
1.12.1.2 Online Redefinition Supports Parallel Execution for LONG-to-LOB Migration ... 1-53
1.12.1.3 Performance Improvements for Loading LOBs ... 1-53
1.12.2 Other Performance Improvements ... 1-53
1.12.2.1 Asynchronous Commit ... 1-53
1.12.2.2 Checkpointing Improvements ... 1-53

1.12.2.3 Database Change Notification ... 1-54
1.12.2.4 Fast Partition Split for Partitioned Index Organized Tables ... 1-54
1.12.2.5 Improved Aggregation Performance ... 1-54
1.12.2.6 Parallel Join Bitmap Filtering ... 1-54
1.13 Server Manageability ... 1-54
1.13.1 Simplified Configuration ... 1-54
1.13.1.1 Automatic Undo Retention Enhancement ... 1-55
1.13.1.2 Net Services Administration in the Enterprise Manager Web Console ... 1-55
1.13.1.3 Self-Tuning Multiblock Read Count ... 1-55
1.13.2 Improved Intelligent Infrastructure ... 1-55
1.13.2.1 Adaptive Alert Thresholds in Enterprise Manager ... 1-56
1.13.2.2 Database Performance Monitoring Enhancements in Enterprise Manager ... 1-56
1.13.2.3 Diagnosis of Hung or Extremely Slow Databases Through Enterprise Manager ... 1-56
1.13.2.4 Display Session Activity History in Enterprise Manager ... 1-56
1.13.2.5 Performance Graph Normalization in Enterprise Manager ... 1-56
1.13.2.6 Read Performance Statistics Directly From the System Global Area (SGA) ... 1-57
1.13.2.7 Cleanup System Statistics ... 1-57
1.13.2.8 Enhanced Active Session History ... 1-57
1.13.2.9 Enhanced Automatic Database Diagnostic Monitor (ADDM) Performance Analysis ... 1-57
1.13.2.10 Progress Monitoring for Advisors ... 1-57
1.13.2.11 Enhanced End-to-End Performance Monitoring ... 1-58
1.13.2.12 Active Session History (ASH) Report ... 1-58
1.13.2.13 Automatic Workload Repository (AWR) Compare Periods Report ... 1-58
1.13.2.14 Configurable AWR SQL Collection ... 1-58
1.13.2.15 Transportable Automatic Workload Repository (AWR) ... 1-58
1.13.3 Automatic Storage Management (ASM) Manageability ... 1-59
1.13.3.1 ASM Command-Line Interface ... 1-59
1.13.3.2 ASM Wait on Rebalance Process ... 1-59
1.13.3.3 Additional ASM Manageability Enhancements ... 1-59
1.13.3.4 Batch Selection of Multiple Disks for Operations Like DELETE, RESIZE, and CHECK ... 1-60
1.13.4 Automatic Storage Management (ASM) Extensions ... 1-60
1.13.4.1 ASM Support for Multiple Database Releases ... 1-60
1.13.4.2 Enhanced Protocol Support for Database to ASM Connection ... 1-60
1.13.5 File Support in Automatic Storage Management (ASM) ... 1-60
1.13.5.1 ASM/XDB FTP Support ... 1-61
1.13.6 Enhanced Automatic Database Diagnostic Monitor (ADDM) ... 1-61
1.13.6.1 Automatic Shared Memory Management of Streams Pool ... 1-61
1.13.7 Enhanced Instance Tuning ... 1-61
1.13.7.1 Automatically Enabled Resource Manager ... 1-61
1.13.7.2 Enhanced Ability to Diagnose Memory Allocation ... 1-61
1.13.7.3 Enhanced Automatic Shared Memory Management ... 1-62
1.13.7.4 Enhanced Resource Manager Monitoring ... 1-62
1.13.8 Enhanced SQL Manageability ... 1-62
1.13.8.1 Transportable SQL Tuning Sets ... 1-62

1.13.8.2 Display SQL History in Enterprise Manager 1-63
1.13.8.3 Interruptible SQL Access Advisor 1-63
1.13.8.4 SQL Access Advisor Recommends Function-Based Indexes 1-63
1.13.8.5 SQL Profile Enhancements 1-63
1.13.8.6 SQL Statistics Enhancement 1-63
1.13.8.7 Timeout for Flushing SQL Statistics 1-64
1.13.9 Enhanced Space Management 1-64
1.13.9.1 Default Automatic Segment Space Management (ASSM) 1-64
1.13.9.2 Automatic Segment Advisor 1-64
1.13.9.3 Online Segment Shrink Enhancements 1-64
1.13.9.4 Size-Based Tablespace Space Usage Alerts 1-64
1.13.9.5 Drop Empty Datafile 1-65
1.14 Windows 1-65
1.14.1 Improved Availability and Scalability With Oracle Real Application Clusters 1-65
1.14.1.1 ODP.NET: Fast Application Notification (FAN) 1-65
1.14.1.2 ODP.NET: Integration with Real Application Clusters Load Balancing Advisory 1-65
1.14.2 Ease of Development 1-66
1.14.2.1 Common Language Runtime (CLR) Stored Procedure Visual Studio Integration 1-66
1.14.2.2 ODP.NET: .NET Framework 1.1 Functionality Support 1-66
1.14.2.3 ODP.NET: Database Change Notification Service 1-67
1.14.2.4 ODP.NET: Derived Parameters Support 1-67
1.14.2.5 ODP.NET: Enhanced Connection Pool Administration 1-67
1.14.2.6 ODP.NET: Query Cancel/Timeout 1-67
1.14.2.7 Server-Side ODP.NET 1-68
1.14.3 Performance Improvements 1-68
1.14.3.1 ODP.NET: Improved LOB Retrieval Performance and Functionality 1-68
1.14.3.2 ODP.NET: Statement Caching 1-68
1.14.4 Security 1-69
1.14.4.1 ODP.NET: Enhanced Virtual Private Database (VPD) Support 1-69
1.14.4.2 Secure Operating System Objects 1-69

2 Oracle Database 10*g* Documentation

2.1 Documentation Titles and Online Documentation Structure 2-1
2.1.1 Oracle Database 10*g* Documentation 2-1
2.1.1.1 General Database Documents 2-1
2.1.1.2 Application Development Documentation 2-2
2.1.1.3 Backup Documentation 2-5
2.1.1.4 Data Mining Documentation 2-5
2.1.1.5 Database Server Documentation 2-5
2.1.1.6 Enterprise Manager Documentation 2-8
2.1.1.7 Gateways Documentation 2-8
2.1.1.8 Installation Documentation 2-11
2.1.1.9 Java Development Documentation 2-16
2.1.1.10 Networking Documentation 2-16
2.1.1.11 OLAP Documentation 2-17

2.1.1.12 Real Application Clusters Documentation 2-18
2.1.1.13 Release Notes Documentation 2-18
2.1.1.14 Text Documentation 2-19
2.1.1.15 Ultra Search Documentation 2-19
2.1.1.16 Windows Documentation 2-20
2.1.1.17 Workflow Documentation 2-21

Index

Send Us Your Comments

Oracle Database New Features Guide 10*g* Release 2 (10.2)

B14214-02

Oracle 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 title and part number of the documentation and the chapter, section, and page number (if available). You can send comments to us in the following ways:

- Electronic mail: infodev_us@oracle.com
- FAX: (650) 506-7227. Attn: Server Technologies Documentation Manager
- Postal service:

  Oracle Corporation
  Server Technologies Documentation Manager
  500 Oracle Parkway, Mailstop 4op11
  Redwood Shores, CA 94065
  USA

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

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

Preface

This Preface contains these topics:

- Audience
- Documentation Accessibility
- Related Documents
- Conventions

Audience

Oracle Database New Features Guide is addressed to people familiar with previous versions of Oracle Database who would like to become familiar with features, options, and enhancements that are new in this release of the database.

Documentation Accessibility

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

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

Accessibility of Code Examples in Documentation

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

Accessibility of Links to External Web Sites in Documentation

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

TTY Access to Oracle Support Services

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

Related Documents

Oracle error message documentation is only available in HTML. If you only have access to the Oracle Documentation CD, you can browse the error messages by range. Once you find the specific range, use your browser's "find in page" feature to locate the specific message. When connected to the Internet, you can search for a specific error message using the error message search feature of the Oracle online documentation.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1

Oracle Database 10*g* New Features

This chapter contains descriptions of all of the features that are new to Oracle Database 10*g* Release 2 (10.2). This chapter contains the following sections:

- Application Development
- Availability
- Business Intelligence
- Clustering
- Content Management
- Database Overall
- Database Security
- Grid Computing
- Grid Management
- Information Integration
- Location Services
- Performance and Scalability
- Server Manageability
- Windows

1.1 Application Development

The following sections describe Application Development features for Oracle Database 10*g* Release 2.

1.1.1 Application Program Interface (API) and Precompilers

These features provide API and functionality improvements to enable ISVs and developers to build more powerful applications with higher levels of availability.

1.1.1.1 Oracle Open Database Connectivity (ODBC) Driver Support for Linux

Support has been added for an Oracle ODBC driver for the Linux platform. Now, ODBC applications can be developed and used on Linux with an Oracle supported ODBC driver.

This support enables ODBC applications to run on the Linux platform, encouraging wider Oracle Database adoption by allowing Linux clients to access Oracle databases.

1.1.1.2 Oracle Call Interface (OCI) Client Version API

This feature adds client versioning macros to support version sensing.

This allows applications to be compiled for multiple versions of the Oracle client from a single source code environment. This leads to easier development of products with customizations based on versions.

1.1.1.3 Start Up and Shut Down of Database OCI API

This feature adds API calls for starting and stopping a database from inside an application.

This enables ISVs and testers to start up and shut down databases without using SQL*Plus or other tools. This streamlines the testing process for application developers.

See also:

Oracle Call Interface Programmer's Guide for details

1.1.1.4 Oracle C++ Call Interface (OCCI) Transparent Application Failover Support

This feature adds extensions to OCCI API to support Transparent Application Failover (TAF).

This enables more robustness with TAF. If a server node goes down, applications automatically reconnect to another surviving node, replaying transactions if possible.

See also:

Oracle C++ Call Interface Programmer's Guide for details

1.1.1.5 Instant Client Light

Instant Client Light reduces the size of the Instant Client package by removing most of the character sets and leaving only English error messages.

This feature allows Instant Client to be used in a wider variety of environments, especially where storage space constraints are prominent.

See also:

Oracle C++ Call Interface Programmer's Guide for details

1.1.1.6 Instant Client Software Development Kit (SDK)

The Instant Client SDK packages the OCI/OCCI header files necessary for developing Oracle client applications outside of an Oracle installation. Now, application developers can develop, test, and deploy Oracle client applications from an Instant Client deployment without the need for an ORACLE_HOME or full client installation.

This feature enables a larger development community to develop client applications for Oracle, encouraging more Oracle Database adoption.

1.1.2 Globalization and Unicode Improvements

Globalization and Unicode support have been enhanced in Oracle Database 10*g* Release 2. Consistent linguistic ordering support has been added for all SQL functions, enhancing the ability to write internationalized applications without regard to locale. The Globalization Development Kit (GDK) for PL/SQL provides new locale mapping tables and Japanese transliteration support. SQL Regular Expression support has been updated to support key Perl conventions used in the Life Science community.

This release also supports the latest version of Unicode 4.0. Identification of potential problems when migrating to another database character set has also been enhanced. The Character Set Scanner (CSSCAN) now offers selective scanning and provides support for new data types. The Language and Character Set Scanner (LCSSCAN) has improved detection quality of shorter text strings and now supports HTML files.

1.1.2.1 Unicode 4.0

Unicode 4.0 is the latest release of the Unicode standard. Oracle's Unicode character sets, AL32UTF8 and AL16UTF16, have been updated to support Unicode 4.0 in Oracle Database 10*g* Release 2. Unicode 4.0 has 1,226 new characters.

This feature provides greater support of the Unicode character set for Oracle Database users.

See also:

Oracle Database Globalization Support Guide for details

1.1.2.2 Unicode Data Type Literal Enhancement

This feature supports SQL string literal INSERT and UPDATE operations on NCHAR without data loss regardless of what the database character set happens to be. Multilingual Unicode data can now be added without restrictions such as having to provide hex Unicode values. The support for this feature is available in SQL, PL/SQL, OCI, and JDBC.

This feature prevents data loss when performing INSERT and UPDATE operations on NCHAR.

See also:

Oracle Database Globalization Support Guide for details

1.1.2.3 Globalization Development Kit (GDK) 2.0

Globalization Development Kit (GDK) has added new locales and common locale mapping information into the GDK for the PL/SQL package.

Among other things, the mappings allow for picking the most appropriate Oracle Linguistic sort when ordering different users' language data.

Additional mapping functions have been added to allow developers to build global PL/SQL applications including Common Sort, Common Time Zone, Common Territory, and Translated Locale Names.

See also:

Oracle Database Globalization Support Guide for details

1.1.2.4 Character Set Scanner (CSSCAN) Enhancements

Two new parameters, QUERY and COLUMN, have been added in CSSCAN in Oracle Database 10*g* Release 2. These parameters offer finer control in performing selective analysis on the user data. Support for some remaining data types that CSSCAN does not handle has been added. They are multilevel VARRAY and multilevel nested tables.

The enhancements in CSSCAN provide more flexibility in defining the criteria for data scanning, which in turn helps to more accurately detect the effect of migrating the database to a new character set.

See also:

Oracle Database Globalization Support Guide for details

1.1.2.5 Language and Character Set File Scanner (LCSSCAN) Enhancements

In this release, enhancements include:

- Improved guessing for shorter length strings
- Support for HTML files

Oracle's Language and Character Set Detection (LCSD) is a high-performance, statistically-based technology for determining the language and the character set for unknown text in a file. It can automatically identify a wide variety of language and character set pairs. Introduced in Oracle Database 10*g* Release 1, LCSD is offered as part of the Migration Utilities and called the Language and Character Set File Scanner (LCSSCAN). LCSD is also offered as part of the Globalization Development Kit through Java API's.

See also:

Oracle Database Globalization Support Guide for details

1.1.2.6 Consistent Linguistic Searching and Sorting

A new NLS_COMP parameter, LINGUISTIC, has been added to support the existing and remaining SQL & PL/SQL operators and functions.

This feature ensures all SQL string comparisons are consistent, and that they follow the linguistic convention as specified in the NLS_SORT parameter.

See also:

Oracle Database Globalization Support Guide for details

1.1.2.7 Transliteration for Japanese Kana Conversion

This feature provides a transliteration mechanism to support various Japanese Kana conversions in different forms. The TRANSLITERATE function for the UTL_I18N PL/SQL package is provided for transliteration in this release.

This feature performs Japanese Kana conversions.

See also:

Oracle Database Globalization Support Guide for details

1.1.3 Java, Java Database Connectivity (JDBC), and Web Services

Oracle Database 10*g* Release 1 brought new major application development and deployment capabilities to Java and Web Services developers. Oracle Database 10*g* Release 2 goes further in terms of JDBC standards support, dynamic load balancing of connections in Real Application Clusters (RAC) and GRID environments, and many performance improvements such as caching XA connections, and DML in the RETURNING clause.

1.1.3.1 JDBC XA Connections Caching

This feature improves JDBC support for caching XAConnections on OracleXADataSources. Applications can now enable Implicit Connection Caching and Fast Connection Caching (in Real Application Clusters environments) on

OracleXADataSources. When XAConnection caching is not enabled, applications obtain XAConnections directly from OracleXADataSource, and then obtain the associated XAResource objects.

This feature improves the performance of distributed transactions through the caching of XA connections in a single instance or multiple instance database environments.

1.1.3.2 JDBC Runtime Connection Load Balancing

This feature provides dynamic routing of connection requests to the least loaded instance. The JDBC connection cache manager uses Real Application Clusters (RAC) workload metrics and the specified load balancing policy to choose the best instance to process the connection request.

The benefit is efficient RAC and Grid database resource utilization by even distribution of the workload across RAC instances based on workload metrics and distribution policies.

1.1.3.3 JDBC Support for Data Manipulation Language (DML) in the RETURNING Clause

This feature enables you to issue a DML command (for example, INSERT, UPDATE, or DELETE) specified in the RETURNING clause upon a SQL query execution.

This feature increases performance by saving a client-to-database round-trip.

1.1.3.4 Support for JDBC 3.0 ResultSet Holdability

By default, the database automatically closes cursors and result sets upon transaction commit. A holdable cursor, or holdable resultset, is not closed when the transaction that contains the cursor is committed thereby reducing hard parses and library latches. JDBC 3.0 ResultSet holdability is specified using the createStatement(), prepareStatement(), or prepareCall() methods and one of the following two constants: HOLD_CURSORS_OVER_COMMIT or CLOSE_CURSORS_AT_COMMIT. Oracle Database supports HOLD_CURSORS_OVER_COMMIT. Any other value passed as an argument is ignored.

Keeping cursors open when a transaction is committed increases the performance of JDBC applications or J2EE components.

1.1.3.5 Support for JDBC 3.0 Retrieval of Auto-Generated Key

During a SQL INSERT command, this feature retrieves the value of the pseudo-column, ROWID, of the newly inserted row.

This feature allows JDBC applications and J2EE components to retrieve the key (ROWID) of the newly inserted row during the same database operation (INSERT). As a result, a network round-trip to the database is saved and the performance is increased.

1.1.4 Oracle HTML DB Improvements

In this release, the Oracle HTML DB user interface has been enhanced and made more consistent to simplify development and deployment. Additional enhancements include:

- The Create Application wizard has been simplified and the generated application made more feature rich.

- The Chart wizard now provides additional chart types, including stacked bar, cluster bar, and dial charts.
- Incorporating web services into your HTML DB application has been made easier through support for UDDI and additional wizards.
- More than ten complete new themes have been added to the built-in themes repository.
- A new Form wizard enables creation of a master detail form without coding.
- Loading of text and spreadsheet data has been enhanced to support larger data sets and loading performance has been improved.
- Tabular forms can now be easily extended to support selecting of rows with check boxes and adding of blank rows.

1.1.4.1 Cluster Bar Charts

SVG charts now include Cluster Bar charts where bars are grouped into columns of data allowing for easy comparison.

The new Cluster Bar charts provide a method of visualizing data in a side-by-side comparison series.

See also:

Oracle HTML DB User's Guide for details

1.1.4.2 Dial Charts

SVG charts now include a dial or gauge chart, allowing you to plot a single value relative to a range defined by a minimum and maximum value represented by a needle on a semicircular surface.

The new Dial chart provides a method of visualizing data commonly used in executive dashboards.

See also:

Oracle HTML DB User's Guide for details

1.1.4.3 Percentage Bar Charts

SVG charts now include a Percent Bar chart used to show cumulative percentages.

Percent Bar charts provide a method of visualizing data that was previously unavailable in Oracle HTML DB.

See also:

Oracle HTML DB User's Guide for details

1.1.4.4 Stacked Bar Charts

SVG charts now include Stacked Bar charts for showing multiple series of absolute values combined into single bars. They are useful when visualizing accumulations or cumulative data.

The addition of Stacked Bar charts allow visualizing of cumulative data in bar charts that was previously not possible.

See also:

Oracle HTML DB User's Guide for details

1.1.4.5 Asynchronous Refresh for All Charts Based on Scalable Vector Graphics (SVG)

SVG charts can now be configured to refresh at an interval defined by the developer. The refresh causes the chart to rerun the query upon which its data set is based.

Every SVG chart used in an HTML DB application can now automatically refresh itself at regular intervals based on the most up-to-date data.

See also:

Oracle HTML DB User's Guide for details

1.1.4.6 Create Master Detail With Choice of Layout

The Master Detail wizard creates a master detail mini-application divided into either 2 or 3 pages with the following features:

- Lost update detection
- Developer-defined navigation through master rows with Next and Previous buttons
- Editing of detail rows on a separate page or on a combined master detail form page
- Single click adding of detail rows in tabular form

The Master Detail wizard automates the manual steps previously required to build a master detail form.

See also:

Oracle HTML DB User's Guide for details

1.1.4.7 Application From a Spreadsheet Wizard

The Create Application wizard has been enhanced in two ways. First, it now enables you to create an application on spreadsheet data. Previously, importing of spreadsheet data had to be done in a separate step using the Data Workshop. This task has now been integrated into the wizard. Secondly, the generated application both on imported spreadsheet data and on an existing table has more features. Those features include:

- Optional read-only version of the generated application
- Record-by-record navigation
- More analysis screens with summary reports and charts

The enhanced Create Application wizard simplifies the process of creating an application based on spreadsheet data and generates a more feature rich application.

See also:

Oracle HTML DB User's Guide for details

1.1.4.8 Create Forms and Reports Based on a Web Service Reference

With the enhanced Web Services wizard, you can create forms that invoke web services and reports based on the results returned by the web service.

The enhancements to the Web Services wizard allow you to incorporate web services in an application without any programming.

See also:

Oracle HTML DB User's Guide for details

1.1.4.9 Discover Web Services Through UDDI Repository

A Universal Description, Discovery, and Integration (UDDI) registry is a directory where businesses register their web services. Through UDDI, the new Web Services wizard can search for registered web services by name.

The UDDI feature of the new Web Services wizard simplifies the process of registering a web service for consumption by an Oracle HTML DB application.

See also:

Oracle HTML DB User's Guide for details

1.1.4.10 Simplified and More Consistent User Interface

The user interface for Oracle HTML DB has been streamlined to improve usability. Specifically, a large number of options have been removed or reorganized to eliminate confusion. Additionally, commonly used tasks such as Run Application or Edit Application have moved to improve consistency.

Oracle HTML DB's simplified and more consistent user interface makes HTML DB easier to navigate and reduces the number of clicks to create and edit applications and components.

See also:

Oracle HTML DB User's Guide for details

1.1.5 PL/SQL Improvements

PL/SQL in Oracle Database 10*g* Release 2 adds these new features:

- Conditional compilation to allow automatic selection of source code according to the version of Oracle Database.
- A new Web Gateway to allow the Oracle instance itself to respond directly to an HTTP request with a response generated by a stored procedure.
- A package to support matrix math.

1.1.5.1 PL/SQL Conditional Compilation

Conditional compilation enables PL/SQL programmers to use new language features in a PL/SQL compilation unit without sacrificing the ability to compile the unit in an older database version that does not support these features. Other typical uses include embedding conditionally compiled debugging and tracing support code in PL/SQL programs.

The benefit is that one body of source code can be deployed in a range of Oracle Database versions to take advantage of new features where they are available and to use fallback approaches where they are not. Debugging and tracing code can be left in place without penalizing the size of the executable or runtime efficiency. Such code can be easily activated by a simple recompilation. This provides a feature in PL/SQL that programmers who use other languages have come to expect.

See also:

Oracle Database PL/SQL User's Guide and Reference for details

1.1.5.2 Obfuscation of Dynamically Generated PL/SQL Source Code

The existing DBMS_DDL supplied package is enhanced with two new subprograms. The command-line WRAP utility takes a plain text CREATE or REPLACE command that specifies creation of a PL/SQL unit and returns a CREATE or REPLACE command where the text of the PL/SQL unit has been obfuscated. The obfuscated representation is identical to what is produced if the same CREATE or REPLACE command is processed with the command-line WRAP utility. The CREATE_WRAPPED command is a shortcut for and has the same semantics as the EXECUTE IMMEDIATE DBMS_DDL.WRAP (input) command.

This feature allows the text of a PL/SQL unit that is created dynamically to be obfuscated. This is critical when an installed product generates new PL/SQL units, for example as part of a customization process, and when the generated code embodies methods that the product vendor considers to be valuable intellectual property. The command-line WRAP utility, available in Oracle Database 10*g* Release 1 and earlier, allowed the code that a vendor shipped to be obfuscated. This new feature allows custom code generated by shipped vendor code to be obfuscated.

See also:

Oracle Database PL/SQL User's Guide and Reference for details

1.1.5.3 Relaxation of Line Length and Overall Limits for the DBMS_OUTPUT PL/SQL Package

The declaration of the package type DBMS_OUTPUT.CHARARR is enhanced to support the increased maximum of 32,767 bytes. Previously, the VARCHAR2 constraint was 255 bytes. This determines the maximum line size that can be written by the DBMS_OUTPUT.PUT and DBMS_OUTPUT.PUT_LINE procedures. The procedure DBMS_OUTPUT.ENABLE (buffer_size integer defaults to 20000) has a new allowed value of NULL for the buffer_size. In this case, NULL means "unlimited." SQL*Plus, through the SET SERVEROUTPUT ON command, enables you to take advantage of the new "unlimited" meaning.

PL/SQL programmers frequently use DBMS_OUTPUT and, in Oracle Database 10*g* Release 1 and earlier, were constrained by the 255 byte limit. When using SQL*Plus, most programmers are regularly caught by the small default overall limit and sometimes by the current 1,000,000 maximum overall limit. In Release 2, the line length limit is increased to 32,767 bytes and the overall limit is removed altogether.

See also:

Oracle Database PL/SQL Packages and Types Reference for details

1.1.5.4 UTL_NLA PL/SQL Package for Matrix Math

C and Fortran APIs for matrix math are freely available. Moreover, efficient platform-specific implementations have been written for many of the most popular platforms. The new UTL_NLA package exposes a matrix PL/SQL data type and wrapper PL/SQL subprograms for two of the most popular of these C APIs; namely BLAS and LAPACK.

To do matrix manipulation in PL/SQL in Oracle Database 10*g* Release 1 and earlier, it was necessary to invent a matrix representation based on PL/SQL native data types and then write matrix manipulation routines from scratch. This required substantial

programming effort and the performance of the resulting implementation was limited. Oracle Database 10*g* Release 2, through the new UTL_NLA package, removes the programming effort and delivers a fast implementation.

See also:

Oracle Database PL/SQL Packages and Types Reference for details

1.1.6 SQL Language Improvements

SQL Language improvements in Oracle Database 10*g* Release 2 include:

- Performance improvements for collection operators.
- Added support for accessing remote LOBs.
- Usability and performance enhancements to Expression Filter.
- Perl compatibility in Regular Expression.
- Rules Manager for Expression Filter can manage and evaluate an unlimited number of rules, provide efficient synchronization of multiple events and threads, and support a single management environment for rules and application data.

1.1.6.1 Rules Manager

Rules Manager is a new feature of Oracle Database 10*g* Release 2. It enables developers to create applications that process and respond to events of any complexity using rules and policies defined in the database. It can evaluate events using data from the application and from database tables. It stores intermediate results to quickly evaluate the next event in a long running composite event (an event made up of two or more simple events). Rules are defined using XML and SQL and can have complex conditions using conjunctions and disjunctions, and specify a set of events, time, and non-occurrence of events with or without a deadline. Event policies control how rules are processed and the duration of an event. Rules can trigger actions that are user-defined procedures running inside an Oracle Database or actions in another application.

The benefit of this feature is that rules that are managed in Oracle Database keep pace with changing business conditions and are always up-to-date; rules are easily changed with SQL and are not included in your application or loaded into a memory-based rules repository. Rules can be evaluated efficiently with the complete business context stored in your Oracle Database and data provided by your application. Event response is flexible; rules can trigger actions in Oracle Database or your application.

Additional advantages of Oracle Database over other approaches include:

- Manageability, by storing rules and event policies with your application data.
- Performance, by evaluating rules and coordinating multiple events and application threads with the full capabilities of Oracle Database.
- Scalability, by evaluating sets of rules of any size.

See also:

Oracle Database Application Developer's Guide - Rules Manager and Expression Filter for details

1.1.6.2 Regular Expression Enhancements

In Oracle Database 10*g* Release 2, Regular Expression is enhanced for Perl compatibility. Support for the common Perl REGEXP extensions that are not included and do not conflict with the POSIX standard are added.

Oracle's SQL Regular Expression provides a simple yet powerful mechanism for rapidly describing patterns and greatly simplifies the way in which you search, extract, format, and otherwise manipulate text in the database.

See also:

Oracle Database Application Developer's Guide - Fundamentals for details

1.1.6.3 Collection Operator Performance Improvements

Collection operators operate on collections of data objects. For Oracle Database 10*g* Release 2, the EQUALITY and MEMBER operators for collections have been improved.

These enhancements have improved the performance for collection operators.

1.1.6.4 Distributed Large Objects (LOB) Support

Distributed LOBs support provides easy-to-use and efficient support for accessing unstructured data in a distributed environment. The data interface for LOBs can now INSERT, UPDATE, and SELECT LOBs across dblinks.

The benefit of this new support is the ability to access remote LOBs.

See also:

Oracle Database Application Developer's Guide - Large Objects for details

1.1.6.5 Expression Filter Enhancements

Expression Filter usability and performance enhancements include:

- Providing extensible optimizer support for the EVALUATE operator.
- Allowing users to specify default values for a subset of attributes in an attribute set.
- Allowing users to specify name spaces in stored XPath expressions.
- Indexing XPath range predicates.
- Allowing users to collect statistics on a set of XPath expressions stored in a table to create an optimal Expression Filter index structure.

Expression Filter features for Oracle Database 10*g* Release 2 include better usability and improved performance. You can specify default values for expression attributes and specify name spaces in XPath expressions.

See also:

Oracle Database Application Developer's Guide - Rules Manager and Expression Filter for details

1.1.7 XML Application Development

New for this release is the Enterprise XML Developer's Kit, which brings full XML support for XMLType in Java, C, and C++. This enables developers to use existing XDK-based applications directly with the database thus improving performance.

Also included for Java is XSLT 2.0 support with XPath functions and operators, bringing new power and efficiency to style sheets.

Finally, application developers can incorporate XML support into applications using the JAXB compiler. Its customization feature enables modification of Java classes to speed development and reduce support costs.

1.1.7.1 XQuery

This feature provides a native implementation of XQuery, the World Wide Web Consortium's (W3C) emerging standard for querying XML, integrated into Oracle Database. It introduces two new functions, XMLQUERY and XMLTABLE.

You can use the XQuery language from an SQL client to query Oracle XML DB repository documents, XML views over relational data, and XMLType data, whether persistent or generated.

XMLTABLE maps the result of an XQuery to relational rows and columns so that results can be queried using SQL.

The benefit of this feature is that Oracle now supports querying of XML, relational, object-relational and repository data using XQuery, the W3C emerging standard for querying XML.

See also:

Oracle XML DB Developer's Guide for details

1.1.7.2 Advanced XPath Rewrite

Oracle Database 10*g* Release 2 extends the XPath Rewrite feature of XML DB to handle additional operations on more complex Schema Constructs. This feature reduces the number of conditions under which XPath Rewrite does not occur. The net effect is that applications using these constructs run faster without any recoding.

This feature allows faster querying and updating of schema-based XMLType data by extending the number of Schema Constructs and Xpath expressions that can be processed by XPath Rewrite.

See also:

Oracle XML DB Developer's Guide for details

1.1.7.3 New InsertXML(), AppendChildXML(), InsertXMLBefore(), and DeleteXML() Functions

The current UpdateXML() function can only manipulate the content of existing nodes within an XML document. It cannot be used to add or remove nodes from a document. The only way to add or remove nodes from a document is to use the methods provided by the DOM API.

This feature resolves this issue by introducing a family of SQL functions similar to UpdateXML() that make it possible to add and remove nodes from a document.

It also enables XPath Rewrite in the cases where the target of the function is a schema-based XMLType with structured storage. This means that operations can be executed in a more efficient manner than would be possible with the DOM API.

This feature improves programmer productivity by providing functions that make it possible to perform complex operations on the content of XML documents directly from SQL. Typically, these functions reduce the number of statements required to perform a particular operation when compared to typical DOM-based operations.

This new set of functions also improves application performance by allowing more complex manipulation of XML content to be performed by XPath Rewrite. XPath Rewrite makes it possible to perform operations on XML documents faster than can be achieved with conventional DOM or SAX-based programming techniques. It is now possible to do in one statement what previously took 10 to manipulate XML eliminating the need to use the DOM API to manipulate XML documents. These new functions improve application performance by allowing complex operations on XML documents to be performed using XPath Rewrite.

See also:

Oracle XML DB Developer's Guide for details

1.1.7.4 Simple Object Access Protocol (SOAP) 1.2 Implementation in C/C++

Implementation of the SOAP 1.2 W3C specification in C and C++.

Developers can build C or C++ SOAP Services to consume and generate SOAP messages using the Oracle XML stack.

See also:

Oracle XML Developer's Kit Programmer's Guide for details

1.1.7.5 SQL/XML Standard (SQL:2003 Standard Part 14) Compliance

This feature extends Oracle's support for the SQL:2003 Standard, adding support for comments, processing instructions, and CDATA and XML prologues. This helps ensure that the SQL/XML operators become the defacto approach to generated XML from a relational query, eventually replacing both XSU and the DBMS_XMLGEN package.

SQL/XML provides an intuitive, standards-based method for creating XML directly from a SQL query. The SQL/XML specification is Part 14 of the SQL:2003 Standard. This feature extends Oracle's support for the SQL/XML operators to bring it in compliance with the SQL:2003 Standard.

1.1.7.6 SQL/XML Standard Compliance

SQL/XML lets you query XML using SQL in a standard way and defines ways of importing and storing XML data in a SQL database, manipulating it within the database, and publishing both XML and conventional SQL data in XML form.

Oracle Database 10*g* Release 2 includes support for all the functions in SQL/XML 2003 (for example, XMLSERIALIZE and XMLPARSE), and some Oracle-proposed extension functions.

See also:

Oracle XML DB Developer's Guide for details

1.1.7.7 Easy XML Application Development Using JAXB Customization

The Java Architecture for XML Binding (JAXB) customization feature permits the XML-generated class names to be replaced with application or Java-specific names to ease manipulation of Java objects making development and maintenance more efficient and straightforward.

Application developers can more easily incorporate XML support into applications using JAXB. The customization feature enables modification of Java classes to speed development and reduce support costs.

See also:

Oracle XML Developer's Kit Programmer's Guide for details

1.1.7.8 Oracle XML Developer's Kit (XDK) PL/SQL Document Object Model (DOM) API Deprecation

This feature completes the work required to make it possible to deprecate the XDK PL/SQL packages and replace them with synonyms equivalent to Oracle XML DB PL/SQL packages without having a negative effect on applications that were originally developed using the XDK implementation.

This feature allows applications developed with the legacy XDK PL/SQL DOM API package (XMLDOM) to leverage the improved performance and functionality offered by the Oracle XML DB PL/SQL DOM API (DBMS_XMLDOM) without requiring significant application changes.

This release completes the deprecation of the XMLDOM, XMLPARSER and XSL_PROCESSOR packages in favor of DBMS_XMLDOM, DBMS_XMLPARSER and DBMS_XSLPROCESSOR.

See also:

Oracle XML DB Developer's Guide for details

1.1.7.9 XPath 2.0 and XQuery 1.0 Functions and Operators Support in Java

The World Wide Web Consortium (W3C) XPath 2.0 and XQuery 1.0 Functions and Operators adds support for the XML Schema data types along with a rich set of functions and operators to XPath that can be used in advance XSLT 2.0 style sheets and XQuerys.

This feature improves the performance and range of applications for style sheets and XQuerys. This reduces the need for compiled code, easing development and maintenance.

See also:

Oracle XML Developer's Kit Programmer's Guide for details

1.1.7.10 Custom Metadata for Oracle XML DB Repository Resources

This feature allows an application developer to define an XML Schema that defines the custom metadata and associate it with a particular XML DB resource. Once the XML Schema has been defined, application developers can create richer content management applications by adding customer metadata to the standard metadata recorded by the Oracle XML DB repository.

This feature makes better use of information contained in documents by making it possible to associate custom metadata with documents stored in the Oracle XML DB

repository. This metadata improves your ability to classify, organize, manage, process, and search the vital corporate information stored in documents.

See also:

Oracle XML DB Developer's Guide for details

1.1.7.11 Secure Sockets Layer (SSL) Support

This feature enables the use of the secure HTTP protocol (HTTPS), in addition to the HTTP protocol, with Oracle XML DB.

SSL adds support for security standards to the Oracle XML protocols servers and allows the HTTPS protocol to be used to secure the transport of information between client and server.

See also:

Oracle XML DB Developer's Guide for details

1.1.7.12 Transportable Tablespace Support

This feature adds new support for XMLType to the Transportable Tablespace capability.

The benefit of this new feature is that it allows the Transportable Tablespace feature of Oracle Database to be used to move tablespaces containing schema-based XMLType tables and columns from one database to another. This method moving XML between different Oracle databases is faster than alternatives such as FTP and Import/Export.

See also:

Oracle Database Administrator's Guide for details

1.1.7.13 Oracle XML DB Administration in the Enterprise Manager Web Console

Enterprise Manager can be used to manage the following Oracle XML DB features:

- XML DB parameters
- XML DB resources
- XML DB Access Control Lists (ACL)
- XML schemas
- XMLType tables and columns

It is now convenient to manage Oracle XML DB in the same place where Oracle Database is managed, which is in the Enterprise Manager Web console.

1.2 Availability

The following sections describe Availability features for Oracle Database 10*g* Release 2.

1.2.1 Fast-Start Failover to Standby Database

In Oracle Database 10*g* Release 2, Oracle Data Guard offers a compelling set of capabilities that provide unparalleled levels of data availability, data protection, and data recovery. For example, on an outage of the primary server, site, or network, the new Fast-Start Failover feature allows Data Guard to rapidly and automatically fail over to the standby database, without requiring manual intervention. This significantly improves the degree of high availability as well as the disaster resilience

for the system infrastructure. This capability makes Data Guard particularly attractive as a high availability/disaster recovery solution for 24x7 applications demanding the highest levels of uptime.

1.2.1.1 Fast-Start Failover

On a failure or outage of the primary server or site, or of the network, Data Guard can now rapidly and automatically fail over to a previously chosen standby database, without requiring manual intervention.

The benefit of fast-start failover is that it increases the degree of high availability, as well the robustness of disaster recovery, for the system infrastructure. It also reduces manual operations, thereby reducing management costs. This makes Data Guard particularly attractive as a high availability/disaster recovery solution for 24x7 applications where it is extremely critical to maintain uptime transparently, despite failures and outages.

See also:

Oracle Data Guard Broker for details

1.2.1.2 Automatic Conversion of the Primary Database to a Standby Database Upon Failover

After a failover, this feature allows the old primary database to be automatically reinstated as a standby database, without requiring re-creation of the primary, or manually performing the Flashback Database operation. This feature is supported only for synchronous redo transports that enabled zero data loss (that is, Maximum Protection and Maximum Availability modes). It also requires Flashback Database to be enabled on both the primary and target standby databases. It is also possible to do this conversion using a single command or a simple push-button interface through Enterprise Manager.

The benefit of this feature is that it enables the Data Guard configuration to easily restore disaster protection in the configuration, without complex manual steps, improving the robustness of the disaster recovery features of Data Guard, as well as improving Data Guard manageability.

See also:

Oracle Data Guard Broker for details

1.2.1.3 Optimized Asynchronous Redo Transmission

Asynchronous redo transmission using the log writer process is no longer limited by the size of the network buffer. The log writer process is no longer blocked on network time outs during asynchronous redo transmission. This allows a greater number of archiver processes to archive redo logs remotely.

The benefit of this feature is that it reduces any effect on the primary database during asynchronous redo transmission. It allows faster redo transmission to standby databases for bulk batch updates on the primary database.

1.2.1.4 Faster Redo Apply Failover

This feature allows the user to activate a physical standby database (that is, fail over to a primary database) that has never been opened read-only, without doing a database restart.

The benefit of faster Redo Apply failover is that it enables customers to recover from a failure or outage much faster, increasing the degree of high availability of the system.

See also:

Oracle Data Guard Concepts and Administration for details

1.2.1.5 Faster SQL Apply Failover

Failover to a logical standby database can now be completed much faster. It is no longer necessary to restart SQL Apply as part of the failover operation.

The benefit of faster SQL Apply failover is that it enables recovery from a failure or outage much faster, increasing the degree of high availability of the system.

See also:

Oracle Data Guard Concepts and Administration for details

1.2.2 Support for Database Transport Across Same Endian Platforms

Oracle Database 10*g* Release 2 introduces database transport across the same Endian platforms which provides a fast and easy way to transport the whole database across platforms with the same Endian format. This feature uses RMAN to quickly convert the datafiles and create the new database. The newly created database is a copy of the source database but on a different operating system/hardware platform; for example, fast Windows to Linux or Solaris to HP-UX migration.

1.2.2.1 Database Transport Across Same Endian Platforms

This transport feature provides a fast method to move a database from one platform to another.

The ability to migrate databases from one platform to another has become more important for Oracle customers. Oracle Data Pump can export the entire database, but it may be time consuming because it uses the SQL layer to process data. Database transport across the same Endian platforms provides a faster and easier way to move Oracle databases from one platform to another.

See also:

Oracle Database Backup and Recovery Advanced User's Guide for details

1.2.3 High Availability Improvements

Oracle Database 10*g* Release 2 continues to enhance the built-in database features for recovery and testing operations to be expeditious, simple, and automatic. Employing Recovery Manager (RMAN) and Flashback Database further reduces the DBA's time spent on the important task of database backup and recovery and creation of test databases. In this release, Data Guard has also optimized the network transmission of redo data, enabling standby databases to be synchronized with the primary database at a faster rate than in previous releases, with minimal effect on the primary database. Similarly, Data Guard has considerably improved its manageability and ease-of-use compared to previous releases, allowing administrators to monitor the state and performance of the disaster recovery configuration at a higher granular level.

1.2.3.1 Add and Create a Transportable Tablespace to a Database

This feature allows a DBA to:

- Generate a tablespace set from the source database and plug it into the target database, essentially automating movement of tablespaces between the databases.
- Generate a tablespace set which can be used to transfer a tablespace to any compatible database.
- Plug-in the tablespace set (provided by the user) into the target database and perform an import operation.

Using this tool, database administrators can automate the process of transporting tablespaces across databases. Tablespaces can be transported between databases running on different platforms using different supported versions. This user-friendly tool performs checks and validations, which are needed to ensure valid and compatible tablespace sets for plugging into the target databases.

1.2.3.2 Additional Data Type Support in LogMiner and SQL Apply

SQL Apply, LogMiner, and Streams now support IOTs with overflow segments and LOBs.

This feature allows Data Guard SQL Apply and LogMiner to be used for protection of a broader range of data types. Similarly, Streams can be used for replication of a broader range of data types.

See also:

Oracle Database Utilities for details

1.2.3.3 Automatic Deletion of Applied Archive Logs

Archived logs, once they are applied on the logical standby database, will be automatically deleted by SQL Apply.

This feature reduces storage consumption on the logical standby database and improves Data Guard manageability.

See also:

Oracle Data Guard Concepts and Administration for details

1.2.3.4 Data Guard: Improved Manageability for Redo Transport, Log Apply, and Broker

This feature deprecates unnecessary initialization parameters and attributes for Oracle Database 10*g* Release 2.

This feature simplifies the usage of initialization parameters and attributes, as well as certain SQL statements.

1.2.3.5 Dynamic RMAN Channel Allocation for Backing Up and Recovering RAC Database

By configuring parallelism when backing up or recovering a RAC database, RMAN channels are dynamically allocated across all RAC instances. This is similar to dynamic RMAN channel allocation for a single instance.

Channel failover allows a failed operation on one node to be continued on another node.

This feature eliminates manual allocation of RMAN channels for each RAC node. It also provides better tolerance of media failures (that is, when one node fails, another node can continue backup work).

See also:

Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide for details

1.2.3.6 RMAN Backup Encryption

RMAN now creates encrypted backups that cannot be restored by unauthorized people. There are 3 modes of backup encryption:

- Transparent encryption
- Password encryption
- Dual-mode encryption using either transparent or password encryption

Backup encryption is a regulatory requirement for many customers. The presence of this capability in Oracle obviates the need to purchase it elsewhere.

See also:

Oracle Database Backup and Recovery Advanced User's Guide for details

1.2.3.7 Easy Conversion of a Physical Standby Database to a Reporting Database

A physical standby database may be activated to be a primary database, opened read/write for reporting purposes, then flashed back in the past to be easily converted back to a physical standby, upon which Data Guard will automatically synchronize the standby database with the primary database. This mechanism also allows a physical standby database to be used as part of a clone, test, or development solution. Note that while this database is open for reporting, it will not be able to receive any redo data from the primary database. Therefore, such a configuration may be more suitable for customers with less demanding disaster recovery requirements.

This feature allows the physical standby database to be utilized for reporting, read/write, and cloning activities, enabling customers to extract more value out of their Data Guard investment.

See also:

Oracle Data Guard Concepts and Administration for details

1.2.3.8 Eliminate Control File Re-Creation

With the control file enhancements, there is no longer a requirement to re-create the control file when changes in the configuration parameters are made. These include the MAXLOGFILE, MAXLOGMEMBERS, MAXLOGHISTORY, MAXDATAFILES, and MAXINSTANCES parameters.

This feature eliminates database downtime when a configuration parameter change is made in the control file.

See also:

Oracle Database Administrator's Guide for details

1.2.3.9 Enhanced Memory Corruption Checking

A new value of MAX has been added to the database parameter, DB_BLOCK_CHECKSUM, to compute checksums on the reading, updating, and deleting of a block.

This feature enhances the ability to validate the block when it is in memory and not just on disk.

1.2.3.10 Enhanced RMAN Backup Job Views

The new RMAN views correlate and simplify the metadata of an RMAN backup job. For example, a backup job can be associated with its backup sets, and a backup set with its composite control file, archived logs, or datafiles. Additional details such as I/O file sizes and compression ratio can be found in the new views.

This features enhances the manageability of RMAN backups by providing a SQL interface to find additional statistics on the backup job and files.

1.2.3.11 Enhanced Server Side Shutdown

When using the SHUTDOWN ABORT command with Oracle Database 10*g* Release 2 and higher clients, all processes (including, for Unix, the foreground process) will be terminated.

Also, when issuing the STARTUP FORCE command, in previous versions, "Starting ORACLE instance (FORCE)" was written to the alert log. In Oracle Database 10*g* Release 2 and higher, "Shutting down instance (ABORT)" followed by "Starting ORACLE instance (NORMAL)" are now written to the alert log.

Transparently handled both in the server and the client, all processes will be cleaned up after a SHUTDOWN is issued. This allows a DBA to startup in the same session.

1.2.3.12 Enhanced Space Management Verification

Additional checking has been added to the DBMS_SPACE_ADMIN package to provide better tools to debug space management problems. Extended syntax is added to the ALTER SYSTEM command to dump the space headers of locally managed tablespace files for added analysis.

This feature enhances the tools used to debug space management-related problems to detect corrupt metadata.

See also:

Oracle Database Administrator's Guide for details

1.2.3.13 Failover Improvements for Real Application Clusters (RAC) Distributed Transactions

Distributed Transactions in a RAC environment detect failures and start the failover and failback processes automatically.

The benefit of this feature is reduced administration and development time because Oracle manages the distributed transaction, therefore, there is no need to write manual procedures.

See also:

Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide for details

1.2.3.14 Faster Instance Startup for Large Memory Configurations

This feature decreases the time it takes to open Oracle Database instance.

Accelerated instance startup makes Oracle Database available sooner when databases have multigigabyte shared memory configurations.

1.2.3.15 Fine-Grained, Automated Tracking of SQL Apply Runtime Performance

Several additional views have been provided, and existing views have been enhanced to enable administrators to dynamically track the progress and performance of SQL Apply at a granular level. Enterprise Manager also provides historical trend analysis of the progress and performance data.

This feature makes it considerably easier to maintain and tune SQL Apply configurations and improves Data Guard manageability.

See also:

Oracle Data Guard Concepts and Administration for details

1.2.3.16 Flashback Across Data Guard Switchovers

It is now possible to flashback the primary and standby databases to an SCN or a point in time prior to the switchover operation. Using flashback in this way on a physical standby database preserves the standby role. Using flashback in this way on a logical standby database changes the role of the standby database to what it was at the target SCN/time.

This extends the flashback window; that is, the possible SCN/time in the past to which any database can be flashed back to. This provides administrators more flexibility to detect and correct human error situations.

See also:

Oracle Data Guard Concepts and Administration for details

1.2.3.17 Flashback Database Through Resetlogs

This feature supports flexible data recovery that enables movement of the database back and forth in time even after you perform a resetlogs operation.

If you ever need to do a resetlogs operation for recovery purposes, you can take the database back to a point in time before the resetlogs if you find that you made a mistake. This provides administrators more flexibility to detect and correct human error situations.

See also:

Oracle Database Backup and Recovery Basics for details

1.2.3.18 Hardware Assisted Resilient Data (HARD) Compliant RMAN Backups

RMAN backup pieces can be stored on and validated by HARD-compliant disk hardware.

This prevents corruptions in RMAN backup pieces from being written to permanent storage.

1.2.3.19 Increase Supported Real Application Clusters Instances for a Database

For Oracle Database 10*g* Release 2, the control file supports 1,055 instances for a RAC database.

This feature provides enhanced support and flexibility for Grid Management.

1.2.3.20 Incremental Roll Forward of Database Copy

This feature can be used to create an RMAN incremental backup from a specified SCN. This incremental backup can be used to quickly move a standby database forward in time or load new data into a clone database for testing.

This feature minimizes the time data is unprotected by creating an incremental backup that can be applied to a standby database to quickly make the standby database transactionally consistent with the primary database. This also eliminates the need to continuously re-create a clone database with newer data values.

See also:

Oracle Database Backup and Recovery Advanced User's Guide for details

1.2.3.21 Online Redefinition Supports Clustered Tables, ADT's, MV Logs, AQ Tables, and Preserves Database Statistics

Online redefinition can now be performed on clustered tables, tables containing Abstract Data Types (ADT), Advanced Queuing (AQ) tables, Materialized View (MV) logs, and it retains all statistics on a table. It also clones statistics, checks and non-NULL constraints, and dependent objects on nested tables. PL/SQL packages do not have to be recompiled if the number and ordering of columns and data types in the table remain unchanged.

This feature means less downtime as more tables are now supported for online redefinition. Analysis of the table after reorganization is no longer required and PL/SQL packages remain available following an online redefinition, thus maintaining application availability. These enhancements are especially beneficial to customers using Oracle Applications.

See also:

Oracle Database Administrator's Guide for details

1.2.3.22 Online Redefinition of a Single Partition

Individual partitions may now be reorganized online.

With this feature, system downtime is reduced and, during reorganization, storage requirements and rollback space requirements are significantly reduced.

See also:

Oracle Database Administrator's Guide for details

1.2.3.23 Optimized Creation of Logical Standby Database

Creation of a logical standby database no longer requires the creation of a specialized logical standby control file, which could not be used by RMAN. Logical standby databases can now be created easily from a physical standby database.

This feature reduces specialized manual operations for creating a logical standby database and improves Data Guard manageability.

See also:

Oracle Data Guard Concepts and Administration for details

1.2.3.24 Unused Block Compression

By optimizing database backups, only blocks currently used by the database are backed up. In previous releases, NULL compression was utilized, whereby only never-used blocks were excluded from backup. With this new feature, all unused blocks, whether they have been used or not in the past, will be excluded from backup.

This enables faster backups and smaller backup files.

See also:

Oracle Database Backup and Recovery Basics for details

1.2.3.25 Reorganized Object Types Improves User Interface

This feature supports reorganization of additional database object types that are currently not supported (for example, nested tables), enhances capabilities like launch-in-context, and improves usability of the Re-org wizard.

Currently if a tablespace contains objects like nested tables, the tablespace cannot be reorganized. This new feature enhances support for more object types for tablespace reorganization. In addition, this feature provides enhanced support for launch-in-context so that applications such as Policy Violations can pass recommendations to the wizard directly and do not depend on the user to enter the data manually.

1.2.3.26 Restore Point

This feature provides the ability to associate a user-defined name, Restore Point, with a point-in-time. The Restore Point can be used with a Flashback Database or Flashback Table during recovery operations. Restore Point can be specified such that it guarantees the database can be recovered to a particular point-in-time.

Restore Point eliminates the need to manually record an SCN or timestamp to use for Flashback Database and Flashback Table operations. In addition, a special 'guaranteed' restore point ensures that the database can be flashed back to that particular point-in-time. Using guaranteed restore points can achieve efficient disk space usage, as only those flashback logs necessary to meet the guaranteed restore point are required to be retained.

Restore points can also be created as 'guaranteed', meaning it provides enhanced Flashback Database efficiency in recovery and disk storage requirements. And, it guarantees the database can be rewound to a specified point-in-time.

See also:

Oracle Database Backup and Recovery Basics for details

1.2.3.27 Temporary Datafiles Are Re-Created on RMAN Recovery

Temporary datafiles that belong to locally managed temporary tablespaces are automatically re-created during recovery operation.

This eliminates the need to manually create temporary tablespaces after recovery.

See also:

Oracle Database Backup and Recovery Basics for details

1.3 Business Intelligence

The following sections describe new features for Oracle Database 10*g* Release 2 in the Business Intelligence area.

1.3.1 SQL Support for Analytic Applications

In this release, Oracle extends the SQL language to support more powerful analytics. The MODEL clause of the SELECT statement, which provides a powerful new approach to analytic calculations, adds to its expressiveness and scalability in this release. Bitmap indexes, a vital structure for efficient Business Intelligence query processing, also gain enhanced scalability.

1.3.1.1 Enhanced Upsert Abilities in MODEL Clause

This feature provides new behavior supporting upsert of cells. Rules are now able to insert new cells even when the left hand side includes existential predicates such as ANY, IN, >, <, and LIKE.

The benefit is more expressive rules and greater productivity in specifying the MODEL clause.

See also:

Oracle Database Data Warehousing Guide for details

1.3.1.2 Enhancements to MODEL Rules Which Use FOR Loops

This feature moves evaluation of MODEL clause rules using FOR loops from compile time to runtime. This permits rules with FOR loops to use reference spreadsheets, nested references, and the ITERATION_NUMBER variable. In certain situations, subqueries within FOR loops can have shared cursors and refer to tables defined in the WITH clause.

The benefit of this feature is that MODEL clauses with certain types of rules perform better, handle larger model sizes, and include more powerful rule expressions.

See also:

Oracle Database Data Warehousing Guide for details

1.3.1.3 Use of Analytic Functions in the SQL MODEL Clause

This feature provides direct use of analytic functions like moving average and cumulative sums in SQL MODEL clause rules.

Native support of analytic functions is significantly faster than alternative techniques and simplifies expression of complex calculations, such as allocations and forecasts, in the MODEL clause.

See also:

Oracle Database Data Warehousing Guide for details

1.3.1.4 Intra-Partition DML in the Presence of Bitmap Indexes

Intra-Partition DML (PDML) can now be fully leveraged in the presence of bitmap indexes.

With this feature in Oracle Database 10*g* Release 2, the last main restriction for PDML is lifted.

1.3.2 Improved Very Large Database (VLDB) Support

Oracle is designed to meet all data requirements for the largest databases. VLDB environments, such as data warehouses, depend on partitioned tables. This release enables more:

- Partitions for each table
- More efficient partition management
- More partition-based query optimization

1.3.2.1 Increased Maximum Number of Partitions Per Object

The maximum number of partitions is now 1024K-1, increased from a previous limit of 64K-1.

Increasing the maximum number of possible partitions to objects enables finer granularities for partitioned objects and even broader flexibility in choosing the right partitioning approach for a business scenario.

See also:

Oracle Database Administrator's Guide for details

1.3.2.2 Enhanced Dynamic Partition Pruning

This feature enhances partition pruning for complex queries.

This feature provides enhanced performance for a broader set of complex queries.

See also:

Oracle Database Data Warehousing Guide for details

1.3.2.3 Resource Optimized DROP TABLE for Partitioned Tables

When dropping a large partitioned table in the nonrecoverable PURGE mode, the DROP operation is internally split to drop chunks of partitions.

Dropping large partitioned tables can affect tens of thousands of partitions that all have to be logically removed from the system. The capability of transparently dropping such an object in an incremental fashion optimizes the resource consumption and positively affects the run-time behavior.

See also:

Oracle Database Administrator's Guide for details

1.3.3 Reduced Information Cycle Time

A key requirement for business intelligence is to be able to analyze data more rapidly after an event occurs. Oracle provides a wide range of features to support efficient ETL (extraction, transformation, load). This release reduces the information cycle time with enhancements ranging from data loading to query processing. New features minimize rollbacks during bulk DML, enable more Change Data Capture environments, simplify materialized view refresh, and enhance query rewrite.

1.3.3.1 Distributed Asynchronous Change Data Capture (CDC)

The internal Change Data Capture (CDC) tasks of Streams Capture and Apply are decoupled, thereby enabling a heterogeneous CDC setup with different operating systems and Oracle versions. The propagation between source and target database is a

pure Streams implementation, enabling asynchronous CDC to leverage any existing Oracle9*i* Release 2 system as a source.

Asynchronous Change Data Capture (CDC) no longer requires the same operating system for source and target. Furthermore, distributed asynchronous CDC capabilities expands the reach of CDC source systems to include Oracle9*i* Release 2.

See also:

Oracle Database Data Warehousing Guide for details

1.3.3.2 DML Error Logging Table

This feature allows bulk DML operations to continue processing, when a DML error occurs, with the ability to log the errors in a DML error logging table. On a DML statement level, you specify whether to log errors or not, the level of detail information, as well as a maximum error threshold, similar to external table and SQL*Loader.

This feature combines the power and speed of bulk processing with the functionality of row level error handling.

See also:

Oracle Database Concepts for details

1.3.3.3 Partition Change Tracking Refresh Without Materialized View Logs

Partition Change Tracking Refresh (PCT) no longer requires materialized view logs. It can be combined with conventional fast refresh to enable a refresh, where one table is using PCT and the other table uses materialized view logs.

This feature improves performance because materialized views can now be created in situations where they were previously not used due to the overhead of materialized view logs. This feature also provides faster refresh due to the mix of PCT and conventional refresh.

See also:

Oracle Database Data Warehousing Guide for details

1.3.3.4 Query Rewrite Using Multiple Materialized Views

Query Rewrite may now use multiple materialized views.

This feature improves performance as more queries are eligible for query rewrite and reduced maintenance costs because fewer materialized views are required.

See also:

Oracle Database Data Warehousing Guide for details

1.3.3.5 Skip Level Dimension Hierarchies

Levels in the dimension can now be skipped. For example, this feature allows a geography dimension where the state level might be skipped for cities such as Washington, DC.

The benefit is more complete modeling of the dimensions and performance benefits, because improved query optimizations are now possible.

See also:

Oracle Database Data Warehousing Guide for details

1.3.4 Online Analytical Processing (OLAP) Support for Analytic Applications

Oracle OLAP provides a complete OLAP server that is fully integrated into Oracle Database. In this release, Oracle OLAP makes cube aggregation more flexible and enhances the OLAP API, enabling an even broader set of analytic applications than in earlier versions. Optimized OLAP cube refresh supports applications demanding the very latest data.

1.3.4.1 Incremental Aggregation of Compressed OLAP Cubes

Incremental aggregation of compressed cubes occurs automatically within the multidimensional engine. The multidimensional engine tracks detail level changes to the cube and aggregates only the summary level data that would be affected by the newly loaded data.

This new feature improves performance for the process of aggregating an OLAP cube that has been stored using cube compression technology. This performance improvement is seen when the compressed cube has been incrementally updated with new data (for example, when adding a new time period to the cube).

See also:

Oracle OLAP DML Reference for details

1.3.4.2 OLAP API Supports Calculated (Virtual) Dimension Members

Within the OLAP API model, application developers can define calculated (virtual) dimension members. For example, a calculated dimension member such as 'New England' in a geography dimension could be created as the sum of stored dimension members 'ME + VT + NH + MA + RI + CT'.

Calculated member support in the OLAP API allows application developers using the OLAP API to create calculated (virtual) dimension members within an OLAP API query. This extends the analytic capabilities of applications that use the OLAP API.

See also:

Oracle OLAP Developer's Guide to the OLAP API for details

1.3.4.3 OLAP Compressed Cubes Support Semiadditive Aggregation Methods

Semiadditive aggregations are those where the aggregation method on one or more dimensions is something other than sum. Examples include first, last, hierarchical sum, and hierachical weighted average. In Oracle Database 10*g* Release 1, compressed cubes could not be used with semiadditive aggregation. In Oracle Database 10*g* Release 2, semiadditive measures can be used with compressed OLAP cubes.

Support for semiadditive aggregation methods with compressed cubes provides the opportunity for significant aggregation performance improvements.

See also:

Oracle OLAP DML Reference for details

1.3.5 Data Mining Support for Analytic Applications

Oracle Data Mining provides integrated mining capabilities directly in Oracle Database. Oracle enriches its data mining for the full spectrum of application developers with enhancements targeted to Java specialists, PL/SQL developers, and data mining newcomers. Highlights include decision tree models, simplified model development, and a new SQL function for predictive data mining.

1.3.5.1 Decision Tree Algorithm

This feature implements a state-of-the-art decision tree algorithm leveraging in-database parallelism.

The benefit of this feature is a transparent, rule-producing algorithm. The presentation of a decision tree in the form of human-interpretable rules increases the confidence, acceptance, and value of the algorithm. The decision tree is optimized to allow fast and scalable data scoring.

See also:

Oracle Data Mining Application Developer's Guide for details

1.3.5.2 Predictive Analytics Package

The Predictive Analytics Package provides an easy-to-use PL/SQL interface enabling straightforward automation of the data-mining process.

The Predictive Analytics Package makes reasonable assumptions about the data and the mining process and greatly simplifies the use of data mining for end users. This feature makes data mining more accessible to non-data analysts.

See also:

Oracle Data Mining Concepts for details

1.3.5.3 SQL-Level Data Mining Scoring Capability

This feature provides SQL built-in data mining scoring functions for classification, regression, clustering, and feature extraction data mining models.

This capability allows customers and applications to:

- Ease deployment of data mining models within the context of existing applications
- Dramatically improve scoring performance
- Enable pipelining of results, which can accelerate the presentation of findings to end users

See also:

Oracle Data Mining Application Developer's Guide for details

1.3.5.4 Java Data Mining (JDM) Compliant Java API

Oracle Database 10*g* Release 2 provides a Java Data Mining (JDM) JSR-73 compliant Java API. Oracle Data Mining 10*g* Release 2 requires a Java API implemented on top of the 10.1 PL/SQL API. This unifies the overall product, enabling interoperability between models.

This feature provides vendor neutral, open standard data-mining application development.

See also:

Oracle Data Mining Application Developer's Guide for details

1.4 Clustering

The following sections describe new features for Oracle Database 10*g* Release 2 in the clustering area.

1.4.1 Real Application Clusters (RAC) Performance

This release provides session management capabilities that improve the ability to balance workloads across instances in a clustered environment. Workload policies are defined which enable connection pools to route work requests to the appropriate instance based on real-time information.

1.4.1.1 Load Balancing Advisory

Real Application Clusters includes a Load Balancing Advisory that monitors the work for each service in an instance. The advisory creates Fast Application Notification (FAN) events to provide feedback that enables routing by the connection pool to respond quickly to changing conditions in the system. In steady state, the system approaches equilibrium with optimal service times across all Oracle Database 10*g* RAC instances. JDBC and ODP.NET Connection Pools are integrated to provide the best connection to an application request instead of a random connection from the pool. Oracle has integrated its connection pool clients (OCI, ODP.NET, and JDBC) with the Load Balancing Advisory. Runtime connection load balancing enables connection pools to route work requests to the connections that can best serve the work based on the FAN events received from RAC. Requests can be routed according to different policies.

The benefit is improved performance and use of resources.

See also:

Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide for details

1.4.1.2 Oracle Cluster File System for Solaris

Oracle Cluster File System (OCFS) will be made available for Sun Solaris environments. This allows for easy management of data that needs to be shared across nodes in a clustered environment. OCFS is designed to support all Oracle Database files and ORACLE_HOME binaries. Note that OCFS for Sun Solaris environments will be released after the release of Oracle Database 10*g* Release 2.

This feature provides enhanced cluster manageability for clusters in Sun Solaris environments.

1.4.2 Real Application Clusters Robustness

This release improves the RAC experience by verifying the environment before installation, providing tools to assist in diagnosing problems that may occur, and ensuring there are redundant copies of Clusterware Files.

1.4.2.1 Clusterware File Redundancy

There are 2 files, the Oracle Cluster Registry (OCR) and the Voting Disk, that are critical to running a cluster using Oracle Clusterware. Oracle Clusterware

automatically maintains two copies of the Oracle Cluster Registry (OCR) file and three or more copies of the Voting Disk files. Mechanisms are also provided for the online replacement of failed OCR devices.

The benefit is that redundant copies of critical clusterware management files are automatically maintained to protect against single points of failure and eliminate the need for third-party storage redundancy solutions.

See also:

Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide for details

1.4.2.2 Cluster Verification

Oracle Clusterware provides a Cluster Verification Utility that verifies the availability and integrity of a wide range of cluster elements in Oracle Clusters and Oracle Database 10*g* RAC environments. These facilities include both a standalone Cluster Verification Tool for use by DBAs and callable routines invoked automatically by Oracle installation, configuration, and system management tools. The Cluster Verification Utility can be invoked at any stage of a cluster configuration.

Cluster verification makes the installation, configuration, and overall management of Oracle Clusterware environments including Oracle Database 10*g* RAC environments simpler by quickly and easily identifying problems in complex cluster systems.

See also:

Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide for details

1.4.2.3 Cluster Tracing and Diagnostic Enhancements

New facilities in Oracle Database 10*g* Release 2 provide improved diagnostics of cluster wait/hang conditions and more effective logging of Oracle Clusterware error/alert messages. In addition, diagnostic information for faster debugging and problem analysis is provided.

The benefit of these enhancements is greater code reliability and faster problem resolution by Oracle Support Services and the development community.

1.4.3 Application High Availability With Oracle Clusterware

This release extends Oracle Clusterware by providing an open API which allows additional (non-Oracle) components running on the database server to be managed by Oracle Clusterware.

1.4.3.1 High Availability API

An open API is provided to Oracle Clusterware to enable customers and partners to integrate their application system components into Oracle high availability environments. Application system components registered with Oracle Clusterware are managed along with Oracle system components to ensure that failures of any component can be detected and recovery procedures initiated to quickly restore the overall system to an operational state.

This feature ensures that complex cluster systems involving Oracle components and components from customers and partners can be managed together by Oracle Clusterware to provide high overall system availability.

See also:

Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide for details

1.5 Content Management

The following sections describe Content Management features for Oracle Database 10*g* Release 2.

1.5.1 Oracle Text Improvements

Improvements for Text include a new set of document filters that provide better quality and performance, Asian language support for CTXRULE index type, and better manageability of text indexes. Also, the new KWIC feature (also called snippet) returns query keywords marked up in their surrounding text, allowing the user to evaluate them in context.

1.5.1.1 CTXRULE Asian Language Support

Enhancements in this release to CTXRULE enable rule-based classification on Japanese, Chinese, and Korean languages.

CTXRULE is a filtering model for Asian language search and rules and enables rule-based classification.

See also:

Oracle Text Application Developer's Guide for details

1.5.1.2 Filtering

A new AUTO_FILTER filter has been added to Oracle Text in this release. This new filter takes advantage of the automatic filtering technology licensed from Verity, Inc.

Oracle Text's automatic filtering technology, licensed from Verity, Inc., enables you to index most document formats. This technology also enables you to convert documents to HTML for document presentation.

See also:

Oracle Text Reference for details

1.5.1.3 Key Word in Context (KWIC)

The Key Word in Context (KWIC) feature has been enhanced in this release with two new procedures, CTX_DOC.SNIPPET and CTX_DOC.POLICY_SNIPPET, that return text fragments containing keywords found in documents.

These new procedures enable you to see the keywords in their surrounding text, allowing you to evaluate them in context.

See also:

Oracle Text Reference for details

1.5.1.4 Online Indexing for Local Partitioned Index

With online indexing for local partitioned indexes, you can now create an index on a partition while updating it; the same as with unpartitioned indexes.

This features improved the manageability and scalability for text indexing.

See also:

Oracle Text Reference for details

1.5.2 *inter*Media Improvements

*inter*Media now includes support for the Digital Imaging and Communications in Medicine (DICOM) standard, the most common medical imaging format. Also provided are:

- APIs that extract embedded content metadata in IPTC, EXIF, and XMP formats from a number of popular image formats and return this metadata as a collection of XML documents.
- APIs that embed XMP formatted metadata into images.

1.5.2.1 *inter*Media DICOM Medical Imaging Format Support

*inter*Media now supports the most common medical imaging format, DICOM version 3. Applications can now use *inter*Media Java and PL/SQL APIs to extract metadata about patients, physicians, diagnoses, treatments, tests and procedures, and other relevant information included in the DICOM format. Metadata extracted from DICOM images is represented as an XML document returned in an XMLType object. Each returned document conforms to an XML schema that is registered with the database. Once this metadata has been extracted from the image, it can be stored in an Oracle Database, indexed, searched, and made available to applications using the standard mechanisms of Oracle Database. Not all metadata is extracted, but metadata generally present in all DICOM images will be extracted.

This new *inter*Media capability allows Oracle Database 10*g* Release 2 to serve as the repository for shared multiterabyte medical imaging archives used by HIMMS and PACS systems, life sciences and medical research centers, and government agencies.

See also:

Oracle interMedia User's Guide for details

1.5.2.2 *inter*Media Support for Image Content Metadata and XML DB Integration

In addition to the current support for format metadata (height, width, mimetype, etc.), *inter*Media now includes Java and PL/SQL APIs that extract embedded content metadata (in IPTC or IIM, EXIF, and XMP format) from TIFF, GIF, and JPEG image formats. Information supplied by digital cameras, image editing software programs, news and wire services, such as photographer, location, date, time, and other application data can now be extracted and represented as a collection of XML documents returned in XMLType. Once metadata is extracted, it can be stored in an Oracle Database, indexed, searched and made available to applications using the standard mechanisms of Oracle Database.

*inter*Media Java and PL/SQL methods also allow any Adobe XMP (tm) formatted, user-defined metadata to be embedded into TIFF, GIF, and JPEG formats.

By placing the metadata and the image data in the same containing binary image, both types of data can be shared and exchanged reliably as a unit.

See also:

Oracle interMedia User's Guide for details

1.5.2.3 Java Advanced Imaging 1.1.2 Support

Java Advanced Imaging 1.1.2 provides support for the new release of Sun's Java Advanced Imaging Package (JAI).

Oracle supports the latest production release of the Sun Microsystem Java Advanced Imaging (JAI) package for *inter*Media applications that use Oracle to store and process image content.

1.6 Database Overall

The following sections describe new features for Oracle Database 10*g* Release 2 for the database in general.

1.6.1 Database Utilities

The following sections describe new Database Utilities features.

1.6.1.1 Easier Manual Upgrades

Manual database upgrades are now easier to do. For any database version being upgraded to Oracle Database 10*g* Release 2, only one common SQL script has to be invoked when performing a database upgrade. Specifying a version-specific script is no longer required, because Oracle now automatically determines what version is being upgraded and runs the appropriate upgrade scripts for that database and all of its included components.

The advantage is that less time is required to do manual database upgrades.

See also:

Oracle Database Upgrade Guide for details

1.6.1.2 Enhanced Pre-Upgrade Information Utility

The Pre-Upgrade Information Utility has been enhanced to provide improved resource estimations for tablespace space usage and elapsed upgrade runtime.

This feature improves the manageability of Data Pump jobs and allows users to get accurate information about possible issues with their specific database upgrade prior to running DBUA or performing a manual upgrade.

1.6.1.3 Default DATA_PUMP_DIR Directory Object

This feature creates a default DATA_PUMP_DIR directory object. This simplifies the use of Data Pump, because the DBA does not have to manually create a directory object before trying to use Data Pump Export.

This allows users to quickly start using Data Pump because a default directory object is automatically created.

See also:

Oracle Database Utilities for details

1.6.1.4 Dumpfile Compression

Data Pump Export dump files can contain database metadata objects in a compressed format, allowing dump files to consume less operating system disk space. This metadata is automatically uncompressed during Import.

This feature reduces the amount of disk space consumed by dump files.

See also:

Oracle Database Utilities for details

1.6.1.5 Extended Job Status Display

This feature has extended the job status display so that users are able to see the degree of parallelism used by the Data Pump Workers in Export/Import jobs.

This provides improved manageability of Data Pump jobs.

1.6.1.6 FILESIZE Value for Dumpfiles

This feature allows users to indicate the size of the dumpfile being created when using the Export command line interactive ADD_FILE command by specifying a FILESIZE value. In previous releases, the size of the file was unlimited.

With this feature, if there are storage restrictions, users can control the size of dumpfiles.

1.6.1.7 Monitoring of Data Pump Jobs Started Outside of Enterprise Manager

This new feature helps to:

- Monitor Data Pump jobs started outside of Enterprise Manager
- Support Data Pump Export from a read-only database
- Move to new server APIs

In previous releases, Enterprise Manager only monitored Data Pump jobs that had been created using Enterprise Manager. However, it is a common activity for a DBA to create Data Pump jobs without using Enterprise Manager and it is possible for these jobs to take a very long time (hours or days) to complete. Enterprise Manager is a database monitoring tool and has been enhanced to monitor all Data Pump jobs, whether created by Enterprise Manager or not.

1.6.1.8 New DBMS_DATAPUMP API Calls

The WAIT_FOR_JOB API provides the simplest mechanism for waiting for the completion of a Data Pump job. The GET_DUMPFILE_INFO API analyzes a specified dumpfile and returns dumpfile information such as version, creation date, and character set.

This new features improves the manageability of Data Pump jobs. You can get more information about Data Pump job completion and specific dumpfiles.

1.6.1.9 New PCTSPACE Transform Parameter on Data Pump Import

With this feature, you can reduce the amount of space required for tablespaces by shrinking tablespace storage allocation. Large datafiles in production systems can be reduced so that smaller datafiles are created during Import on test systems.

The benefit of this feature is the ability to reduce the amount of tablespace storage required for datafiles.

1.6.1.10 Object ID (OID) Transform Parameter

With this enhancement to Data Pump Import, you can suppress the generation of the Export OID clause for object types using the new OID transform parameter.

This is useful for duplicating schemas across databases, using Export and Import, when you cannot guarantee that object types will have identical OID values in those databases.

1.6.1.11 SQL*Plus AUTOTRACE EXPLAIN Conversion to DBMS_XPLAN

SQL*Plus SET AUTOTRACE EXPLAIN PLAN_TABLE queries are replaced by a call to the DBMS_XPLAN package.

This feature provides a single Oracle standard for EXPLAIN PLAN_TABLE output.

See also:

*SQL*Plus User's Guide and Reference* for details

1.6.1.12 SQL*Plus DBMS_OUTPUT Buffer and Line Size Limit Increase

In Oracle Database 10*g* Release 2, the SQL*Plus SET SERVEROUTPUT command uses the newly increased line size (32767 bytes) and unlimited number of lines of DBMS_OUTPUT.

The benefit of this feature is increased usability of SQL*Plus.

See also:

*SQL*Plus User's Guide and Reference* for details

1.6.1.13 SQL*Plus XQuery Support

You can enter XQuery 1.0 (an XML Query Language) statements and retrieve results in SQL*Plus.

This feature allows XQueys to be created and tested in a highly available tool (SQL*Plus).

See also:

*SQL*Plus User's Guide and Reference* for details

1.7 Database Security

The following sections describe new features for Oracle Database 10*g* Release 2 in the Security and Directory areas.

1.7.1 Privacy and Confidentiality

Privacy, confidentiality, and the protection of personally identifiable information (PII) is a global concern in today's economy. Oracle provides a wealth of server-enforced protections for private and confidential information. Oracle Database 10*g* Release 2 introduces Transparent Data Encryption that quickly and easily encrypts sensitive information. Existing applications generally continue working without any modification because the data is transparently decrypted when accessed through SQL.

1.7.1.1 Transparent Data Encryption

Oracle uses industry standard encryption algorithms including AES and 3DES to encrypt columns that have been marked to be encrypted. Key management is handled by the Database.

This feature provides transparent data encryption and key management. Businesses can protect sensitive data in Oracle databases without any changes to their application.

SQL interfaces to Key Management hide the complexity of encryption. This feature mitigates liability due to storage theft.

See also:

Oracle Database Advanced Security Administrator's Guide for details

1.7.1.2 Key Management for Data Encryption

This feature provides a user interface to the transparent data encryption feature.

The benefit is transparent key management for data encryption.

See also:

Oracle Database Advanced Security Administrator's Guide for details

1.7.2 Improved Manageability

The ability to focus more on security policies and less on low level implementation details is key to securing the enterprise. Oracle Database 10*g* Release 2 Transparent Data Encryption manages the encryption keys for the application administrator, simplifying the implementation and management of encrypted data. In addition, Oracle Database 10*g* Release 2 auditing enhancements include writing audit data in XML format to the operating system and providing additional protections for audit data at the operating system level.

1.7.2.1 Enhanced Proxy Authentication

Proxy authentication in Oracle Database has always been available using OCI or JDBC programming. This feature extends this capability to the CONNECT command and enhances the current OCI/JDBC interface. For example, an application user/schema could be managed as a directory user, improving manageability of application schemas across the enterprise.

This feature extends proxy authentication capability to tools such as SQL*Plus which are used by applications for running batch jobs. Application security is improved as the application schema password is not exposed at connection time. It also enhances the existing proxy capability by allowing proxy users to be directory users.

See also:

Oracle Database Enterprise User Administrator's Guide for details

1.7.2.2 Audit Trail in Standard XML Format

Audit trail can be written to the operating system files in standard XML format. These logs can also be viewed using dynamic views providing better access to information in the operating system audit logs.

This feature enhances database operating system audit capability to support XML audit trail format. This allows the DBA to build custom tools to trace abnormal activities and view and mine audit logs. It also enhances the ability of the DBA to view operating system audit information when authenticated to the database.

See also:

Oracle Database Security Guide for details

1.7.3 Secure By Default

A major objective for Oracle Database 10*g* Release 2 is to ensure you can easily install and configure Oracle Database in a hardened configuration. This includes locking default accounts, providing database roles for privilege management, and protection of sensitive information. Oracle Database 10*g* Release 2 provides additional security for passwords used in fixed database links. In addition, the much used CONNECT role has been modified to include only those privileges required to connect to the database, namely the CREATE SESSION privilege.

1.7.3.1 Connect Role Privilege Reduction

The connect role privilege reduction feature reduces the number of privileges granted to the connect role to one, the CREATE SESSION privilege. The privileges have been removed from the connect role:

- CREATE CLUSTER
- CREATE DATABASE LINK
- CREATE SEQUENCE
- ALTER SESSION
- CREATE SYNONYM
- CREATE TABLE
- CREATE VIEW

This feature helps you deploy secure configurations by helping enforce the least privileged principle.

1.7.4 Improved Integration and Interoperability

The ability to leverage security with Oracle features, such as Real Application Clusters, is key to deploying secure and scalable solutions. Oracle Label Security has been enhanced in Oracle Database 10*g* Release 2 to support Oracle Real Application Clusters. This provides customers additional capabilities for protecting sensitive information.

In addition, Oracle Wallets have been enhanced to provide the ability to store username and passwords for command line operations. For example, batch jobs using Oracle SQL*Plus can use a wallet to connect to the database versus having to supply a username and password on the command line.

In addition, Oracle Database 10*g* Release 2 introduces a new command line proxy capability that allows authorized users to proxy to different accounts upon authentication to the database.

1.7.4.1 Secure External Password Store

Oracle Wallet functionality has been extended beyond PKI credential store to include username and password storage. This enables a user to store his or her username and password in an Oracle Wallet and have the password picked up automatically at connection time. For example, batch jobs using SQL*Plus could type `sqlplus /` and authenticate using the credentials in the wallet.

Administrators and application developers now have a secure alternative to hardcoding usernames and passwords into scripts and other programs.

See also:

Oracle Database Security Guide for details

1.7.4.2 Enhanced Centralized Password Policies for the Database

This feature enforces password age, disabled accounts, and account lockouts for directory-managed users who access Oracle databases.

This feature allows security administrators to enforce centralized password policies for enterprise users when accessing one or more Oracle databases. It provides for reduced administrative costs as there is a single point of administration. This enhancement allows organizations to better manage security across the enterprise.

See also:

Oracle Database Enterprise User Administrator's Guide for details

1.7.4.3 Oracle Label Security Support for Real Application Clusters

Oracle Label Security support for RAC means that security policies will be synchronized between nodes. Oracle Label Security session security settings can fail over to other nodes automatically.

Oracle customers using RAC can now leverage Oracle Label Security to meet organizational security requirements for protection of sensitive data.

See also:

Oracle Label Security Administrator's Guide for details

1.8 Grid Computing

The following sections describe new features for Oracle Database 10*g* Release 2 in the Grid Computing area.

1.8.1 Data Provisioning

This release provides features to simplify data provisioning within a Grid environment, making it easier and faster to archive, move, and copy large sets of database data. These new features enable performing these bulk provisioning operations with no effect on any production database.

1.8.1.1 File Groups Repository

This feature enables storage of point-in-time copies of file and tablespace data for reporting, auditing, and information sharing.

Groups of files (for example, tablespaces) can be copied to the repository, or from the repository to another database or file system, without effecting access to production data.

See also:

Oracle Streams Concepts and Administration for details

1.8.1.2 Transportable Tablespace from Backup

This feature provides the ability to create a transportable tablespace without making it read-only.

This enables a DBA to copy or move a tablespace of data using the transportable tablespaces feature without making the tablespace read-only in the source database.

See also:

Oracle Database Backup and Recovery Advanced User's Guide for details

1.8.2 Oracle Scheduler

The Oracle Scheduler has been enhanced to provide additional critical functionality to meet the scheduling needs of enterprise customers. In addition to functionality and usability improvements, jobs can now be chained and scheduled based on external events.

1.8.2.1 Event Scheduling

This feature addresses the need to perform an action based on the occurrence of an event. The Scheduler can perform actions based on external events or actions can be performed based on the Scheduler generated events.

The benefit of this feature is that you can schedule jobs based on events.

See also:

Oracle Database Administrator's Guide for details

1.8.2.2 Job Chaining

A job chain describes a set of jobs that are interrelated. Rules are used to define the dependencies between the jobs.

This feature enables you to define dependencies between jobs.

See also:

Oracle Database Administrator's Guide for details

1.8.2.3 Scheduler Enhancements

This feature enables you to define a custom calendar such as a business fiscal year. Jobs can be scheduled to run on the last work day of every fiscal quarter. This feature also enables the use of rules (for example, Memorial Day is the last Monday of the month of May) to identify dates and save them as a schedule. Existing schedules can be combined to create composite schedules. For example, a HOLIDAYS schedule can be excluded from a DAILY schedule to create a WORKDAYS schedule.

The Scheduler enhancements enable customers to define complex schedules such as the last working day of every fiscal quarter.

See also:

Oracle Database PL/SQL Packages and Types Reference for details

1.9 Grid Management

The following sections describe new features for Oracle Database 10*g* Release 2 in the Grid Management area.

1.9.1 Configuration Plug-ins and Assistants

This release includes enhancements to configuration assistants or plug-ins and post-installation configuration.

1.9.1.1 Enterprise Manager Configuration Plug-in

This feature helps configure Enterprise Manager Database Control with Oracle Database 10*g* Release 2 Grid Control. EMCA is an executable within the ORACLE_HOME but can also be called as a plug-in from other configuration assistants (for example, DBCA).

The Enterprise Manager Configuration plug-in (EMCA) helps configure Enterprise Manager Database Control under various conditions:

- Out-of-box configuration using a seed database
- Out-of-box configuration with a custom database
- Post-install configuration using DBCA
- Post-install configuration using custom scripts
- Migration to Oracle Database 10*g* Release 2 from earlier versions

1.9.1.2 Migrate Non-ASM Databases to ASM

Automatic Storage Management (ASM) provides simplified storage management within Oracle Database 10*g*. ASM is integrated with DBCA and Enterprise Manager. DBCA offers functionality to create an ASM instance and configure the disk groups for database creation. Enterprise Manager offers functionality to manage and monitor ASM and disk groups as well as automating the various stages involved in converting an existing non-ASM database onto ASM-based storage.

Enterprise Manager provides an interface to help customers migrate their existing databases onto ASM from non-ASM databases.

1.9.1.3 Standalone Automatic Storage Management (ASM) Configuration

This feature adds a high-level option in DBCA, called ASM Disk Group Management, to create an ASM instance as well as configure disk groups independent of the database configuration.

In Oracle Database 10*g* Release 1, an ASM configuration using DBCA was supported only as part of database creation. In previous releases, setting up ASM required cancelling out of DBCA after ASM setup was done. In Oracle Database 10*g* Release 2, ASM Disk Group Management has made creating an ASM instance and configuring disk groups easier.

1.9.1.4 Standalone Disk Group Management

This feature administers ASM disk groups as standalone entities, independent of the creation of a database.

In Oracle Database 10*g* Release 2 (10.2), you can create an Automatic Storage Management (ASM) instance using an option in the Database Configuration Assistant (DBCA). You can, therefore, prepare an ASM environment for an existing database using the DBCA without concurrently creating a new and unneeded database. However, you can continue to use the DBCA to create a new database and build an ASM instance, if one does not already exist, to manage the new database's storage requirements.

See also:

Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide for details

1.9.1.5 Upgrade 10.1 Database to 10.2

The Database Upgrade Assistant guides the DBA through the upgrade process from supported releases of Oracle Database 10*g* Release 1 to Oracle Database 10*g* Release 2.

This feature provides a wizard interface to upgrade databases to the new release.

See also:

Oracle Database Upgrade Guide for details

1.9.2 Database Management

This release provides new and improved features to Enterprise Manager that facilitate easier management of the database.

1.9.2.1 Ability to Discover and Update Host Cluster Targets in Enterprise Manager

This feature provides a way to add or configure cluster targets either through auto discovery or manual configuration.

Currently you cannot manually add cluster targets in Enterprise Manager. This feature provides a way to create and configure cluster targets in case cluster discovery fails during installation.

1.9.2.2 Access Advisor Enhancements in Enterprise Manager

Enhancements to Access Advisor include:

- Template support
- Quick tune (tune a single statement)
- Journaling
- Recommendation implementation status
- Action implementation status

The Access Advisor user interface in Enterprise Manager has been enhanced to improve the user experience and expose more powerful features of the Access Advisor.

1.9.2.3 Backup and Recovery Configuration Enhancements in Enterprise Manager

Enhancements to Backup include:

- RMAN script jobs
- Backup job notifications
- Recovery catalog configuration

The Backup screens in Enterprise Manager have been enhanced to improve the user experience and expose more powerful RMAN features.

1.9.2.4 Clone External Objects as Part of the Cloning Operation

External object cloning enhancement in the Enterprise Manager database cloning feature ensures that external objects, which are tables stored as external files, get

copied over to the destination systems during the clone operation. This ensures that no data is left behind.

The ability to clone external objects extends the database cloning to ensure that all data is cloned, not just objects residing within the database.

1.9.2.5 Database Administration Enhancements in Enterprise Manager

Enhancements to database administration include:

- Redo log file size advisory
- All possible actions for an object are provided on the View/Edit object page
- Search for PL/SQL objects based on object status
- Delete multiple objects simultaneously
- Compile multiple PL/SQL objects simultaneously
- Enable or disable multiple triggers simultaneously

The database administration area has been enhanced through the addition of some minor features that increase the DBA's productivity.

1.9.2.6 Database Auditing Administration in Enterprise Manager

Enterprise Manager provides the following administration capabilities for database auditing:

- Configuration of audit-related initialization parameters
- Administration of audited objects: statement auditing and schema object auditing

This feature makes it convenient for the DBA to manage database auditing in the Enterprise Manager Web console, where all other aspects of Oracle Database are managed.

1.9.2.7 Enable or Disable Undo Retention Guarantee

Oracle Database 10*g* lets you guarantee undo retention. In Enterprise Manager 10*g* Release 2, you can enable and disable undo retention through the Enterprise Manager user interface. When you enable this option, the database never overwrites unexpired undo data, that is, undo data whose age is less than the undo retention period. This option is disabled by default, which means that the database can overwrite the unexpired undo data to avoid failure of DML operations if there is not enough free space left in the undo tablespace.

By enabling the guarantee option, you instruct the database not to overwrite unexpired undo data even if it means risking failure of currently active DML operations. Therefore, use caution when enabling this feature.

A typical use of the guarantee option is when you want to ensure deterministic and predictable behavior of Flashback Query by guaranteeing the availability of the required undo data.

1.9.2.8 Enhanced Backup Management

Enterprise Manager has added new functionality to Backup Management that provides a single view of all backups run against the databases in a domain. The details of the backup are available by clicking on the specific backup for the database.

By providing a combined view of all database backups and their status, a DBA can quickly pinpoint when a backup job did not complete. A drill down to the log allows the DBA to fix the problem and restart the backup.

See also:

Oracle Database 2 Day DBA for details

1.9.2.9 Enhanced Diagnostics for Cache Coherency Statistics

This release provides better drill down and diagnostics around cache coherency statistics.

Enhanced diagnostics for cache coherency helps users do more diagnostics around the cache coherency statistics.

1.9.2.10 Integrate Segment Advisor with Segment Findings

This feature provides better integration between the Segment Advisor and segment findings which provides a more unified interface with consistent findings and recommendations.

This release supports the new server features provided by the 10.2 Segment Advisor. This feature helps integrate the Segment Advisor with the segment findings. These enhancements help support non-ASSM tablespaces and detection of chained rows in segments.

1.9.2.11 Lock and Unlock Statistics

In previous releases, lock statistics would lock the state of statistics based on object usage in the application. When statistics on a table were locked, all the statistics depending on the table, including table statistics, column statistics, histograms, and statistics on all the dependent indexes were locked.

For Oracle Database 10*g* Release 2, the database server provides the lock and unlock feature in the optimizer statistic area. Enterprise Manager provides support for these features with user-friendly wizards to help assist DBAs with the operations.

You can now lock the statistics to a state where you know the statistic is good and will not change. For example, you can collect statistics for a volatile table while it is fully populated and freeze the statistics as a set of representative statistics.

You can also lock the state of the statistic when there are no statistics. This is useful to keep the statistics on a table uncollected in order to use Dynamic Sampling.

When you no longer want the statistics to be locked, use the Unlock wizard to unlock the statistics.

1.9.2.12 Management of Directory Objects in the Enterprise Manager Web Console

This feature provides management of Directory Objects in Enterprise Manager including search, view details, and life-cycle operations such as create, delete, and update.

The DIRECTORY object facilitates administering access and usage of BFILE data types. It is convenient for the DBA to manage Directory Objects in the Enterprise Manager Web console where all other aspects of Oracle Database are managed.

1.9.2.13 Miscellaneous Undo Advisor User Interface Enhancements

The following enhancements have been made in this release:

- The required undo tablespace size is computed as the maximum of the auto tuned value of undo retention and the low threshold value.
- In the Undo Advisor page, the field name "New Undo Retention" has been changed to "New Low Threshold Undo Retention."
- "Undo Retention (minutes) No Recommendation" is displayed if auto tuning is disabled. It is not displayed if auto tuning is enabled (the default).
- In the Undo Management page under the System Activity and Tablespace Usage, the default heading has changed to "Last One Hour" from "Last Seven Days".

This feature provides some user interface-related enhancements to the Undo Advisor pages.

1.9.2.14 New Storage and Configuration Policies

For Oracle Database 10*g* Release 2, the database Storage and Configuration policies have been enhanced as follows:

- Storage Policy: Database default temporary tablespace not set
- Storage Policy: Database permanent tablespace not set
- Storage Policy: Tablespace not using Automatic Segment Space Management
- Config Policy: Not using Dynamic SGA and PGA
- Config Policy: Not using the latest optimizer version
- Config Policy: STATISTICS_LEVEL parameter not TYPICAL

This feature enhances the existing storage and configuration policies by adding new database storage and configuration policies and consolidating any overlap in the existing policies.

This feature also enhances some existing metrics. The storage policies help the end user ensure that their tablespaces follow best practice rules. For example, they help the user identify dictionary-managed tablespace or tablespace that contains mixed data segments. This helps the end user in correctly setting up tablespace storage, thus avoiding space and performance problems.

The configuration policies check basic database settings to ensure the end user is taking advantage of the latest database features. These policies also help ensure that the databases are configured in the manner recommended by Oracle.

By adding this policy support, Enterprise Manager becomes more competitive in the marketplace.

1.9.2.15 Oracle Spatial Administration in the Enterprise Manager Web Console

Management of Oracle Spatial in Enterprise Manager includes creation and management of spatial tables, metadata, and spatial indexes.

Oracle Spatial provides a SQL schema and functions that facilitate the storage, retrieval, update, and query of collections of spatial features in an Oracle Database. It is convenient for the DBA to manage Oracle Spatial in the Enterprise Manager Web console where all other aspects of Oracle Database are managed.

1.9.2.16 Performance Monitoring for Large Clusters

This feature incorporates additional cache coherency monitoring and improves scalability for large clusters. This feature also provides general improvements in RAC database performance monitoring.

The performance monitoring capabilities have been expanded for RAC environments. RAC DB Home pages scale for 64-128 nodes/instances. There is additional monitoring of cache coherency for RAC instances.

The presented performance data is more meaningful, performance problems are easier to detect, and helps identify the root cause of cluster database performance issues.

1.9.2.17 Restore Statistics From Historic Snapshots

The Restore Statistics wizard enables restoration of statistics of the databases or their related objects to a specified timestamp in the past within the retention period.

For Oracle Database 10*g* Release 2, the database server provides the restore statistics feature in the optimizer statistic area. Enterprise Manager provides support with a user-friendly wizard to help assist DBAs with the operations.

If the statistics collected in the management window have caused a sub-optimal plan to be generated, you can revert to past statistics using the Restore Statistics wizard.

1.9.2.18 Restore and Recovery Enhancements in Enterprise Manager

Enhancements for Restore and Recovery include:

- Recover in context of new Oracle home
- Enhanced Guided Recovery wizard
- Dry-run restore

Restore and Recovery operations have been further automated in Enterprise Manager.

1.9.2.19 Segment Advisor to Work on the Whole Database

This feature helps support a predefined Segment Advisor task for the entire database.

This feature can be scheduled in a maintenance window.

1.9.2.20 Show Segment Details of Tablespace

For a selected tablespace, this feature displays the number of segments, segment names, segment types, number of extents, total size of extents for each segment (in blocks), the maximum number of extents that can be allocated to each segment, and the percent of the tablespace used.

The ability to display more details on tablespace segments will ease administration.

1.9.2.21 Statistics Collection for Fixed and Dictionary Objects

This feature supports Oracle Database 10*g* new feature of statistics collection for fixed and dictionary objects.

For Oracle Database 10*g* Release 2, the database server provides the capability to gather statistics for fixed and dictionary objects in the optimizer statistic area. Enterprise Manager provides a user-friendly wizard to help assist DBAs with the operations through the Gather Statistics wizard.

If your database suffers from performance issues due to out-of-date statistics for fixed and dictionary objects, use the 10.2 Enterprise Manager Gather Statistics wizard to allow statistics collection on these objects.

1.9.2.22 Tabular Display of Tablespace Segments

This feature displays a table of segments with the ability to sort based on individual columns.

The benefit of this feature is an improved display of tablespace segments and the ability to sort into a tabular format.

1.9.2.23 Track Interconnect Information Statistics

Now, you can collect and display additional cluster target properties including interconnect information statistics.

This feature allows Enterprise Manager to monitor the interconnect on a RAC instance.

1.9.2.24 Use RMAN Compression During Cloning

RMAN has a compression capability built in, and Enterprise Manager has been enhanced to offer the use of this capability within the database cloning feature.

The ability to use RMAN compression during database clone operations reduces the network bandwidth requirements for database cloning operations. This benefits anyone who is cloning over low bandwidth network links, and saves disk space when the database clones are saved for later reuse.

1.9.2.25 Wizard-Based Loading of Flat Files in Enterprise Manager

The Enterprise Manager Load Data wizard helps you specify a data file and then steps through some simple screens where you can provide information about the structure of the file (both character delimited and fixed length files are supported). Enterprise Manager uses the supplied information to automatically generate a SQL*Loader control file, and subsequently creates a SQL*Loader job for loading the data file into the database.

SQL*Loader is a powerful and scalable tool for importing data from files into Oracle Database. However, the process of creating SQL*Loader control files is time-consuming and error-prone. The Enterprise Manager Load Data wizard automates this process, thereby making the data loading capabilities of SQL*Loader more accessible.

See also:

Oracle Database 2 Day DBA for details

1.9.3 Patching

These features include enhancements to the patching tool used to patch databases and application server software.

1.9.3.1 OPatch the Patching Tool

OPatch is a tool for patching Oracle Database that is available from Oracle9*i* Release 2 and later and Application Server (starting 9.0.4.1) ORACLE_HOMEs. It also serves as a backend of Enterprise Manager patching.

This feature helps you apply interim patches to your environments. OPatch serves as the backend for Enterprise Manager patching and can therefore achieve large scale patching in Grid environments.

1.9.4 Software Install and Cloning

This release includes enhancements to Oracle Universal Installer (OUI) framework and individual product installs that are used for deploying Oracle software.

1.9.4.1 Improve Automatic Storage Management (ASM) Installation Process

The improved install process with Database Configuration Assistant (DBCA) and Installer can now place ASM in its own ORACLE_HOME and then configure ASM disk groups. Upon completion, you can run additional Oracle Universal Installer (OUI) sessions to install and configure one or more Oracle databases in separate ORACLE_HOMEs (or database environments) that would then use ASM.

Using the new installation process, you can set up the recommended configuration when more than one database instance runs on a single server or node (rather than having the ASM instance and a database instance share a single ORACLE_HOME). The improved installation process makes it easier to manage multiple database environments on a single server.

See also:

Oracle Database Installation Guide for Microsoft Windows (32-Bit) for details

1.9.4.2 Real Application Clusters Cloning

This feature provides the ability to clone RAC nodes to add a node to the cluster or to build a new cluster.

The RAC deployment model supports simple cloning mechanisms for extending RAC horizontally. The ability to clone and extend RAC horizontally is the key element of the Grid message.

1.9.4.3 Real Application Clusters Install Framework

The RAC install framework includes framework application programming and command-line interfaces that are needed for Oracle Clusterware and RAC installations on a new cluster or onto a newly added node to an existing cluster.

This feature helps you to provision RAC software and add nodes to an existing RAC environment with minimal manual effort.

1.9.4.4 Silent Install for Real Application Clusters

In this release, Real Application Clusters supports silent installation and operations, and the add note feature.

Silent operations enable you to use non-interactive installs. The add node feature is critical for increasing and decreasing capacity on demand. In summary, RAC provides a simple, scalable, and stable deployment framework.

1.10 Information Integration

The following sections describe new features for Oracle Database 10*g* Release 2 in the Information Integration area.

1.10.1 Oracle Streams Performance and Usability Improvements

These features provide functionality and performance enhancements, as well as usability improvements, to make Oracle Streams the most powerful, highest performing, and easiest-to-use information sharing solution for Oracle data.

1.10.1.1 Streams Performance

The following Streams components have been improved:

- Capture processes
- Propagation
- Apply processes

This feature improves the performance of Streams.

See also:

Oracle Streams Concepts and Administration for details

1.10.1.2 Support for Hot Mining With Downstream Capture

Downstream capture, the ability to off load the capture process from the production source database, can now be configured for near real-time processing (hot mining) of the primary source redo logs at the alternate database.

A capture process can run on a database other than the source database. The redo log files from the source database are copied to the other database, called a downstream database, and the capture process captures changes in these redo log files at the downstream database. To improve the performance of certain configurations, downstream capture can now be configured for near real-time processing (hot-mining) of the primary source redo logs at the alternate database.

See also:

Oracle Streams Concepts and Administration for details

1.10.1.3 Buffered Messaging

This feature enables non-persistent message queuing for user applications. This will enable you to build high performance message queuing applications.

This feature provides a faster, less resource-intensive queuing implementation.

See also:

Oracle Streams Advanced Queuing User's Guide and Reference for details

1.10.1.4 Commit-Time Ordering of Messages

This feature enables you to use commit-time to control the order in which user-enqueued events in a queue are browsed or dequeued.

Commit-time queues are useful when your environment needs to support transactional dependency ordering during dequeue, or when you need to support consistent browsing of events in a queue.

See also:

Oracle Streams Concepts and Administration for details

1.10.1.5 Apply Handler Dependency Declaration

Oracle Streams can apply changes made at the source site in parallel at the destination site while ensuring that dependent changes are applied in the same order as they were done at the source site. Occasionally, there are complex dependencies among tables that cannot be specified through database objects (such as referential constraints).

This feature enhances Oracle Streams to allow for the specification of these complex dependencies and the enforcement of these dependencies when applying the changes in parallel.

See also:

Oracle Streams Replication Administrator's Guide for details

1.10.1.6 Simplified Large Object (LOB) Handling in Streams Apply Process

With this feature, DML Handlers can receive LOB logical change records (LCR) as a single column, rather than iterating through multiple LOB LCRs (LOB chunks) for the column.

This feature improves ease-of-use by simplifying access to LOBs within a customized Streams apply process.

See also:

Oracle Streams Replication Administrator's Guide for details

1.10.1.7 Declarative LCR Transformations

The declarative logical change record (LCR) transformation feature enables simple configuration for DML LCRs of the most common transformation scenarios.

The ability to rename columns of a table, tables, or the table owner (schemas), as well as adding or removing columns from the LCR, is enabled external to user-supplied PL/SQL functions.

See also:

Oracle Streams Concepts and Administration for details

1.10.1.8 Notification Enhancements

Oracle Database 10*g* Release 2 provides advanced notification features to support cache invalidation.

This feature extends the notification capabilities available with Oracle Streams Advanced Queuing.

See also:

Oracle Streams Advanced Queuing User's Guide and Reference for details

1.10.1.9 Oracle Java Message Service (OJMS) Nonpersistent Messaging

This feature utilizes new AQ extensions and exposes "true" nonpersistent messaging through OJMS APIs. Previously, nonpersisent messaging was implemented on persistent, AQ messaging.

The benefit is improved performance and functional correctness for OJMS nonpersistent messaging.

1.10.1.10 Simplified Configuration APIs

In this release, common tasks such as performing an application upgrade or database migration has been simplified.

This feature improves ease-of-use by providing simplified API's for common tasks such as performing an application upgrade or database migration. Also, it completes RAC transparency by providing queue-to-queue propagation.

See also:

Oracle Streams Replication Administrator's Guide for details

1.10.1.11 Oracle Streams Administration in the Enterprise Manager Web Console

Enterprise Manager now supports the configuration, administration, and monitoring of Streams in Oracle Database.

This feature makes it convenient for the DBA to manage Oracle Streams in the Enterprise Manager Web console, where all other aspects of Oracle Database are managed.

1.10.2 Improved Client/Server and Server/Server Communications

Client/server and server/server communications have been improved in this release. Notification of failures will speed application failover providing better availability. Other improvements provide streamlined and more robust communications.

1.10.2.1 Local Naming Support for JavaNet

This release provides local naming capability, using tnsnames.ora, for JavaNet. With this feature, you can use either tnsnames.ora or LDAP naming adapters.

The benefit of this feature is that JavaNet performs name lookups the same as the C Net stack. This makes the system administrator's work easier and more consistent when configuring both JavaNet and C Net stack.

Support will continue for the current LDAP resolution that requires you to specify the LDAP server URL and DN in the CONNECT string. However, now you can use an alias for LDAP support which is the same as for the C Net stack.

1.10.2.2 Oracle Streams AQ Administration in the Enterprise Manager Web Console

Enterprise Manager provides the following administration features for Oracle Streams Advanced Queuing (AQ):

- Manage queue tables
- Manage queues
- Manage transformations

Oracle Streams AQ provides database-integrated message queuing functionality. It is now convenient for the DBA to manage Oracle Streams AQ in the Enterprise Manager Web console where he manages all other aspects of Oracle Database.

1.10.2.3 Transparent Application Failover (TAF) Notification

Failover will now be instantaneous, regardless of the state of the client when the failure occurs.

In the event of a failure of a database instance, the server will send a notification to clients to initiate failover.

See also:

Oracle Call Interface Programmer's Guide for details

1.11 Location Services

The following sections describe new features for Oracle Database 10*g* Release 2 in the Location Services area.

1.11.1 Extend Business Geographics Features to Support Europe

This release extends the Spatial routing engine and geocoder to work with European countries and improve the quality and robustness of these features (including Spatial Analysis and Mining).

1.11.1.1 Oracle Spatial Routing Engine Support for Western Europe

The Spatial routing engine introduced in Oracle Database 10*g* Release 1 has now been extended to provide driving distances, times, and directions between addresses for over a dozen Western European countries, including (but not limited to) Germany, United Kingdom, and France.

The Spatial routing engine enhancements enable logistics, transportation, and location-based services applications to build driving direction services for over a dozen Western European countries.

See also:

Oracle Spatial User's Guide and Reference for details

1.11.2 Extend Coordinate System in Core Spatial

This release supports the EPSG-based coordinate systems model for the Oil & Gas Industry.

1.11.2.1 Oracle Spatial Support for EPSG Coordinate Systems

Oracle Spatial now supports the European Petroleum Survey Group (EPSG) coordinate systems definitions.

Oracle Spatial coordinate systems support the European Petroleum Survey Group (EPSG) data model and data set that is essential for oil and gas exploration and production systems. EPSG support provides benefits of standardization, expanded support, and flexibility for oil and gas companies, georaster data vendors, and GIS users in general.

See also:

Oracle Spatial User's Guide and Reference for details

1.11.3 GeoRaster Compression Standards

This release supports industry standard compression techniques for Spatial GeoRaster.

1.11.3.1 Oracle Spatial GeoRaster Compression

Spatial GeoRaster supports industry standard compression techniques for raster image data, including the JPEG baseline (lossy) and DEFLATE (lossless) standards. All GeoRaster functions now work on both compressed and uncompressed images.

Remote sensing imagery results in very large data sets, growing at the rate of a terabyte or more for each day. The capability to store and manage these images in compressed form is key for users and DBAs alike. You save money on storage costs when image sizes are reduced by up to 80 percent. This is essential for defense/security, agricultural, and environmental monitoring applications.

See also:

Oracle Spatial GeoRaster for details

1.11.4 Improved Performance of Network Data Model

This release provides performance and scalability improvements to support analysis on very large networks.

1.11.4.1 Graph Partitioning for Very Large Networks in Oracle Spatial

The size of a network analyzed using Oracle Spatial network data model is no longer limited by the amount of physical memory available. This release includes graph partitioning, which divides a large network into logical partitions based on partition ID. Logical partitions can be incrementally loaded into memory for analysis.

Oracle Spatial network data model graph partitioning creates scalability that is independent of physical memory. Applications with large networks such as utility networks, street networks, and biological pathway data in life sciences will benefit from this important enhancement.

1.11.5 Improved Topology Operations

This release improves the performance, quality, and robustness of Spatial Topology according to the requirements of US CENSUS.

1.11.5.1 Whole Feature Inserts and Updates for the Oracle Spatial Topology Data Model

The Spatial topology data model has been extended to support feature-level spatial transactions against persistent topology in the database. In the previous release, multiple operations were required to insert or update all of the node, edge, and face elements of a feature. Now, a feature insert or update occurs as a single operation.

Whole feature-level inserts and updates dramatically reduce the code required for updating and maintaining topology data sets. This meets the need of major mapping agencies and land management applications.

See also:

Oracle Spatial Topology and Network Data Models for details

1.12 Performance and Scalability

The following sections describe new features for Oracle Database 10*g* Release 2 in the Performance and Scalability areas.

1.12.1 Improved Large Object (LOB) Performance

Oracle Database 10*g* Release 2 includes performance optimizations for commonly used features such as LOBs, PL/SQL, and Index Organized Tables. Performance gains are

obtained by redesign and code streamlining and do not require changes in customer applications using these features, unless explicitly stated otherwise.

1.12.1.1 LOB Access Statistics in Dynamic Performance Views

Dynamic performance views are essential tools for database administrators to monitor and tune database performance. With the additional LOB access statistics in dynamic performance views, database administrators are offered a clear picture of LOB access activities.

This feature provides LOB access statistics in dynamic performance views.

See also:

Oracle Database Application Developer's Guide - Large Objects for details

1.12.1.2 Online Redefinition Supports Parallel Execution for LONG-to-LOB Migration

Online Redefinition now performs parallel execution when converting from the LONG data type to the LOB data type on non-partitioned tables.

The benefit is reduced maintenance time required to convert to the LOB data type. This is especially beneficial to users of SAP and Oracle Applications.

See also:

Oracle Database Application Developer's Guide - Large Objects for details

1.12.1.3 Performance Improvements for Loading LOBs

Direct path load in SQL*Loader for out-of-line LOBs has been improved significantly.

Depending on the sizes of out-of-line LOBs, LOB loading performance has improved up to five times faster than in previous releases.

1.12.2 Other Performance Improvements

Oracle Database 10*g* Release 2 includes a number of additional performance optimizations for long running queries, background activities, and other commonly used features.

1.12.2.1 Asynchronous Commit

Asynchronous commit allows database clients to do more work while the database writes the redo log to disk.

This feature improves transaction throughput by reducing database commit time.

See also:

Oracle Database Application Developer's Guide - Fundamentals for details

1.12.2.2 Checkpointing Improvements

In Oracle Database 10*g* Release 2, improvements have been made to checkpointing.

This feature offers better performance and easier configuration of checkpointing.

1.12.2.3 Database Change Notification

You can now receive notification whenever a change occurs in the database on the result set of registered queries.

This feature enables any cache or object holding query results to ensure that it contains the very latest data.

See also:

Oracle Database Application Developer's Guide - Fundamentals for details

1.12.2.4 Fast Partition Split for Partitioned Index Organized Tables

The SPLIT command is used to split a single partition of an index organized table into two partitions. In earlier releases, the SPLIT command moved each of the rows from the existing partition into one of the two new partitions.

Index organized table performance is improved when a partition of an index organized table is split into two partitions. Rows are not moved if all of the rows in the existing partition fall into one of the new partitions.

See also:

Oracle Database Administrator's Guide for details

1.12.2.5 Improved Aggregation Performance

This feature provides enhancements to aggregation performance.

The benefit is faster execution of SQL operations containing aggregations.

See also:

Oracle Database Performance Tuning Guide for details

1.12.2.6 Parallel Join Bitmap Filtering

This feature reduces amount of data sent by the right side of a join based on the bitmap created by left side of a join.

This improves parallel join performance, especially on RAC, by reducing the amount of data traffic.

1.13 Server Manageability

The following sections describe new features for Oracle Database 10*g* Release 2 in the Server Manageability area.

1.13.1 Simplified Configuration

A number of new features in database and Enterprise Manager make initial database creation and setup even easier. Building on the work done in Oracle Database 10*g* Release 1, Oracle Database 10*g* Release 2 automates the configuration of more initialization parameters, such as multiblock read count. In addition, Enterprise Manager can now easily configure and manage a wide variety of advanced net services components, such as advanced security, network data encryption, and Transparent Application Failover (TAF).

1.13.1.1 Automatic Undo Retention Enhancement

This feature tunes undo retention to provide maximum retention for a fixed size undo tablespace.

This feature tunes the system to provide maximum undo retention for the fixed size undo tablespace. This further simplifies Automatic Undo Management, as you no longer need to configure the UNDO_RETENTION parameter even for flashback purposes.

See also:

Oracle Database Administrator's Guide for details

1.13.1.2 Net Services Administration in the Enterprise Manager Web Console

Enterprise Manager can now be used to manage the following Net Services components and features:

- Configuration of Oracle advanced security strong authentication.
- Configuration of network data encryption and integrity for Oracle servers and clients.
- Configuration of secure socket layer authentication.
- Transparent Application Failover (TAF) configuration.
- Search functionality for Listeners and Net Service names.
- Editable location of tnsnames.ora for import to Oracle Internet Directory.
- Listener target support for IMAP and POP presentation.
- Rename Listeners.
- Rename Net Services names.

Oracle Net Services provides enterprise-wide connectivity solutions in distributed, heterogeneous computing environments. It is convenient for the DBA to manage the various components of Oracle Net Services in the Enterprise Manager Web console, where all other aspects of Oracle Database are managed.

1.13.1.3 Self-Tuning Multiblock Read Count

The DB_FILE_MULTIBLOCK_READ_COUNT parameter controls the amount of block prefetching done in the buffer cache during scan operations, such as full table scan and index fast full scan. The value of this parameter can have a significant impact on the overall database performance. This feature enables Oracle Database to automatically select the appropriate value for this parameter depending on the operating system optimal I/O size and the size of the buffer cache.

This feature simplifies manageability by automating the tuning of DB_FILE_ MULTIBLOCK_READ_COUNT initialization parameter.

See also:

Oracle Database Performance Tuning Guide for details

1.13.2 Improved Intelligent Infrastructure

Enhancements to Intelligent Self-Management Infrastructure enable more accurate collection of workload and performance statistics, allow quick detection of causes behind the change in system behavior across two comparable time intervals, and make

it possible to easily diagnose transient performance problems. Additionally, Enterprise Manager now allows monitoring and troubleshooting of slow or "hung" systems by retrieving necessary information directly from SGA whenever a normal connection is not possible. The alert thresholds can now be configured to automatically adapt to changes in system behavior thanks to Enterprise Manager's new dynamic baselining capability.

1.13.2.1 Adaptive Alert Thresholds in Enterprise Manager

Enterprise Manager now computes statistical alert thresholds using either static (user-defined) or dynamic (self-adjusting) baselines.

Dynamic statistical baselines can significantly improve the accuracy of performance alerting while also reducing exposure to the false positives commonly incurred under fixed threshold schemes.

1.13.2.2 Database Performance Monitoring Enhancements in Enterprise Manager

The enhancements for Performance Monitoring include:

- System-level activity for any duration in real-time, recent real-time, and historical time
- CPU-level and Wait Class-level activity for any duration in real-time, recent real-time and historical time
- Ability to run Automatic Database Diagnostic Monitor (ADDM) on demand
- Display of Alert thresholds in the Performance pages

These enhancements greatly increase the ability of a DBA to quickly detect performance problems both proactively (using alerts, real-time data, and ADDM) and reactively (using recent real-time, historical data, and ADDM).

1.13.2.3 Diagnosis of Hung or Extremely Slow Databases Through Enterprise Manager

Enterprise Manager provides screens for monitoring slow systems. These screens are populated using the Direct SGA Attach mechanism. An intelligent processing layer transforms the collected raw statistics into useful information that can be used by DBAs to efficiently perform root-cause analysis and take appropriate action.

This feature makes it easier for DBAs to manage databases. In particular, it facilitates analysis of slow systems, and detection and analysis of hung systems. In both cases, the feature provides access to diagnostic information that would otherwise be impossible or extremely difficult to retrieve.

1.13.2.4 Display Session Activity History in Enterprise Manager

Enterprise Manager now provides the ability to view the history of database session activity for both realtime (less than one hour) and historical time contexts.

This feature greatly increases the ability of a DBA to diagnose problems related to a particular database session, as it permits the DBA to go back in time and view the history of the session activity.

1.13.2.5 Performance Graph Normalization in Enterprise Manager

Enterprise Manager now provides the ability to normalize database performance graphs by metric baseline statistics.

When performance information is normalized by the baseline context, it makes it easier for the DBA to detect if something unusual is happening. The definition of "unusual" depends on using a baseline as the measure of "normal."

1.13.2.6 Read Performance Statistics Directly From the System Global Area (SGA)

Enterprise Manager can now collect key performance statistics of the database without making SQL queries. This involves attaching to the SGA directly and reading statistics from the shared memory. The data collected by these direct SGA reads is used to populate the Enterprise Manager database performance screens.

When a database is hung or slow, it is typically facing severe contention in resources. Using SQL to collect performance statistics at such a time is not acceptable.

1.13.2.7 Cleanup System Statistics

This feature rationalizes some of the system statistics especially those concerning I/Os.

The benefit of this feature is easier performance tuning.

See also:

Oracle Database Performance Tuning Guide for details

1.13.2.8 Enhanced Active Session History

This feature provides additional Active Session History (ASH) information to find blocker and waiter identities and their associated transaction IDs and SQL.

This feature simplifies manageability by providing a more accurate view of recent system activity. This information helps DBAs investigate transient problems due to waiters or blockers that are not detected during a periodic performance analysis.

See also:

Oracle Database Performance Tuning Guide for details

1.13.2.9 Enhanced Automatic Database Diagnostic Monitor (ADDM) Performance Analysis

The intelligent self-management infrastructure has been enhanced to make Automatic Database Diagnostic Monitor (ADDM) analysis more accurate in the areas related to CPU, paging, and integrated cache. In addition, the scope of ADDM has itself been broadened to include more server components such as Streams, AQs, RMAN, and RAC. These enhancements allow ADDM to detect performance bottlenecks more accurately and provide better recommendations to resolve them.

This feature enhances manageability by making the database resident, self-diagnostic engine more powerful and accurate.

1.13.2.10 Progress Monitoring for Advisors

In this release, the advisor framework has been enhanced to include APIs for progress monitoring of advisors.

You can now externally monitor the progress of advisor tasks execution by querying the V$ADVISOR_PROGRESS view.

See also:

Oracle Database Performance Tuning Guide for details

1.13.2.11 Enhanced End-to-End Performance Monitoring

This feature enhances the End-to-End Performance Monitoring functionality introduced in Oracle Database 10*g* Release 1. The usability of SQL tracing has been improved by allowing a complete programmatic access at the session level. In addition, the sessions with tracing enabled can now be identified from V$SESSION view and a new dynamic view DBA_ENABLED_TRACES has been introduced to view the database and instance level settings. Also, a new procedure has been added to DBMS_MONITOR package to provide a single switch for enabling and disabling SQL tracing for all sessions.

This feature enhances manageability by adding ease-of-use functionality to the End-to-End Performance Monitoring feature introduced in Oracle Database 10*g* Release 1.

See also:

Oracle Database Performance Tuning Guide for details

1.13.2.12 Active Session History (ASH) Report

This feature uses Active Session History (ASH) data to analyze the performance of the system based on activity in the last few minutes.

This feature enhances manageability by making it easier to diagnose transient performance problems. This feature is intended to help administrators diagnose transient problems that are not detected during a period performance analysis.

1.13.2.13 Automatic Workload Repository (AWR) Compare Periods Report

This feature allows users to generate Automatic Workload Repository (AWR) reports identifying differences between two pairs of snapshots. This report is helpful to users in better understanding changes in workload and correlating it to performance problems pointed out by ADDM or validating tuning actions as a result of implementing ADDM's recommendations.

This feature enhances manageability by enabling simpler and more accurate performance diagnosis.

1.13.2.14 Configurable AWR SQL Collection

This feature allows users to customize how the top resource consuming SQL (TOP N SQL) statements are collected in the Automatic Workload Repository (AWR). By default, Oracle uses a set of predefined values to determine how many SQL statements should be collected for different measurement criteria. You can now change these values or direct AWR to capture all the statements in the cursor cache.

This feature enhances manageability by providing a more flexible mechanism to track the top resource-consuming SQL statements. Using this feature, administrators can capture a large number of SQL statements for historical performance trending and analysis purposes. This feature is also useful for SQL Workload performance comparison across major change activities such as database or application upgrade.

See also:

Oracle Database Performance Tuning Guide for details

1.13.2.15 Transportable Automatic Workload Repository (AWR)

This feature provides functionality to:

- Extract AWR snapshot data specified by snapshot range into a flat file.

- Load the extracted flat file into a specified staging schema on the same release of the database as the one the from which data was extracted.
- Copy over into SYS schema, if necessary, from the specified schema. Once the data is loaded into SYS schema, support can run the usual AWR reports or perform ADDM analysis.

If there's a collision in data in terms of DB ID and snapshot IDs in the current SYS schema or staging schema, the data is ignored. Snapshot data that does not collide with existing data is loaded. This feature is supported on Oracle Database 10*g* Release 2 and higher. The source and destination databases should be on the major and minor same release, but patch set releases could differ.

The primary goal of the transportable AWR snapshot data feature in this release is to support advanced AWR/ADDM problem diagnosis in consultation with support.

1.13.3 Automatic Storage Management (ASM) Manageability

This release expands the manageability capabilities for Automatic Storage Management (ASM) to both the GUI and command-line interface to provide the database administrator with greater flexibility.

1.13.3.1 ASM Command-Line Interface

The ASM Command-Line Interface (ASMCMD) utility provides an easy alternative for accessing the files and directories within ASM diskgroups.

Adding an additional means by which to access and manage files within ASM makes it simpler to manage. The ASMCMD interface provides easy manipulation of files within Automatic Storage Management (ASM) disk groups.

See also:

Oracle Database Utilities for details

1.13.3.2 ASM Wait on Rebalance Process

ASM ALTER DISKGROUP command, which results in a rebalance, now has the option of specifying the option to WAIT. This enables multiple actions to be done at one time with the rebalance process not starting until a later time. The Power_Limit can also be set to a value of 0 which will result in a wait before rebalance is started.

This makes it easier to manage the disks and rebalance process in ASM. The WAIT keyword allows a script that adds or removes disks to wait for the disk group to be rebalanced before proceeding.

See also:

Oracle Database Administrator's Guide for details

1.13.3.3 Additional ASM Manageability Enhancements

Additional ASM enhancements provide information about the status and space usage of ASM to the DBA. Adding two new fixed V$ASM views provides more access for Enterprise Manager to also reflect status and usage of disk space.

These enhancements make it easier to monitor and manage ASM disks. Enhancements to existing views and Enterprise Manager make information about check disk group and free space more concise and easier to understand.

1.13.3.4 Batch Selection of Multiple Disks for Operations Like DELETE, RESIZE, and CHECK

This feature allows multiple disks to be selected for DELETE, RESIZE, and CHECK operations.

The batch DELETE operation prevents the unnecessary rebalancing that occurs when disks are deleted one by one. It also improves usability and enhances the existing support in Enterprise Manager for ASM.

1.13.4 Automatic Storage Management (ASM) Extensions

This release expands the capabilities of ASM to support mixed Oracle environments. For example, support for multiple Oracle Database versions and support for multiple database-to-ASM connection types.

1.13.4.1 ASM Support for Multiple Database Releases

When mixing software releases, Automatic Storage Management (ASM) functionality reverts to the functionality of the lowest release being used. For example, a 10.1.0.2 database talking to a 10.1.0.3 ASM instance will not be able to exploit any new features in 10.1.0.3 and behaves as if it were talking to a 10.1.0.2 ASM instance. Conversely, a 10.1.0.3 database talking to a 10.1.0.2 ASM instance behaves as if it were a 10.1.0.2 database.

Two new text columns, SOFTWARE_VERSION and COMPATIBLE_VERSION, have been added to V$ASM_CLIENT.

This new support provides more flexibility for ASM in a multirelease database environment. To ensure that ASM transparently supports both older and newer software releases of the database, both forward-compatibility and backward-compatibility are maintained between all Oracle Database 10*g* Release 1 and all Oracle Database 10*g* Release 2 releases, and for as long as possible with later releases.

See also:

Oracle Database Administrator's Guide for details

1.13.4.2 Enhanced Protocol Support for Database to ASM Connection

This release provides enhanced protocol for connections between the database and Automatic Storage Management (ASM) for greater connectivity and compatibility. This also allows one clusterwide ASM instance to manage storage for many databases using several types of connection mechanisms.

The benefit is greater flexibility for ASM to manage a clustered pool of storage. Allowing more compatibility for database to ASM connections enables easier consolidation and compatibility of shared storage with a cluster.

See also:

Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide for details

1.13.5 File Support in Automatic Storage Management (ASM)

This release provides a user interface that allows the upload and extraction of files into or out of an ASM managed storage pool.

1.13.5.1 ASM/XDB FTP Support

This feature adds new functionality for both XDB and ASM providing a user interface for getting files into and out of an ASM managed pool of storage.

This feature extends ASM to support environments that require movement of application data.

See also:

Oracle Database Administrator's Guide for details

1.13.6 Enhanced Automatic Database Diagnostic Monitor (ADDM)

The self-diagnostic engine of Oracle Database 10*g*, Automatic Database Diagnostic Monitor (ADDM), has been enhanced to make its analysis more accurate in the areas related to CPU, Paging, and Integrated Cache. In addition, the scope of ADDM has itself been broadened to include more server components such as Streams, AQ, RMAN, and RAC.

1.13.6.1 Automatic Shared Memory Management of Streams Pool

When Streams is used in a single database, memory is allocated from a pool in the System Global Area (SGA) called the Streams pool. The Streams pool contains buffered queues and is used for internal communications during parallel capture and apply. The automatic shared memory management feature is enabled by setting the SGA_TARGET parameter.

The Streams pool is now an automatically sized SGA component to improve manageability.

See also:

Oracle Streams Concepts and Administration for details

1.13.7 Enhanced Instance Tuning

Automatic Instance Tuning capabilities of Oracle Database 10*g* have been further strengthened by enhancements made to Automatic Shared Memory Management and Database Resource Manager features. The Automatic Shared Memory Management feature now automatically manages the Streams pools component of the System Global Area (SGA). Also, the Database Resource Manager is now enabled out-of-the-box and new enhancements have been made to simplify its monitoring and tuning.

1.13.7.1 Automatically Enabled Resource Manager

Starting with this release, the Database Resource Manager will be enabled by default.

This feature enhances manageability by automatically activating the Database Resource Manager.

See also:

Oracle Database Administrator's Guide for details

1.13.7.2 Enhanced Ability to Diagnose Memory Allocation

Oracle Database provides a memory manager to internal subsystems for their dynamic memory allocation needs. Virtually all allocation and deallocation requests include a comment which can be used to maintain statistics, and which Oracle Support Services

can use to match memory usage to source code. This feature improves the ability to diagnose memory leaks and out-of-memory errors by maintaining more information, making component usage visible to DBAs and developers, and allowing Oracle Support Services to see memory usage in a more understandable format.

This feature enhances the information needed and improves the ability to diagnose memory allocation by Oracle Support Services.

See also:

Oracle Database Performance Tuning Guide for details

1.13.7.3 Enhanced Automatic Shared Memory Management

The Automatic Shared Memory Management feature now automatically sizes the Streams pool based on workload demand. The internal algorithm of the feature has also been enhanced to make the transfer of memory across SGA components more effective.

Enhancements to the Automatic Shared Memory Management feature further strengthen the manageability enhancements introduced in Oracle Database 10*g* Release 1.

See also:

Oracle Streams Concepts and Administration for details

1.13.7.4 Enhanced Resource Manager Monitoring

A number of new enhancements are being introduced to help assess the effectiveness of Database Resource Manager and tune Resource Plans if required. The Database Resource Manager is now fully integrated with the database self-management infrastructure (for example, ADDM, Automatic Workload Repository and Active Session History, Server Generated Alerts) in order to benefit from the self-management capabilities of Oracle Database 10*g*. In addition, a number of new statistics have been introduced and it is now possible to monitor the functioning of Database Resource Manager at the session level.

These features enhance manageability by making it easier to monitor and tune the functioning of the Database Resource Manager.

See also:

Oracle Database Administrator's Guide for details

1.13.8 Enhanced SQL Manageability

Oracle Database 10*g* Release 2 further enhances the Automatic SQL Tuning capability introduced in Release 1. SQL Profiles can now be shared among SQL statements having different literal values and SQL Tuning Sets can now be transported across different databases to enable application performance comparison during upgrade testing. SQL Access Advisor has also been enhanced to make it interruptible and provide recommendations about function-based indexes.

1.13.8.1 Transportable SQL Tuning Sets

This feature provides the capability to import and export SQL Tuning Sets (STS) from one system to another.

This feature allows the transfer of SQL workloads from one database to another for remote performance diagnostics and tuning. When poorly performing SQL statements

are encountered on a production system, you may not want to give developers direct access to it for the purpose of investigation and tuning. This feature allows the DBA to transport the offending SQL statements to a test system where they can be analyzed and tuned.

See also:

Oracle Database Performance Tuning Guide for details

1.13.8.2 Display SQL History in Enterprise Manager

Enterprise Manager provides the ability to view the execution history of a SQL statement.

This feature greatly increases the ability of a DBA to diagnose problems related to a particular SQL statement, as it permits the DBA to go back in time and view the history of the SQL statement's execution plans.

1.13.8.3 Interruptible SQL Access Advisor

The SQL Access Advisor may be interrupted while generating recommendations allowing the results to be reviewed.

The benefit is that recommendations can be viewed immediately rather than having to wait for the process to complete, which could be long for a complex database and workload.

See also:

Oracle Database Data Warehousing Guide for details

1.13.8.4 SQL Access Advisor Recommends Function-Based Indexes

The SQL Access Advisor now recommends function-based indexes.

This enhancement provides more indexing types recommended by SQL Access Advisor.

See also:

Oracle Database Data Warehousing Guide for details

1.13.8.5 SQL Profile Enhancements

This feature allows SQL Profile to match SQL text after literal values have been normalized into bind variables.

SQL Profile has been enhanced to allow applications to use literal values rather than bind variables.

See also:

Oracle Database Performance Tuning Guide for details

1.13.8.6 SQL Statistics Enhancement

A new view, V$SQLSTATS, has been added in this release.

This view has been added for improved performance, monitoring, and diagnostics.

See also:

Oracle Database Performance Tuning Guide for details

1.13.8.7 Timeout for Flushing SQL Statistics

This feature updates SQL statistics every 5 seconds for a SQL statement that is still in progress.

This feature makes it easy to monitor long running SQL statements that are still in progress by keeping SQL statistics up-to-date.

1.13.9 Enhanced Space Management

Space management operations have been further simplified in Oracle Database 10*g* Release 2 as it can automatically identify tables and indexes that can be compressed to reclaim wasted space. It is now possible to drop empty datafiles and the Automatic Segment Space Management (ASSM) feature is now enabled by default.

1.13.9.1 Default Automatic Segment Space Management (ASSM)

In this release, Automatic Segment Space Management (ASSM) is enabled by default. The tablespace attribute, Segment Space Management, now has the default value of AUTO.

All new tablespaces created using default attributes benefit from the manageability and performance advantages of ASSM over Freelist-based space management.

See also:

Oracle Database Administrator's Guide for details

1.13.9.2 Automatic Segment Advisor

Automatic Segment Advisor proactively identifies segments on a nightly basis that have significant wasted space due to data fragmentation and therefore are good candidates for Online Segment Shrink.

The key benefit is that the database administrators are proactively made aware of segments that need reorganization. They do not need to run Segment Advisor to determine what segments have data fragmentation. The advice is already available on the Enterprise Manager home page for immediate implementation.

See also:

Oracle Database Administrator's Guide for details

1.13.9.3 Online Segment Shrink Enhancements

This feature extends Online Segment Shrink capability to large object (LOB) segments and index-organized table (IOT) overflow segments.

Online Segments Shrink capability is now extended to all types of segments. This will further reduce the need to perform other more advanced reorganization methods such as Online Redefinition.

See also:

Oracle Database Administrator's Guide for details

1.13.9.4 Size-Based Tablespace Space Usage Alerts

This feature provides the ability to specify space pressure alerts based on the number of free bytes in the tablespace. Oracle Database 10*g* Release 1 provided alerts based on the percentage of the tablespace that was full; with the default being 85% for warning and 97% for critical alerts. For large tablespaces, you may want to specify thresholds in

absolute values since a tablespace that is 97% full may still have substantial free space available which does not warrant an alert.

With this feature, DBAs can exactly specify how much free space must be available in the tablespace before they get an alert.

See also:

Oracle Database Administrator's Guide for details

1.13.9.5 Drop Empty Datafile

This feature provides the ability to drop an empty datafile. It drops the datafile from the database and deletes it from the file system.

If a DBA adds a datafile in the wrong location or to the wrong tablespace, this feature provides a simple way to correct the error by dropping the datafile.

See also:

Oracle Database Administrator's Guide for details

1.14 Windows

The following sections describe new features for Oracle Database 10*g* Release 2 in the Windows area.

1.14.1 Improved Availability and Scalability With Oracle Real Application Clusters

ODP.NET has been enhanced with greater ability to intelligently use RAC. When optimizing their application, administrators can allocate the percentage of database connections for each instance, as well as have Oracle automatically detect downed instances and clean up bad connections. Additionally, Oracle clients provide RAC instance affinity for distributed transactions to make managing these transactions easier, especially during failover.

1.14.1.1 ODP.NET: Fast Application Notification (FAN)

ODP.NET connection pool subscribes to notifications that indicate whether nodes and services are available or unavailable. Based on these notifications, the ODP.NET connection pool can free idle connections that were connected to bad nodes and create new ones to healthy nodes, if possible.

The fast connection failover feature is a proactive way of migrating connections from downed services or downed nodes to responsive ones, leading to better availability with a minimal amount of application management needed.

See also:

Oracle Data Provider for .NET Developer's Guide for details

1.14.1.2 ODP.NET: Integration with Real Application Clusters Load Balancing Advisory

The ODP.NET connection pool subscribes to notifications provided by the RAC Load Balancing Advisory. These events provide information on the current service level provided by each instance and a recommendation on how connections are to be directed in a RAC database for a service that is enabled on multiple instances in the cluster. Connections are dispensed or removed from the connection pool based on the current service level.

The load balancing feature is a proactive way of routing work to responsive nodes based on the performance data gathered by the database. This feature allows efficient utilization of database resources to ensure optimum performance by limiting bottlenecks.

See also:

Oracle Data Provider for .NET Developer's Guide for details

1.14.2 Ease of Development

Integration with Microsoft Development Tools and implementation of MS Data Access specifications make it very easy for Windows Developers to develop with Oracle Database.

Enhancements in this area include:

- Integration with Visual Studio
- Added .NET Framework 1.1 functionality
- Database change notification
- Support for derived parameters
- Enhanced connection pool administration
- Control of query cancellation
- Support for server-side ODP.NET

1.14.2.1 Common Language Runtime (CLR) Stored Procedure Visual Studio Integration

This release provides a wizard in Visual Studio that automatically handles the data type mapping and PL/SQL wrapper code generation. This feature provides a way to deploy the built assembly into the database.

Integration with Visual Studio makes it easy for developers to build and deploy CLR stored procedures into an Oracle Database.

See also:

Oracle Database Extensions for .NET Developer's Guide for details

1.14.2.2 ODP.NET: .NET Framework 1.1 Functionality Support

The HasRows property indicates whether a result set is empty unless the application requests a row. The EnlistDistributedTransction method can explicitly decide when and in which transaction it wishes to participate. In .NET Framework 1.0, connections were only allowed to enlist at connection time to the transaction context that was associated with the component and thread. In .NET Framework 1.1, the transaction context can be passed around components (through the ITransaction interface) and have connections enlist in them explicitly.

The HasRows property and EnlistDistributedTransction method that were added to the .NET Framework in its 1.1 release are now supported by ODP.NET, allowing developers to use this additional functionality in an easy-to-use manner.

See also:

Oracle Data Provider for .NET Developer's Guide for details

1.14.2.3 ODP.NET: Database Change Notification Service

Oracle Database Change Notification Service enables applications to receive notifications when there is a change in the query result set of a statement. Using the Database Change Notification Service, the application can specify a notification request for a statement execution, and create a notification registration using the notification framework. When there is a change in the query result set, the notification framework notifies the application about the change. Based on the information provided by the notification framework, the application can then respond to the notification accordingly. For example, the application might need to refresh its own copy of the data for the query result set that is stored locally in the application.

An application can utilize Oracle Database Change Notification Service to be alerted when a database change occurs that affects the results of a query. The changes are propagated to the client to ensure that an accurate, up-to-date data set is used by the user.

See also:

Oracle Data Provider for .NET Developer's Guide for details

1.14.2.4 ODP.NET: Derived Parameters Support

ODP.NET supports the DeriveParameters method on the OracleCommandBuilder which populates the parameter collection for a command, given the stored procedure or function name. This feature is intended to be used at design time, not runtime, as it incurs a database round trip to execute.

The DeriveParameters method automatically populates metadata information needed to bind parameters to a stored procedure or function, saving the user from binding this information explicitly.

See also:

Oracle Data Provider for .NET Developer's Guide for details

1.14.2.5 ODP.NET: Enhanced Connection Pool Administration

This feature provides a way for applications to clear connections in one pool or all pools in an application domain. It provides the facility to refresh the connections and clear the pool of any invalid connections. It also allows applications to destroy the pool all together.

This feature allows developers to explicitly create or destroy connection pools or reset the attributes of existing connection pools to provide greater control in managing ODP.NET connection resources.

See also:

Oracle Data Provider for .NET Developer's Guide for details

1.14.2.6 ODP.NET: Query Cancel/Timeout

This feature allows an ODP.NET command to be canceled if results are not returned within a specified time. Otherwise, a command can continue to block other database calls that need to be executed.

Developers can control how long to wait for a command to complete before canceling the call, providing greater control over an application's service level.

See also:

Oracle Data Provider for .NET Developer's Guide for details

1.14.2.7 Server-Side ODP.NET

The server-side ODP.NET is the same as the client-side ODP.NET with a few exceptions. These exceptions are defined in Oracle Data Provider for .NET Developer's Guide.

A .NET stored procedure or function residing in Oracle may need to access the database. The usual interface to Oracle from .NET is the Oracle Data Provider for .NET. But if the .NET procedure or function is running on the server, there are certain optimizations and limitations that apply. For example, the implicit database session can be used to improve performance. Other features, like connection pooling are not supported.

See also:

Oracle Data Provider for .NET Developer's Guide for details

1.14.3 Performance Improvements

ODP.NET includes a number of improvements that make queries and LOB data retrieval performance faster. Query metadata is now cached so that repeated queries can be executed faster because statements only need to be parsed once. When retrieving LOBs, Oracle has reduced the number of database round trips necessary to retrieve the data to the client.

1.14.3.1 ODP.NET: Improved LOB Retrieval Performance and Functionality

ODP.NET has been improved to make fewer round trips when retrieving LOB data:

- To ensure full LOB functionality when InitialLOBFetchSize is used.
- To remove the requirement of the primary key, ROWID, or unique columns to be present in the query to fetch the remaining data for the LOB column when InitialLOBFetchSize is used.
- To allow InitialLOBFetchSize to fetch greater than 32 KB of data in the first round-trip.

LOB data retrieval now requires fewer database round trips, which improves performance. Additionally, LOBs now have fewer limitations making them easier to use and tune for performance.

See also:

Oracle Data Provider for .NET Developer's Guide for details

1.14.3.2 ODP.NET: Statement Caching

This feature provides and manages a cache of statements for each session. In the server, it means that cursors are ready to be used without the need to parse the statement again. Statement caching can be used with connection pooling and with session pooling, and improves performance and scalability. It can also be used without session pooling.

When identical SQL statements are executed, this feature improves performance by preventing parsing of the SQL statement repeatedly. The SQL statement is parsed just once.

See also:

Oracle Data Provider for .NET Developer's Guide for details

1.14.4 Security

This feature improves and enhances the security of Oracle Database and other components on windows.

1.14.4.1 ODP.NET: Enhanced Virtual Private Database (VPD) Support

This feature allows applications to set the Client Identifier in the Application Context for every database session using ODP.NET.

ODP.NET developers can more easily use Oracle's VPD support by allowing the Client Identifier, which acts like an application user proxy, to be set at the connection string level.

See also:

Oracle Data Provider for .NET Developer's Guide for details

1.14.4.2 Secure Operating System Objects

This feature creates a secure Win32 session environment for the Oracle instance so the process and it's internal state are not accessible from an external environment including sessions with the same privileges.

This enhancement prevents older clients from connecting when they are running on the same machine as the database.

This feature improves security for Oracle on Windows.

2

Oracle Database 10*g* Documentation

This chapter lists the books available online that document Oracle Database 10*g*.

2.1 Documentation Titles and Online Documentation Structure

All books are available in electronic format. The following subcategories of documentation are available:

- General Database Documents
- Application Development
- Backup
- Data Mining
- Database Server
- Enterprise Manager
- Gateways
- Installation
- Java Development
- Networking
- OLAP
- Real Application Clusters
- Release Notes
- Text
- Ultra Search
- Windows
- Workflow

2.1.1 Oracle Database 10*g* Documentation

The following sections provide a brief overview of Oracle Database 10*g* Release 2 documentation.

2.1.1.1 General Database Documents

Table 2–1 provides information about the *Oracle Database Read Me* and *Oracle Database Licensing Information* for Oracle Database 10*g* Release 2.

Table 2–1 General Database Documents

Documentation	Part Number	Description
Oracle Database Readme	B14233-01	A description of important last minute features and changes not included in Oracle Database Documentation Library.
Oracle Database Licensing Information	B14199-01	A description of Oracle Database licensing.

2.1.1.2 Application Development Documentation

Table 2–2 provides a list of Application Development documentation included in Oracle Database 10*g* Release 2.

Table 2–2 Application Development

Documentation	Part Number	Description
*Oracle SQL*Module for Ada Programmer's Guide*	A58231-03	Provides a complete description of Module Language, an ANSI/ISO SQL standard for developing applications that access data stored in a relational database. Module Language uses parameterized procedures to encapsulate SQL statements. The procedures can then be called from an Ada application. This guide also describes how you can use SQL*Module to call PL/SQL procedures stored in Oracle Database. A number of complete examples using Module Language, Ada code, and stored database procedures are provided.
*Pro*COBOL Programmer's Guide*	A96109-03	Describes how to develop COBOL programs that use the SQL and PL/SQL database languages to access and manipulate Oracle data.
Oracle Globalization Development Kit Java API Reference	B14224-01	Describes APIs for the Oracle Globalization Development Kit.
Oracle Database Application Developer's Guide - Large Objects	B14249-01	Describes use of large object (LOB) datatypes BLOB, CLOB, NCLOB, and BFILE in application development. You can use these datatypes to store and manipulate unstructured and semi-structured data in binary or character format. This book also describes APIs for working with LOBs in supported programmatic environments. New APIs that allow you to use LOBs 8 terabytes in size or larger, depending on your configuration, are described.
Oracle Call Interface Programmer's Guide	B14250-01	Presents the Oracle Call Interface (OCI), an application programming interface (API) that enables applications written in C or C++ to interact with one or more Oracle databases.
Oracle Database Application Developer's Guide - Fundamentals	B14251-01	Introduces the features needed to develop applications for Oracle Database 10*g*, particularly triggers and stored procedures. This book also gives a high-level overview of topics covered in more detail in other books, such as PL/SQL, Java, and OCI.
Oracle XML Developer's Kit Programmer's Guide	B14252-01	Introduces you to the Oracle XML Developer's Kit (XDK) and how the various language components of the XDK can work together to generate and store XML data in a database or in a document outside the database. Examples and sample applications are introduced where possible.
Oracle Database Application Developer's Guide - Workspace Manager	B14253-01	Describes how to use Oracle Database Workspace Manager to work with long transactions. Workspace management refers to the ability of the database to hold different versions of the same record (that is, row) in one or more workspaces. Users of the database can then change these versions independently. This manual includes conceptual, usage, and reference information.

Table 2–2 (Cont.) Application Development

Documentation	Part Number	Description
Oracle Spatial GeoRaster	B14254-01	Describes GeoRaster, a feature of Oracle Spatial that lets you store, index, query, analyze, and deliver GeoRaster data, that is, raster image data and its associated Spatial vector geometry data, plus metadata.
Oracle Spatial User's Guide and Reference	B14255-01	Provides usage and reference information for indexing and storing spatial data and for developing spatial applications. Includes an appendix about Oracle Locator 10*g* and its relationship to Oracle Spatial.
Oracle Spatial Topology and Network Data Models	B14256-01	Provides usage and reference information about the topology data model and network data model capabilities of Oracle Spatial.
Oracle Database PL/SQL Packages and Types Reference	B14258-01	Describes the PL/SQL packages and defined types supplied with Oracle Database. Packages are listed alphabetically with syntax, procedures and functions, and parameters described for each package.
Oracle XML DB Developer's Guide	B14259-01	Describes the database native XML support for this release. Oracle XML DB features a hierarchical repository with advanced foldering, versioning, and security mechanisms. This manual describes methods of storing, generating, accessing, searching, validating, transforming, and indexing XML and other data, including how to use FTP or HTTP/WebDav to access XML in the database. This manual also covers using Oracle XML DB with Oracle Text, Oracle Streams Advanced Queuing, and Oracle Enterprise Manager.
Oracle Database Application Developer's Guide - Object-Relational Features	B14260-01	Describes user-defined object datatypes and how to use these datatypes to model complex real-world entities as objects in the database.
Oracle Database PL/SQL User's Guide and Reference	B14261-01	Presents PL/SQL, the Oracle procedural extension of SQL, an advanced fourth-generation programming language. Explains the concepts behind PL/SQL and illustrates every facet of the language.
Oracle Database Application Developer's Guide - Rules Manager and Expression Filter	B14288-01	Provides usage and reference information about how to manage, index, and evaluate conditional expressions in relational tables.
Oracle Database Data Cartridge Developer's Guide	B14289-01	Describes how to implement custom indexing and query optimization services and how to package and use these as a server extension called a data cartridge.
Oracle Database Data Cartridge Java API Reference	B14292-01	Describes the Java APIs for the Data Cartridge.
Oracle Database XML Java API Reference	B14293-01	Describes the Java APIs for XML.
Oracle C++ Call Interface Programmer's Guide	B14294-01	Presents the Oracle C++ Call Interface (OCCI), an application program interface (API) that enables applications written in C++ to interact with one or more Oracle databases. To extend the functionality of code generated by the Object Type Translator (OTT), programmers may want to add code in the OTT-generated file. The OTT can distinguish between OTT-generated code and user-generated code by looking for some predefined markers (tags). Support for these tags has been added in this release.
Oracle interMedia Reference	B14297-01	Describes how to enable Oracle Database 10*g* to store, manage, and retrieve images, audio, video, or other heterogeneous media data in an integrated fashion with other enterprise information.

Table 2–2 (Cont.) Application Development

Documentation	Part Number	Description
Oracle interMedia Java Classes for Servlets and JSP API Reference	B14300-01	Describes how to use servlets and JSPs to facilitate the retrieval and uploading of multimedia data from and to Oracle Database.
Oracle interMedia Java Classes API Reference	B14301-01	Describes how to enable Oracle Database to store, manage, and retrieve multimedia data in an integrated fashion with other enterprise information. Oracle *inter*Media extends Oracle Database reliability, availability, and data management to multimedia content in Internet, electronic commerce, and media-rich applications.
Oracle interMedia User's Guide	B14302-01	Describes how to develop and deploy multimedia content applications running on Oracle Application Server or Oracle Database that enables the database to store, manage, and retrieve image, audio, video, and heterogeneous media data using the relational and object interfaces.
Oracle HTML DB User's Guide	B14303-02	Oracle HTML DB is a browser-based development environment that enables users to build data-driven applications called flows. A flow consists of multiple HTML pages linked together with tabs, buttons, or hypertext links. Oracle HTML DB dynamically renders and processes these pages from data stored in tables. Oracle HTML DB consists of three core components: Data Workshop. Enables users to import data into and export data from a hosted database. SQL Workshop. Enables uses to view and manage database objects from a Web browser. Flow Builder. Provides users with an online development environment in which to build the pages that comprise of flow.
*Pro*FORTRAN Supplement to the Oracle Precompilers Guide*	B14352-01	Describes how to write FORTRAN programs that use SQL to access and manipulate Oracle data.
*Pro*PL/1 Supplement to the Oracle Precompilers Guide*	B14353-01	Describes how to write PL/1 programs that use SQL to access and manipulate Oracle data. An understanding of the material in the Programmer's Guide to the Oracle Precompilers is assumed.
Oracle Database Programmer's Guide to the Oracle Precompilers	B14354-01	This guide is a comprehensive user's guide and on-the-job reference to the Pro*COBOL and Pro*FORTRAN Precompilers. It shows you step-by-step how to develop applications that use the powerful database language SQL to access and manipulate Oracle data.
Oracle Spatial Java API Reference	B14373-01	Consists of interfaces and classes that support features available with Oracle Spatial, including GeoRaster and the topology and network data models, and with MapView, which is packaged with Oracle Application Server.
Oracle HTML DB 2 Day Developer	B14377-01	Contains tutorials with step-by-step instructions that explain how to create a variety of application components and complete applications using the Oracle HTML DB development environment.
*Pro*C/C++ Programmer's Guide*	B14407-01	Describes how to develop C++ programs that use the SQL and PL/SQL database languages to access and manipulate Oracle data.
Oracle Database XML C++ API Reference	B16206-01	Provides information about XML related C++ APIs. It primarily lists the syntax of functions, methods, and procedures associated with these API.
Oracle Database XML C API Reference	B16207-01	Provides information about XML related C APIs. It primarily lists the syntax of functions, methods, and procedures associated with these APIs.

2.1.1.3 Backup Documentation

Table 2–3 provides a list of Backup documentation included in Oracle Database 10*g* Release 2.

Table 2–3 Backup

Documentation	Part Number	Description
Oracle Database Backup and Recovery Advanced User's Guide	B14191-01	Provides in-depth information on the mechanics of backup and recovery, and a guide to performing complex and less frequently performed backup and recovery tasks, including user-managed backup and recovery and performance tuning of backup and recovery.
Oracle Database Backup and Recovery Basics	B14192-01	Provides introductory background information on backup and recovery and a guide on how to perform common backup and recovery tasks.
Oracle Database Backup and Recovery Quick Start Guide	B14193-01	This book is the starting point for all users of the backup and recovery documentation. It presents the basics of backup and recovery through the Recovery Manager, provides a guide to using the rest of the backup and recovery documentation, and includes a quick reference for frequently used Recovery Manager commands and views.
Oracle Database Backup and Recovery Reference	B14194-01	This book describes RMAN syntax and recovery catalog views.

2.1.1.4 Data Mining Documentation

Table 2–4 provides a list of Data Mining documentation included in Oracle Database 10*g* Release 2.

Table 2–4 Data Mining

Documentation	Part Number	Description
Oracle Data Mining Administrator's Guide	B14338-01	Describes how to install the Oracle Data Mining software and how to perform other administrative functions common to all Oracle Data Mining environments on both UNIX and Windows platforms.
Oracle Data Mining Concepts	B14339-01	Provides an overview of basic Oracle Data Mining concepts.
Oracle Data Mining Application Developer's Guide	B14340-01	Describes how to use the Oracle Data Mining PL/SQL and Java APIs and the SQL scoring functions to create data mining applications.
Oracle Data Mining Java API Reference	B14341-01	Describes the classes and methods in the Oracle Data Mining Java API, the JDM JSR-73-compliant API for data mining. The Java API supports a full range of data mining activities, including model building and scoring, data preparation, and import/export of models.

2.1.1.5 Database Server Documentation

Table 2–5 provides a list of Database Server documentation included in Oracle Database 10*g* Release 2.

Table 2–5 Database Server

Documentation	Part Number	Description
Oracle HTTP Server Administrator's Guide	B14190-01	Describes how to administer Oracle HTTP Server.
Oracle Database SQL Quick Reference	B14195-01	Contains a high-level description of the SQL used to manage information in Oracle Database.
Oracle Database 2 Day DBA	B14196-01	Provides a single source reference for administering Oracle Database in a small to midsize environment. If you are familiar with computers, but unfamiliar with administering Oracle Database, and your computer has four CPUs or fewer, and you do not expect to support more than 500 users, this book is appropriate for you.
Oracle Database Sample Schemas	B14198-01	Describes the sample schemas that are included in the seed database that ships with Oracle Database 10*g*. These schemas are used in examples in Oracle documentation and curriculum materials and in demos on Oracle Technology Network (OTN).
Oracle Database SQL Reference	B14200-01	Contains a complete description of the SQL used to manage information in Oracle Database. The first volume describes all of the building blocks of SQL statements. The remaining volumes describe the SQL statements alphabetically.
Oracle Database High Availability Overview	B14210-01	Describes Oracle's approach for a highly available database environment. Provides an overview of high availability and helps you to determine your high availability requirements. It also describes Oracle Database products and features that are designed to support high availability and describes the primary database architectures that can help your business achieve high availability.
Oracle Database Performance Tuning Guide	B14211-01	Explains in detail how to enhance performance by writing and tuning SQL properly, by using performance tools, and by optimizing instance performance.
Oracle Database New Features	B14214-01	Describes the new features of Oracle Database 10*g* and lists the documentation for this release.
Oracle Database Utilities	B14215-01	Describes how to use Oracle Database utilities to load data into a database, transfer data between databases, and maintain data. The utilities described include the new Data Pump export and import products, the original export and import products, SQL*Loader, external tables, the Metadata API, LogMiner, DBVERIFY, and DBNEWID.
Oracle Database Error Messages	B14219-01	Online only. Includes all database error messages.
Oracle Database Concepts	B14220-01	Describes how Oracle Database functions and lays a conceptual foundation for more detailed information contained in other Oracle Database books.
Oracle Database Data Warehousing Guide	B14223-01	Provides conceptual, reference, and implementation material for using Oracle Database 10*g* in data warehouses. It covers the full range of data warehousing activities, from physical database design to advanced calculation techniques.
Oracle Database Globalization Support Guide	B14225-01	Describes Oracle globalization support for the database. Explains how to set up a globalization support environment, choose and migrate a character set, customize locale data, do linguistic sorting, program in a global environment, and program with Unicode.

Table 2–5 (Cont.) Database Server

Documentation	Part Number	Description
Oracle Database Advanced Replication	B14226-01	Contains conceptual information about Oracle Advanced Replication. Also includes information about planning an Advanced Replication environment, an introduction to the Replication Management tool in Oracle Enterprise Manager, and information about troubleshooting Advanced Replication problems.
Oracle Database Advanced Replication Management API Reference	B14227-01	Contains reference information about and step-by-step instructions for using the replication management API, which is a set of PL/SQL packages for setting up and managing an Advanced Replication environment. Also includes reference information about data dictionary views that are important for Advanced Replication and security considerations for an Advanced Replication environment.
Oracle Streams Replication Administrator's Guide	B14228-01	Contains conceptual information about Oracle Streams replication and about configuring, administering, and monitoring an Oracle Streams replication environment. Also includes detailed examples of configuring Oracle Streams replication environments.
Oracle Streams Concepts and Administration	B14229-01	Contains conceptual information about Oracle Streams and information about configuring, administering, and monitoring an Oracle Streams environment. Also includes detailed examples of using Oracle Streams for various purposes.
Oracle Data Guard Broker	B14230-01	Describes the Oracle Data Guard broker, a management and monitoring interface that automates many of the tasks involved in configuring and monitoring standby databases in an Oracle Data Guard configuration. This guide provides comprehensive descriptions and examples for using both the command-line interface and the Oracle Data Guard Manager graphical user interface.
Oracle Database Administrator's Guide	B14231-01	Explains how to create and to manage Oracle Database. Includes information about distributed processing.
Oracle Database Heterogeneous Connectivity Administrator's Guide	B14232-01	Describes the Oracle approach to information integration in a heterogeneous environment. Describes Oracle Transparent Gateways and Generic Connectivity and is an administrator's guide for these Oracle products.
Oracle Database Reference	B14237-01	Provides reference information about database initialization parameters, static data dictionary views, dynamic performance views, database limits, and SQL scripts that are part of Oracle Database system.
Oracle Database Upgrade Guide	B14238-01	Explains the process of planning and executing database upgrades on Oracle Database server. In addition, this guide provides information about compatibility, about upgrading applications to the current release of Oracle, and about important changes in the current release, such as initialization parameter and data dictionary changes.
Oracle Data Guard Concepts and Administration	B14239-01	Provides a comprehensive overview of Oracle Data Guard concepts and describes how to configure and implement standby databases that can take over production operations if your production database becomes unusable. This guide includes several database scenarios such as creating, recovering, failing over, switching over, configuring, and backing up standby and primary databases.
Oracle Streams Advanced Queuing User's Guide and Reference	B14257-01	Describes features of application development and integration using the Oracle messaging system, Advanced Queuing (AQ). Includes the PL/SQL, C, Visual Basic, Java, and JMS interfaces to AQ and gateways to non-Oracle messaging systems.

Table 2–5 (Cont.) Database Server

Documentation	Part Number	Description
Oracle Streams Advanced Queuing Java API Reference	B14291-01	Describes the Java APIs for Streams Advanced Queuing.
Oracle HTTP Server mod_plsql User's Guide	B14337-01	Describes how to use `mod_plsql` with HTTP Server.
*SQL*Plus Quick Reference*	B14356-01	Shows iSQL*Plus buttons and icons, and iSQL*Plus and SQL*Plus command syntax.
*SQL*Plus User's Guide and Reference*	B14357-01	Introduces SQL*Plus and its uses, and provides a description of each SQL*Plus command.
Oracle Database Administrator's Reference for UNIX-Based Operating Systems	B15658-01	Explains how to create and to manage Oracle Database on Linux and Solaris operation systems.

2.1.1.6 Enterprise Manager Documentation

Table 2–6 provides a list of Enterprise Manager documentation included in Oracle Database 10*g* Release 2.

Table 2–6 Enterprise Manager

Documentation	Part Number	Description
Oracle Universal Installer and OPatch User's Guide	B16227-01	Describes how to use the Oracle Universal Installer to install Oracle and third-party software.

2.1.1.7 Gateways Documentation

Table 2–7 provides a list of Gateways documentation included in Oracle Database 10*g* Release 2.

Table 2–7 Gateways

Documentation	Part Number	Description
Oracle Transparent Gateway for Microsoft SQL Server Administrator's Guide for Microsoft Windows	B14270-01	Describes the Oracle Transparent Gateway for Microsoft SQL Server, which enables Oracle client applications to access Microsoft SQL Server data through Structured Query Language (SQL).
Oracle Transparent Gateway for Sybase Administrator's Guide for Microsoft Windows	B14271-01	Describes the Oracle Transparent Gateway for Sybase on Microsoft Windows, which enables Oracle client applications to access Sybase data through Structured Query Language (SQL).
Oracle Transparent Gateway for Teradata Administrator's Guide for Microsoft Windows	B14272-01	Describes the Oracle Transparent Gateway for Teradata on Microsoft Windows, which enables Oracle client applications to access Teradata data through Structured Query Language (SQL).
Oracle Transparent Gateway for Informix Administrator's Guide for hp-ux	B14274-01	Describes the Oracle Transparent Gateway for Informix on HP-UX, which enables Oracle client applications to access Informix data through Structured Query Language (SQL).
Oracle Transparent Gateway for Informix Administrator's Guide for Solaris Operating System (SPARC 64-Bit)	B14275-01	Describes the Oracle Transparent Gateway for Informix on Solaris, which enables Oracle client applications to access Informix data through Structured Query Language (SQL).

Table 2–7 (Cont.) Gateways

Documentation	Part Number	Description
Oracle Transparent Gateway for Teradata Administrator's Guide for hp-ux	B14276-01	Describes the Oracle Transparent Gateway for Teradata on HP-UX, which enables Oracle client applications to access Teradata data through Structured Query Language (SQL).
Oracle Transparent Gateway for Ingres II Administrator's Guide for hp-ux	B14277-01	Describes the Oracle Transparent Gateway for Ingres II on HP-UX, which enables Oracle client applications to access Ingress II data through Structured Query Language (SQL).
Oracle Transparent Gateway for Ingres II Administrator's Guide for Solaris Operating System (SPARC 64-Bit)	B14278-01	Describes the Oracle Transparent Gateway for Ingress II on Solaris, which enables Oracle client applications to access Ingress II data through Structured Query Language (SQL).
Oracle Transparent Gateway for Rdb Administrator's Guide for hp OpenVMS	B14279-01	Describes the Oracle Transparent Gateway for Rdb on HP OpenVMS, which enables Oracle client applications to access Rdb data through Structured Query Language (SQL).
Oracle Transparent Gateway for Teradata Administrator's Guide for Solaris Operating System (SPARC 64-Bit)	B14280-01	Describes the Oracle Transparent Gateway for Teradata on Solaris, which enables Oracle client applications to access Teradata data through Structured Query Language (SQL).
Oracle Transparent Gateway for RMS Administrator's Guide for hp OpenVMS	B14281-01	Describes the Oracle Transparent Gateway for RMS on HP OpenVMS, which enables Oracle client applications to access RMS data through Structured Query Language (SQL).
Oracle Transparent Gateway for Sybase Administrator's Guide for hp-ux	B14283-01	Describes the Oracle Transparent Gateway for Sybase on HP-UX, which enables Oracle client applications to access Sybase data through Structured Query Language (SQL).
Oracle Transparent Gateway for Sybase Administrator's Guide for IBM AIX	B14284-01	Describes the Oracle Transparent Gateway for Sybase on IBM AIX, which enables Oracle client applications to access Sybase data through Structured Query Language (SQL).
Oracle Transparent Gateway for Sybase Administrator's Guide for hp Tru64 UNIX	B14285-01	Describes the Oracle Transparent Gateway for Sybase on HP Tru64 UNIX, which enables Oracle client applications to access Sybase data through Structured Query Language (SQL).
Oracle Transparent Gateway for Sybase Administrator's Guide for Solaris Operating System (SPARC 64-Bit)	B14286-01	Describes the Oracle Transparent Gateway for Sybase on Solaris, which enables Oracle client applications to access Sybase data through Structured Query Language (SQL).
Oracle Procedural Gateway for APPC Installation and Configuration Guide for UNIX	B16209-01	Describes installing, configuring, and administering the Oracle Procedural Gateway for APPC on the UNIX platform.
Oracle Procedural Gateway for APPC User's Guide for UNIX	B16210-01	Describes the architecture, uses, and features of the Oracle Procedural Gateway for APPC on the UNIX platform.
Oracle Procedural Gateway for APPC Installation and Configuration Guide for Microsoft Windows	B16211-01	Provides Oracle applications with seamless access to IBM mainframe data and services through Remote Procedure Call (RPC) processing.
Oracle Procedural Gateway for APPC User's Guide for Microsoft Windows	B16212-01	Provides Oracle applications seamless access to virtually any APPC-enabled system, including IBM mainframe data and services through Remote Procedure Call (RPC) processing. The document is intended for Windows users.
Oracle Procedural Gateway for APPC Messages Guide for UNIX and Microsoft Windows	B16213-01	Contains messages for all Oracle Procedural Gateway for APPC products.

Table 2–7 (Cont.) Gateways

Documentation	Part Number	Description
Oracle Procedural Gateway for WebSphere MQ Installation and User's Guide for UNIX	B16215-01	Provides additional information about installing, configuring, and administering the Oracle Procedural Gateway for WebSphere MQ. It is also for developers writing applications that access message queuing systems, particularly those developers who need to access queues owned by both WebSphere MQ and other non-Oracle message queuing systems as well as queues owned by Oracle Advanced Queuing (AQ).
Oracle Procedural Gateway for WebSphere MQ Installation and User's Guide for Microsoft Windows	B16216-01	Describes installing, configuring, and administering the Oracle Procedural Gateway for WebSphere MQ on the Windows platform.
Oracle Transparent Gateway for DRDA Installation and User's Guide for UNIX	B16217-01	Describes installing, configuring, and administering the Oracle Transparent Gateway for DRDA on the UNIX platform.
Oracle Transparent Gateway for DRDA Installation and User's Guide for Microsoft Windows	B16218-01	Provides users with transparent access to DRDA databases as if they were Oracle databases.
Oracle Transparent Gateway for DB2 Installation and User's Guide for IBM z/OS (OS/390)	B16220-01	Describes installing, configuring, and administering the Oracle Transparent Gateway for DB2 on the IBM z/OS platform.
Oracle Transparent Gateway for iWay Installation and User's Guide for IBM z/OS (OS/390)	B16221-01	Describes installing, configuring, and administering the Oracle Transparent Gateway for iWay on the IBM z/OS platform.
Oracle Transparent Gateway for DB2/400 Installation and User's Guide for IBM iSeries OS/400	B16222-01	Describes installing, configuring, and administering the Oracle Transparent Gateway for DB2/400 on the IBM iSeries OS/400 platform.
Oracle Access Manager for AS/400 Installation and User's Guide for IBM iSeries OS/400	B16223-01	Describes installing, configuring, and administering the Oracle Access Manager for AS/400 on the IBM iSeries OS/400 platform.
Oracle Procedural Gateway Visual Workbench for WebSphere MQ Installation and User's Guide for Microsoft Windows (32-Bit)	B19082-01	Describes installing, configuring, and administering the Oracle Procedural Gateway for Visual Workbench for WebSphere MQ on the Windows platform.
Oracle Transparent Gateway Installation Guide for Microsoft Windows (32-Bit)	B19083-01	Describes installing and configuring the Oracle Transparent Gateway for the Windows platform.
Oracle Transparent Gateway Installation Guide for Solaris Operating System, hp-ux, IBM AIX, and hp Tru64 UNIX	B19084-01	Describes installing and configuring the Oracle Transparent Gateway for the Solaris, hp-ux, IBM AIX, and hp Tru64 UNIX platforms.
Oracle Transparent Gateway Installation Guide for hp OpenVMS	B19085-01	Describes installing and configuring the Oracle Transparent Gateway for the hp OpenVMS platform.

2.1.1.8 Installation Documentation

Table 2–8 provides a list of Installation documentation included in Oracle Database 10*g* Release 2.

Table 2–8 Installation

Documentation	Part Number	Description
Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide for AIX Based Systems	B14201-01	Explains Oracle Real Application Clusters (RAC) setup and configuration procedures for AIX as well as how to use Oracle tools for software installation and database creation. This RAC install book is new for Oracle Database10*g*.
Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide for hp-ux	B14202-01	Explains Oracle Real Application Clusters (RAC) setup and configuration procedures for hp HP-UX PA-RISC (64-Bit) as well as how to use Oracle tools for software installation and database creation. This RAC install book is new for Oracle Database10*g*.
Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide for Linux	B14203-01	Explains Oracle Real Application Clusters (RAC) setup and configuration procedures for Linux as well as how to use Oracle tools for software installation and database creation. This RAC install book is new for Oracle Database10*g*.
Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide for Apple Mac OS X	B14204-01	Explains Oracle Real Application Clusters (RAC) setup and configuration procedures for Apple Mac OS X as well as how to use Oracle tools for software installation and database creation. This RAC install book is new for Oracle Database10*g*.
Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide for Solaris Operating System	B14205-01	Explains Oracle Real Application Clusters (RAC) setup and configuration procedures for Solaris Operating System (SPARC 64-Bit) as well as how to use Oracle tools for software installation and database creation. This RAC install book is new for Oracle Database10*g*.
Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide for hp Tru64 UNIX	B14206-01	Explains Oracle Real Application Clusters (RAC) setup and configuration procedures for HP Tru64 as well as how to use Oracle tools for software installation and database creation. This RAC install book is new for Oracle Database10*g*.
Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide for Microsoft Windows	B14207-01	Explains Oracle Real Application Clusters (RAC) setup and configuration procedures for Microsoft Windows Platforms as well as how to use Oracle tools for software installation and database creation. This RAC install book is new for Oracle Database10*g*.
Oracle Database Client Installation Guide for Microsoft Windows (32-Bit)	B14312-01	Describes how to install Windows 32-bit clients.
Oracle Database Client Installation Guide for Microsoft Windows (64-Bit) on Intel Itanium	B14313-01	Describes how to install Windows clients on the Intel Itanium platform.
Oracle Database Client Quick Installation Guide for Microsoft Windows (32-Bit)	B14314-01	Describes how to quickly install Oracle Database Client on the Windows 32-bit systems.
Oracle Database Client Quick Installation Guide for Microsoft Windows (64-Bit) on Intel Itanium	B14315-01	Describes how to quickly install Oracle Database Client for Windows on Intel Itanium systems.

Table 2–8 (Cont.) Installation

Documentation	Part Number	Description
Oracle Database Installation Guide for Microsoft Windows (32-Bit)	B14316-01	Provides platform-specific preinstallation, installation, postinstallation, and component information for your single instance Oracle Database 10*g* database and client software on the Windows 32-bit platform.
Oracle Database Installation Guide for Microsoft Windows (64-Bit) on Intel Itanium	B14317-01	Provides platform-specific preinstallation, installation, postinstallation, and component information for your single instance Oracle Database 10*g* database and client software for Windows on the Intel Itanium platform.
Oracle Database Quick Installation Guide for Microsoft Windows (32-Bit)	B14318-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products for Windows 32-bit systems.
Oracle Database Quick Installation Guide for Microsoft Windows (64-Bit) on Intel Itanium	B14319-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products for Windows on Intel Itanium systems.
Oracle Database Companion CD Quick Installation Guide for Microsoft Windows (32-Bit)	B14322-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products on Windows 32-bit systems.
Oracle Database Companion CD Quick Installation Guide for Microsoft Windows (64-Bit) on Intel Itanium	B14323-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products for Windows on Intel Itanium systems.
Oracle Database Companion CD Installation Guide for Microsoft Windows (32-Bit)	B14324-01	Explains how to install and configure the products available on Oracle Database Companion media for the Windows 32-bit platform.
Oracle Database Companion CD Installation Guide for Microsoft Windows (64-Bit) on Intel Itanium	B14325-01	Explains how to install and configure the products available on Oracle Database Companion media for Windows on the Intel Itanium platform.
Oracle Database Installation Guide for Linux x86	B15660-01	Provides platform-specific preinstallation, installation, postinstallation, and component information for your single instance Oracle Database 10*g* database software on the Linux x86 platform.
Oracle Database Quick Installation Guide for Linux x86	B15661-01	Describes how to quickly install Oracle Database for Linux x86 systems.
Oracle Database Client Installation Guide for Linux x86	B15662-01	Describes how to install Linux x86 clients.
Oracle Database Client Quick Installation Guide for Linux x86	B15663-01	Describes how to quickly install Oracle Database Client on Linux x86 systems.
Oracle Database Companion CD Installation Guide for Linux x86	B15664-01	Explains how to install and configure the products available on Oracle Database Companion media for the Linux x86 platform.
Oracle Database Companion CD Quick Installation Guide for Linux x86	B15665-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products on Linux x86 systems.
Oracle Database Installation Guide for Linux x86-64	B15667-01	Describes how to install Oracle Database 10*g* products on Linux x86-64 systems.

Table 2–8 (Cont.) Installation

Documentation	Part Number	Description
Oracle Database Quick Installation Guide for Linux x86-64	B15668-01	Describes how to quickly install Oracle Database 10*g* products on Linux x86-64 systems.
Oracle Database Client Installation Guide for Linux x86-64	B15669-01	Describes how to install Oracle Database Client on Linux x86-64 systems.
Oracle Database Client Quick Installation Guide for Linux x86-64	B15670-01	Describes how to quickly install Oracle Database Client on Linux x86-64 systems.
Oracle Database Companion CD Installation Guide for Linux x86-64	B15671-01	Describes how to install Oracle Database 10*g* products and Oracle Database 10*g* companion products on Linux x86-64 systems.
Oracle Database Companion CD Quick Installation Guide for Linux x86-64	B15672-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products on Linux x86-64 systems.
Oracle Database Installation Guide for Linux Itanium	B15674-01	Describes how to install Oracle Database 10*g* products on Linux Itanium systems.
Oracle Database Quick Installation Guide for Linux Itanium	B15675-01	Describes how to quickly install Oracle Database 10*g* products on Linux Itanium systems.
Oracle Database Client Installation Guide for Linux Itanium	B15676-01	Describes how to install Oracle Database Client on Linux Itanium systems.
Oracle Database Client Quick Installation Guide for Linux Itanium	B15677-01	Describes how to quickly install Oracle Database Client on Linux Itanium systems.
Oracle Database Companion CD Installation Guide for Linux Itanium	B15678-01	Describes how to install Oracle Database 10*g* products and Oracle Database 10*g* companion products on Linux Itanium systems.
Oracle Database Companion CD Quick Installation Guide for Linux Itanium	B15679-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products on Linux Itanium systems.
Oracle Database Installation Guide for Microsoft Windows (x64)	B15681-01	Provides platform-specific preinstallation, installation, postinstallation, and component information for your single instance Oracle Database 10*g* database and client software on the Windows x64 platform.
Oracle Database Quick Installation Guide for Microsoft Windows (x64)	B15682-01	Explains how to install and configure the products available on Oracle Database Companion media for the Windows x64 platform.
Oracle Database Client Installation Guide for Microsoft Windows (x64)	B15684-01	Describes how to install Windows x64 clients.
Oracle Database Client Quick Installation Guide for Microsoft Windows (x64)	B15685-01	Describes how to quickly install Oracle Database Client on Windows x64 systems.
Oracle Database Companion CD Installation Guide for Microsoft Windows (x64)	B15686-01	Explains how to install and configure the products available on Oracle Database Companion media for the Windows x64 platform.

Table 2–8 (Cont.) Installation

Documentation	Part Number	Description
Oracle Database Companion CD Quick Installation Guide for Microsoft Windows (x64)	B15687-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products on the Windows x64 systems.
Oracle Database Installation Guide for Solaris Operating System (SPARC 64-Bit)	B15690-01	Describes how to install Oracle Database 10*g* products on Solaris 64-bit systems.
Oracle Database Quick Installation Guide for Solaris Operating System (SPARC 64-Bit)	B15691-01	Describes how to quickly install Oracle Database 10*g* products on Solaris 64-bit systems.
Oracle Database Client Installation Guide for Solaris Operating System (SPARC 64-Bit)	B15692-01	Describes how to install Oracle Database Client on Solaris 64-bit systems.
Oracle Database Client Quick Installation Guide for Solaris Operating System (SPARC 64-Bit)	B15693-01	Describes how to quickly install Oracle Database Client on Solaris 64-bit systems.
Oracle Database Companion CD Installation Guide for Solaris Operating System (SPARC 64-Bit)	B15694-01	Describes how to install Oracle Database 10*g* products and Oracle Database 10*g* companion products on Solaris 64-bit systems.
Oracle Database Companion CD Quick Installation Guide for Solaris Operating System (SPARC 64-Bit)	B15695-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products on Solaris 64-bit systems.
Oracle Database Installation Guide for Solaris Operating System (x86)	B15697-01	Describes how to install Oracle Database 10*g* on Solaris x86 systems.
Oracle Database Quick Installation Guide for Solaris Operating System (x86)	B15698-01	Describes how to quickly install Oracle Database 10*g* on Solaris x86 systems.
Oracle Database Client Installation Guide for Solaris Operating System (x86)	B15699-01	Describes how to install Oracle Database Client on Solaris x86 systems.
Oracle Database Client Quick Installation Guide for Solaris Operating System (x86)	B15700-01	Describes how to quickly install Oracle Database Client on Solaris x86 systems.
Oracle Database Companion CD Installation Guide for Solaris Operating System (x86)	B15701-01	Describes how to install Oracle Database 10*g* products and Oracle Database 10*g* companion products on Solaris x86 systems.
Oracle Database Companion CD Quick Installation Guide for Solaris Operating System (x86)	B15702-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products on Solaris x86 systems.
Oracle Database Installation Guide for Solaris Operating System (x86-64)	B15704-01	Describes how to install Oracle Database 10*g* on Solaris x86-64 systems.

Table 2–8 (Cont.) Installation

Documentation	Part Number	Description
Oracle Database Quick Installation Guide for Solaris Operating System (x86-64)	B15705-01	Describes how to quickly install Oracle Database 10*g* on Solaris x86-64 systems.
Oracle Database Client Installation Guide for Solaris Operating System (x86-64)	B15706-01	Describes how to install Oracle Database Client on Solaris x86-64 systems.
Oracle Database Client Quick Installation Guide for Solaris Operating System (x86-64)	B15707-01	Describes how to quickly install Oracle Database Client on Solaris x86-64 systems.
Oracle Database Companion CD Installation Guide for Solaris Operating System (x86-64)	B15708-01	Describes how to install Oracle Database 10*g* products and Oracle Database 10*g* companion products on Solaris x86-64 systems.
Oracle Database Companion CD Quick Installation Guide for Solaris Operating System (x86-64)	B15709-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products on Solaris x86-64 systems.
Oracle Database Installation Guide for hp-ux PA-RISC (64-Bit)	B19068-01	Describes how to install Oracle Database 10*g* on hp-ux PA-RISC 64-bit systems.
Oracle Database Quick Installation Guide for hp-ux PA-RISC (64-Bit)	B19081-01	Describes how to quickly install Oracle Database 10*g* on hp-ux PA-RISC 64-bit systems.
Oracle Database Client Installation Guide for hp-ux PA-RISC (64-Bit)	B19070-01	Describes how to install Oracle Database Client on hp-ux PA-RISC 64-bit systems.
Oracle Database Client Quick Installation Guide for hp-ux PA-RISC (64-Bit)	B19071-01	Describes how to quickly install Oracle Database Client on hp-ux PA-RISC 64-bit systems.
Oracle Database Companion CD Installation Guide for hp-ux PA-RISC (64-Bit)	B19072-01	Describes how to install Oracle Database 10*g* products and Oracle Database 10*g* companion products on hp-ux PA-RISC 64-bit systems.
Oracle Database Companion CD Quick Installation Guide for hp-ux PA-RISC (64-Bit)	B19073-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products on hp-ux PA-RISC 64-bit systems.
Oracle Database Installation Guide for AIX 5L Based Systems (64-Bit)	B19075-01	Describes how to install Oracle Database 10*g* on AIX 5L Based 64-bit systems.
Oracle Database Quick Installation Guide for AIX 5L Based Systems (64-Bit)	B19076-01	Describes how to quickly install Oracle Database 10*g* on AIX 5L Based 64-bit systems.
Oracle Database Client Installation Guide for AIX 5L Based Systems (64-Bit)	B19077-01	Describes how to install Oracle Database Client on AIX 5L Based 64-bit systems.

Table 2–8 (Cont.) Installation

Documentation	Part Number	Description
Oracle Database Client Quick Installation Guide for AIX 5L Based Systems (64-Bit)	B19078-01	Describes how to quickly install Oracle Database Client on AIX 5L Based 64-bit systems.
Oracle Database Companion CD Installation Guide for AIX 5L Based Systems (64-Bit)	B19079-01	Describes how to install Oracle Database 10*g* products and Oracle Database 10*g* companion products on AIX 5L Based 64-bit systems.
Oracle Database Companion CD Quick Installation Guide for AIX 5L Based Systems (64-Bit)	B19080-01	Describes how to quickly install Oracle Database 10*g* products and Oracle Database 10*g* companion products on AIX 5L Based 64-bit systems.

2.1.1.9 Java Development Documentation

Table 2–9 provides a list of Java Development documentation included in Oracle Database 10*g* Release 2.

Table 2–9 Java Development

Documentation	Part Number	Description
Oracle Database Java Developer's Guide	B14187-01	Provides a general overview on how to develop, load, and execute Java applications in the database.
Oracle Database JPublisher User's Guide	B14188-01	Describes the JPublisher utility for Java programmers who want classes in their applications to correspond to SQL or PL/SQL entities or server-side Java classes. In Oracle Database 10*g*, JPublisher also provides features supporting Web services call-ins to the database and call-outs from the database.
Oracle Database JDBC Developer's Guide and Reference	B14355-01	Provides guidance for developers of JDBC-based applications and applets. This document can be read by anyone with an interest in JDBC programming, but assumes at least some prior knowledge of Java, Oracle PL/SQL, and Oracle databases.
Oracle Database SQLJ Developer's Guide and Reference	B16018-01	Describes how SQLJ enables application programmers to embed SQL operations in Java code in a way that is compatible with the Java design philosophy.

2.1.1.10 Networking Documentation

Table 2–10 provides a list of Networking documentation included in Oracle Database 10*g* Release 2.

Table 2–10 Networking

Documentation	Part Number	Description
Oracle Database Net Services Administrator's Guide	B14212-01	Explains how to plan, configure, and manage enterprise-wide connectivity with Oracle Net Services.
Oracle Database Net Services Reference	B14213-01	Contains a complete listing and description of the control utility commands and configuration file parameters available for managing components of Oracle Net Services.
Oracle Database Security Guide	B14266-01	Provides a comprehensive overview of security for this database release, featuring security as a necessary, primary focus for Oracle installations. It provides the conceptual framework for understanding security requirements and threats as well as introducing the features, techniques, and policies for meeting those needs with Oracle products. The audience for this book includes system managers as well as users, database administrators, and application developers.
Oracle Label Security Administrator's Guide	B14267-01	Describes how to use Oracle Label Security to protect sensitive data. It explains the basic concepts behind label-based security and provides examples to show how it is used.
Oracle Database Advanced Security Administrator's Guide	B14268-01	Provides an overview and detailed configuration information about database advanced security features. These features include network encryption, strong authentication (RADIUS, Kerberos, and SSL), and centralized user management that uses the Oracle Identity Management infrastructure (Enterprise User Security). The audience for this guide includes network security and enterprise user security administrators.
Oracle Database Enterprise User Administrator's Guide	B14269-01	Describes how to implement, configure, and administer Oracle Database users in Oracle Internet Directory, the directory service provided by the Oracle Identity Management platform.

2.1.1.11 OLAP Documentation

Table 2–11 provides a list of OLAP documentation included in Oracle Database 10*g* Release 2.

Table 2–11 OLAP

Documentation	Part Number	Description
Oracle OLAP DML Reference	B14346-01	Contains a complete description of the OLAP Data Manipulation Language (OLAP DML) used to define and manipulate analytic workspace objects.
Oracle OLAP Developer's Guide to the OLAP API	B14347-01	Introduces the Oracle OLAP API, a Java application programming interface for Oracle OLAP, which is used to perform online analytical processing of the data stored in Oracle Database. Describes the API and how to discover metadata, create queries, and retrieve data.
Oracle OLAP Java API Reference	B14348-01	Describes the classes and methods in the Oracle OLAP Java API for querying analytic workspaces and relational data warehouses.

Table 2–11 (Cont.) OLAP

Documentation	Part Number	Description
Oracle OLAP Application Developer's Guide	B14349-01	Explains how SQL and Java applications can extend their analytic processing capabilities by using Oracle OLAP in the Enterprise Edition of Oracle Database.
Oracle OLAP Reference	B14350-01	Explains the syntax of PL/SQL packages and types and the column structure of views related to Oracle OLAP.
Oracle OLAP Analytic Workspace Java API Reference	B14351-01	Describes the classes and methods in the Oracle OLAP Java API for building and maintaining analytic workspaces.

2.1.1.12 Real Application Clusters Documentation

Table 2–12 provides a list of Real Application Clusters documentation included in Oracle Database 10*g* Release 2.

Table 2–12 Real Application Clusters

Documentation	Part Number	Description
Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide	B14197-01	Explains how to administer Oracle Real Application Clusters.

2.1.1.13 Release Notes Documentation

Table 2–13 provides a list of Release Notes documentation included in Oracle Database 10*g* Release 2.

Table 2–13 Release Notes

Documentation	Part Number	Description
Oracle Database Release Notes for Microsoft Windows (32-Bit)	B14264-01	Provides important last minute features and changes specific to the Windows 32-bit platform not included in Oracle Database Documentation Library.
Oracle Database Release Notes for Microsoft Windows (64-Bit) on Intel Itanium	B14265-01	Provides a description of important last minute features and changes specific to the Intel Itanium platform not included in Oracle Database Documentation Library.
Oracle Database Release Notes for Linux x86	B15659-01	Provides important last minute features and changes specific to the Linux x86 platform not included in Oracle Database Documentation Library.
Oracle Database Release Notes for Linux x86-64	B15666-01	Provides a description of important last minute features and changes specific to the Linux x86-64 platform not included in Oracle Database Documentation Library.
Oracle Database Release Notes for Linux Itanium	B15673-01	Provides a description of important last minute features and changes specific to the Linux Itanium platform not included in Oracle Database Documentation Library.
Oracle Database Release Notes for Microsoft Windows (x64)	B15680-01	Provides important last minute features and changes specific to the Windows x64 platform not included in Oracle Database Documentation Library.

Table 2–13 (Cont.) Release Notes

Documentation	Part Number	Description
Oracle Database Release Notes for Solaris Operating System (SPARC 64-Bit)	B15689-01	Provides a description of important last minute features and changes specific to the Solaris 64-bit platform not included in Oracle Database Documentation Library.
Oracle Database Release Notes for Solaris Operating System (x86)	B15696-01	Provides a description of important last minute features and changes specific to the Solaris x86 platform not included in Oracle Database Documentation Library.
Oracle Database Release Notes for Solaris Operating System (x86-64)	B15703-01	Provides a description of important last minute features and changes specific to the Solaris x86-64 platform not included in Oracle Database Documentation Library.
Oracle Enterprise Manager Database Control Release Notes	B16248-01	Provides a description of important last minute features and changes specific to Oracle Enterprise Manager Database Control not included in Oracle Database Documentation Library.
Oracle Database Release Notes for hp-ux PA-RISC (64-Bit)	B19067-01	Provides a description of important last minute features and changes specific to the hp-ux PA-RISC 64-bit platform not included in Oracle Database Documentation Library.
Oracle Database Release Notes for AIX 5L Based Systems (64-Bit)	B19074-01	Provides a description of important last minute features and changes specific to the AIX 5L Based 64-bit platform not included in Oracle Database Documentation Library.

2.1.1.14 Text Documentation

Table 2–14 provides a list of Text documentation included in Oracle Database 10*g* Release 2.

Table 2–14 Text

Documentation	Part Number	Description
Oracle Text Application Developer's Guide	B14217-01	Describes how to build an application with Oracle Text, such as a text query application or document classification system. Examples are provided for creating a text table, indexing, and querying. This book also contains information about query tuning, document presentation, and using a thesaurus in your application.
Oracle Text Reference	B14218-01	Contains reference information for Oracle Text, including Oracle Text SQL statements, operators, supplied PL/SQL packages, and views. Examples are provided for using the PL/SQL packages and operators.

2.1.1.15 Ultra Search Documentation

Table 2–15 provides a list of Ultra Search documentation included in Oracle Database 10*g* Release 2.

Table 2–15 Ultra Search

Documentation	Part Number	Description
Oracle Ultra Search API Reference	B14221-01	Describes Java APIs for Ultra Search.
Oracle Ultra Search Administrator's Guide	B14222-01	Describes how to build web-based query applications using Oracle Ultra Search. Topics include crawling, indexing, and searching text content in databases or HTML pages. Java Server Pages (JSP) web-application examples are provided.

2.1.1.16 Windows Documentation

Table 2–16 provides a list of Windows documentation included in Oracle Database 10*g* Release 2.

Table 2–16 Windows

Documentation	Part Number	Description
*Pro*COBOL Getting Started for Microsoft Windows (32-bit)*	A96113-03	Provides an addendum to the Pro*COBOL Programmer's Guide, containing information specific to the Windows 32-bit platform.
Oracle Database Platform Guide for Microsoft Windows (32-Bit)	B14304-01	Provides platform-specific information for configuring and using Oracle Database software on the Windows 32-bit platform. Topics discussed include: architecture, services, tools, configuration, administration, tuning, backup, accounts, passwords, user authentication, registry parameters, and application development.
Oracle Database Platform Guide for Microsoft Windows (64-Bit) on Intel Itanium	B14305-01	Provides platform-specific information for configuring and using Oracle Database software on Intel Itanium. Topics discussed include: architecture, services, tools, configuration, administration, tuning, backup, accounts, passwords, user authentication, registry parameters, and application development.
Oracle Database Extensions for .NET Developer's Guide	B14306-01	Provides introductory, installation, post-installation configuration, and usage information for Oracle Data Provider for .NET.
Oracle Data Provider for .NET Developer's Guide	B14307-01	Describes provider-specific features and properties of ODP.NET. Oracle Data Provider for .NET (ODP.NET) is an implementation of Microsoft's ADO.NET interface and provides access to Oracle databases.
Oracle Objects for OLE C++ Class Library Developer's Guide	B14308-01	Describes the concepts and the object classes. The Oracle Objects for OLE C++ Class Library is a collection of C++ classes that provide programmatic access to the OO4O automation server.
Oracle Objects for OLE Developer's Guide	B14309-01	Describes the concepts important to OO4O and the object classes. Oracle Objects for OLE (OO4O) allows easy access to data stored in Oracle databases with any programming or scripting language that supports the Microsoft COM Automation and ActiveX technology.
Oracle COM Automation Feature Developer's Guide	B14310-01	Describes how Oracle COM Automation Feature provides a mechanism to manipulate COM objects through PL/SQL. This developer's guide provides introductory, installation, postinstallation configuration, and usage information for Oracle COM Automation Feature.
Oracle Provider for OLE DB Developer's Guide	B14311-01	Describes provider-specific features and properties of OraOLEDB. Oracle Provider for OLE DB (OraOLEDB) is based on an open standard that uses a set of Component Object Model (COM) interfaces for data access, in this case to Oracle databases.
Oracle Services for Microsoft Transaction Server Developer's Guide	B14320-01	Provides introductory, installation, configuration, usage, and administration information for using Oracle Services for Microsoft Transaction Server. Microsoft Transaction Server is a proprietary Component Object Model (COM) transaction processing system that runs on an Internet or network server.
*Pro*C/C++ Getting Started for Microsoft Windows (32-bit)*	B14321-01	Provides an addendum to the Pro*C/C++ Programmer's Guide, containing information specific to the Windows 32-bit platform.
Oracle Database Platform Guide for Microsoft Windows (x64)	B15688-01	Provides platform-specific information for configuring and using Oracle Database software on the Windows x64 platform. Topics discussed include: architecture, services, tools, configuration, administration, tuning, backup, accounts, passwords, user authentication, registry parameters, and application development.

2.1.1.17 Workflow Documentation

Table 2–17 provides a list of Workflow documentation included in Oracle Database 10*g* Release 2.

Table 2–17 Workflow

Documentation	Part Number	Description
Oracle Workflow Administrator's Guide	B15852-02	Describes how to administer Oracle Workflow.
Oracle Workflow Developer's Guide	B15853-01	Describes how to develop applications using Oracle Workflow.
Oracle Workflow User's Guide	B15854-02	Describes how to use Oracle Workflow.
Oracle Workflow API Reference	B15855-02	Describes the APIs used by Oracle Workflow.

Index

A

Access Advisor, 1-41
Active Session History (ASH)
  analyzing system performance, 1-58
  finding blocker and waiter identities, 1-57
ADD_FILE command
  FILESIZE value, 1-34
ADDM
  performance analysis, 1-57
aggregation performance improvements, 1-54
ALTER DISKGROUP command
  WAIT option, 1-59
ALTER SYSTEM command, 1-20
analysis
  ADDM performance, 1-57
  of system performance, 1-58
  Predictive Analytics Package, 1-28
  using CSSCAN, 1-3
  using decision tree algorithm, 1-28
analytic functions
  and the MODEL clause, 1-24
API calls
  starting a database, 1-2
  stopping a database, 1-2
application management
  using Rules Manager, 1-10
archive logs
  automatic deletion, 1-18
Asian language support
  CTXRULE, 1-31
ASM
  batch selection of multiple disks, 1-60
  configuring disk groups
    using ASM Disk Group Management, 1-40
  creating an instance
    using ASM Disk Group Management, 1-40
  installing, 1-47
  managing disks, 1-59
  managing files
    using ASMCMD, 1-59
  migrate non-ASM databases to, 1-40
  moving application data, 1-61
  protocol for connection to database, 1-60
  rebalance process, 1-59
  standalone disk group management, 1-40
  status and space usage, 1-59
  support for multiple database releases, 1-60
ASM Disk Group Management
  configuring disk groups, 1-40
  creating an ASM instance, 1-40
ASSM
  AUTO default, 1-64
Asynchronous Change Data Capture, 1-26
asynchronous commit, 1-53
asynchronous redo transmission
  optimized, 1-16
audit trail
  and XML, 1-36
AUTO_FILTER filter
  for Oracle Text, 1-31
Automatic Database Diagnostic Monitor
  See ADDM
Automatic Segment Advisor
  reorganizing data fragmentation, 1-64
Automatic Segment Space Management
  See ASSM
Automatic Shared Memory Management, 1-62
Automatic Storage Management
  See ASM
Automatic Undo Management, 1-55
Automatic Workload Repository (AWR)
  and SQL statement collection, 1-58
  performance diagnosis, 1-58
  transferring data, 1-58
AUTOTRACE EXPLAIN query, 1-35

B

backup
  database, 1-23
  encryption, 1-19
  incremental, 1-22
  storage, 1-21
Backup Management
  in Enterprise Manager, 1-42
batch selection
  for CHECK, 1-60
  for DELETE, 1-60
  for RESIZE, 1-60
bitmap indexes
  and Intra-Partition DML, 1-24

block validation, 1-19
buffered messaging, 1-48
business registry
through Web Services wizard, 1-8

C

cache coherency
monitoring, 1-44
statistics and diagnostics, 1-43
centralized password policies, 1-38
Change Data Capture
distributed asynchronous capabilities, 1-26
channel failover
and RAC instances, 1-18
Character Set Scanner
See CSSCAN
character sets
and Instant Client Light, 1-2
Unicode, 1-3
charts
asynchronous refresh, 1-7
Cluster Bar, 1-6
Dial, 1-6
Percentage Bar, 1-6
Stacked Bar, 1-6
checkpointing improvements, 1-53
cleanup system statistics, 1-57
cloning RAC nodes, 1-47
Cluster Bar charts, 1-6
cluster systems
and application system components, 1-30
and Oracle components, 1-30
high availability, 1-30
management, 1-30
cluster tracing and diagnostics, 1-30
Cluster Verification Utility, 1-30
clusterware file redundancy, 1-30
collection operator
and data objects, 1-11
Common Language Runtime (CLR)
stored procedure integration, 1-66
Common Sort mapping function, 1-3
Common Territory mapping function, 1-3
Common Time Zone mapping function, 1-3
configuration policies, 1-44
connect role privileges, 1-37
control file, 1-19
support for RAC instances, 1-21
Create Application wizard
using spreadsheet data, 1-7
CREATE command
of WRAP utility, 1-9
CREATE_WRAPPED command
of WRAP utility, 1-9
CSSCAN
COLUMN parameter, 1-3
QUERY parameter, 1-3
selective analysis of data, 1-3
CTX_DOC.POLICY_SNIPPET procedure, 1-31
CTX_DOC.SNIPPET procedure, 1-31
CTXRULE Asian language support, 1-31

D

data fragmentation
reorganizaing
using Automatic Segment Advisor, 1-64
Data Guard
and fast-start failover, 1-16
automatic reinstantiation of primary database, 1-16
flashback, 1-21
Redo Apply failover, 1-16
SQL Apply failover, 1-17
tracking runtime performance, 1-21
Data Mining
decision tree algorithm, 1-28
Predictive Analytics Package, 1-28
scoring functions, 1-28
data objects
and collection operator, 1-11
Data Pump
and tablespace storage allocation, 1-34
DATA_PUMP_DIR directory object, 1-33
dumpfile size, 1-34
Export, 1-33
dumpfile compression, 1-33
extended job status display, 1-34
job management, 1-34
monitoring jobs, 1-34
Object ID (OID) transform parameter, 1-34
PCTSPACE transform parameter, 1-34
upgrade management, 1-33
data recovery
using flashback, 1-21
data type conversion
LONG to LOB, 1-53
data types
Unicode, 1-3
DATA_PUMP_DIR directory object, 1-33
database backup, 1-23
database change notification, 1-54
database event management
using Rules Manager, 1-10
database migration, 1-17
non-ASM to ASM, 1-40
database object type
reorganizing, 1-23
database password policies, 1-38
database patching
using OPatch, 1-46
database recovery
Restore Point, 1-23
Database Resource Manager, 1-61, 1-62
database upgrade
release 10.1 to release 10.2, 1-41
datafile
dropping, 1-65
DB_BLOCK_CHECKSUM parameter

MAX value, 1-19
DB_FILE_MULTIBLOCK_READ_COUNT
initialization parameter, 1-55
DBA_ENABLED_TRACES dynamic view
performance monitoring, 1-58
DBMS_DATAPUMP API calls
GET_DUMPFILE_INFO, 1-34
WAIT_FOR_JOB, 1-34
DBMS_OUTPUT package, 1-9, 1-35
DBMS_SPACE_ADMIN package, 1-20
DBMS_XPLAN
replaces AUTOTRACE EXPLAIN, 1-35
decision tree algorithm, 1-28
deprecated
XMLDOM
now DBMS_XMLDOM, 1-14
XMLPARSER
now DBMS_XMLPARSER, 1-14
XMLPROCESSOR
now DBMS_XMLPROCESSOR, 1-14
diagnosing memory allocation, 1-62
Dial charts, 1-6
DICOM
medical imaging format support, 1-32
Directory Objects
management, 1-43
distributed environment
LOB support, 1-11
Distributed Transactions
and RAC, 1-20
DML error logging table, 1-26
downstream database capture, 1-48
dropping an empty datafile, 1-65
dropping partitioned tables, 1-25
dynamic partition pruning, 1-25
dynamic RMAN channel allocation, 1-18

E

End-to-End Performance Monitoring, 1-58
Enterprise Manager
alert thresholds, 1-56
backup and recovery, 1-41
Backup Management, 1-42
database administration, 1-42
database auditing administration, 1-42
Directory Objects, 1-43
displaying activity history, 1-56
displaying SQL statement history, 1-63
external object cloning, 1-42
Load Data wizard, 1-46
managing Streams, 1-50
monitoring slow systems, 1-56
Net Services management, 1-55
Oracle Spatial management of, 1-44
performance graph normalization, 1-56
Performance Monitoring, 1-56
reading performance statistics from SGA, 1-57
restore and recovery, 1-45
Streams Advanced Queuing management, 1-50
undo retention, 1-42
XML management of, 1-15
Enterprise Manager Configuration plug-in, 1-40
event scheduling, 1-39
Expression Filter, 1-11

F

failover
Redo Apply, 1-16
SQL Apply, 1-17
Fast Connection Caching
on OracleXADataSources, 1-5
fast partition split
for index organized tables, 1-54
fast refresh
Partition Change Tracking Refresh (PCT), 1-26
fast-start failover, 1-16
file groups repository, 1-38
flashback
data recovery, 1-21
primary and standby databases, 1-21
functions
AppendChildXML(), 1-12
DeleteXML(), 1-12
InsertXML(), 1-12
InsertXMLBefore(), 1-12
XMLQUERY, 1-12
XMLTABLE, 1-12

G

GDK
locale mapping functions
and PL/SQL packages, 1-3
mapping functions
Common Sort, 1-3
Common Territory, 1-3
Common Time Zone, 1-3
Translated Locale Names, 1-3
Globalization Development Kit
See GDK

H

HARD
disk storage for RMAN backups, 1-21
Hardware Assisted Resilient Data
See HARD
high availability
and fast-start failover, 1-16
of cluster systems, 1-30
host cluster targets
discovering and updating, 1-41
hot mining, 1-48
HTTPS
secure HTTP protocol, 1-15

I

image content metadata support, 1-32

imaging support
content metadata, 1-32
Java Advanced Imaging (JAI), 1-33
medical, 1-32
Implicit Connection Caching
on OracleXADataSources, 1-5
index organized tables
fast partition split, 1-54
indexing
AUTO_FILTER filter, 1-31
local partitioned indexes, 1-31
Instant Client Light
and character sets, 1-2
Instant Client SDK
and OCCI header files, 1-2
and OCI header files, 1-2
interconnect information statistics, 1-46
Intra-Partition DML
and bitmap indexes, 1-24
IOT support, 1-18
and Online Segment Shrink, 1-64

J

Java
and XPath support, 1-14
and XQuery support, 1-14
Java Advanced Imaging (JAI) support, 1-33
Java Architecture for XML Binding
See JAXB
Java Data Mining (JDM), 1-28
Java Database Connectivity
See JDBC
Java objects
and XML class names, 1-14
JavaNet
local naming, 1-50
JAXB
and XML class names, 1-14
JDBC
and caching XA Connections
on OracleXADataSources, 1-5
and ResultSet holdability, 1-5
and retrieval of auto-generated key, 1-5
runtime connection load balancing, 1-5
support for DML, 1-5
job chaining, 1-39

K

key management, 1-35
Key Word in Context
See KWIC
keywords
evaluating in a document, 1-31
KWIC
CTX_DOC.POLICY_SNIPPET procedure, 1-31
CTX_DOC.SNIPPET procedure, 1-31

L

Language and Character Set File Scanner
See LCSSCAN
language and character sets
LCSSCAN, 1-4
large clusters
performance monitoring, 1-44
Large Objects
See LOB
LCSSCAN
and HTML file support, 1-4
linguistic search and sort, 1-4
Linux
and ODBC Driver support, 1-1
Load Balancing Advisory, 1-29
Load Data wizard
in Enterprise Manager, 1-46
LOB
access statistics in dynamic performance views, 1-53
and Online Segment Shrink, 1-64
conversion from LONG data type, 1-53
distributed support, 1-11
handling in Streams apply process, 1-49
improved loading performance, 1-53
local naming
and JavaNet, 1-50
local partitioned indexes
online indexing of, 1-31
lock statistics, 1-43
logical change record (LCR)
transforming, 1-49
logical standby database
automatic deletion of archived logs, 1-18
creating, 1-22
recovery, 1-17
LogMiner
and IOT support, 1-18
LONG
conversion to LOB data type, 1-53
LONG to LOB data type conversion, 1-53

M

managing RMAN backups, 1-20
manual upgrades, 1-33
mapping functions
Common Sort, 1-3
Common Territory, 1-3
Common Time Zone, 1-3
Translated Locale Names, 1-3
master detail mini-application, 1-7
Master Detail wizard, 1-7
materialized views
and fast refresh, 1-26
and Query Rewrite, 1-26
matrix math
for PL/SQL, 1-9
medical imaging format support
DICOM, 1-32

memory allocation
diagnosing, 1-62
memory corruption checking, 1-19
messages
commit-time ordering, 1-48
migrating a database, 1-17
non-ASM to ASM, 1-40
MODEL clause
and upsert of cells, 1-24
using analytic functions, 1-24
using FOR loop
evaluation of, 1-24
monitoring slow systems
using Enterprise Manager, 1-56
multiblock read count
tuning, 1-55

N

Net Services
and Enterprise Manager, 1-55
NLS_COMP
LINGUISTIC parameter, 1-4

O

Object ID (OID) transform parameter, 1-34
objects
maximum number of partitions, 1-25
OCCI
header files and Instant Client SDK, 1-2
support for TAF, 1-2
OCFS
for Sun Solaris, 1-29
OCI
client versioning, 1-2
header files and Instant Client SDK, 1-2
ODBC Driver
support for Linux, 1-1
ODP.NET
connection pool administration, 1-67
Database Change Notification Service, 1-67
derived parameters support, 1-67
Fast Application Notification (FAN), 1-65
fast connection failover, 1-65
integration with RAC Load Balancing
Advisory, 1-65
LOB retrieval, 1-68
.NET Framework support, 1-66
query cancellation, 1-67
query timeout, 1-67
server-side support, 1-68
statement caching, 1-68
Virtual Private Database (VPD) support, 1-69
OJMS
nonpersistent messaging, 1-49
OLAP
calculated dimension members, 1-27
compressed cubes
incremental aggregation of, 1-27
semiadditive aggregation of, 1-27
Online Analytical Processing
See OLAP
Online Redefinition
parallel execution of LONG to LOB data
types, 1-53
online redefinition
of single partitions, 1-22
of tables, 1-22
Online Segment Shrink
extended to LOB and IOT, 1-64
OPatch tool, 1-46
Open Database Connectivity Driver
See ODBC Driver
operating system objects
security, 1-69
Oracle C++ Call Interface
See OCCI
Oracle Call Interface
See OCI
Oracle Cluster File System
See OCFS
Oracle Cluster Registry (OCR)
and clusterware file redundancy, 1-30
Oracle Database
and CSSCAN, 1-3
and LCSSCAN, 1-4
Unicode character set, 1-3
Oracle HTML DB
Cluster Bar charts, 1-6
Dial charts, 1-6
Percentage Bar charts, 1-6
Stacked Bar charts, 1-6
Oracle Java Message Service
See OJMS
Oracle Label Security
for RAC, 1-38
Oracle Wallet
secure external password storage, 1-37
OracleXADataSources
Fast Connection Caching, 1-5
Implicit Connection Caching, 1-5

P

parallel join bitmap filtering, 1-54
Partition Change Tracking Refresh (PCT), 1-26
partition pruning, 1-25
partitioned tables
dropping, 1-25
PCTSPACE transform parameter, 1-34
Percentage Bar charts, 1-6
Performance Monitoring
using Enterprise Manager, 1-56
Perl
Regular Expression and compatibility with, 1-11
physical standby database
activation, 1-16
convert to reporting database, 1-19
PL/SQL

and WRAP utility, 1-9
applications
Common Sort, 1-3
Common Territory, 1-3
Common Time Zone, 1-3
Translated Locale Names, 1-3
BLAS subprogram, 1-9
conditional compilation, 1-8
LAPACK subprogram, 1-9
matrix math, 1-9
obfuscation of dynamically generated source code, 1-9
packages
and GDK, 1-3
DBMS_OUTPUT, 1-9
UTL_I18N
TRANSLITERATE function, 1-4
UTL_NLA, 1-9
Predictive Analytics Package, 1-28
Pre-Upgrade Information Utility, 1-33
primary database
reinstantiation, 1-16
progress monitoring for advisors, 1-57
protocol for connections
between database and ASM, 1-60
proxy authentication, 1-36

Q

Query Rewrite
and materialized views, 1-26
querying XML, 1-12

R

RAC
and Oracle Label Security, 1-38
cloning nodes, 1-47
cluster tracing and diagnostics, 1-30
Cluster Verification Utility, 1-30
clusterware file redundancy, 1-30
Distributed Transactions, 1-20
failure detection, 1-20
install framework, 1-47
Load Balancing Advisory, 1-29
and ODP.NET, 1-65
silent install, 1-47
Real Application Clusters
See RAC
Redo Apply failover, 1-16
Regular Expression
and Perl compatibility, 1-11
reorganizing object types, 1-23
REPLACE command
of WRAP utility, 1-9
Restore Point, 1-23
Restore Statistics, 1-45
RMAN
compression during cloning, 1-46
RMAN backups
encryption, 1-19
incremental, 1-22
managing, 1-20
storage, 1-21
RMAN recovery
temporary datafiles, 1-23
Rules Manager
application management, 1-10
database event management, 1-10
runtime connection load balancing
and JDBC, 1-5

S

Scheduler enhancements, 1-39
Schema Constructs
and XPath Rewrite, 1-12
secure external password storage, 1-37
secure HTTP protocol (HTTPS), 1-15
secure operating system objects, 1-69
Secure Sockets Layer
See SSL
Segment Advisor, 1-45
and segment findings, 1-43
SGA_TARGET parameter
and shared memory management, 1-61
shared memory management
and Streams pool, 1-61
SHUTDOWN ABORT command, 1-20
Simple Object Access Protocol
See SOAP
single partitions
online redefinition, 1-22
skip level dimension hierarchies, 1-26
SOAP
and C, 1-13
and C++, 1-13
space management
verification, 1-20
space usage alerts
for tablespace, 1-65
Spatial
GeoRaster
compression of image data, 1-51
graph partitioning, 1-52
support for European Petroleum Survey Group (EPSG), 1-51
support for Western Europe, 1-51
topology data model
inserts and updates, 1-52
SPLIT command
and index organized tables, 1-54
SQL 2003 Standard compliance, 1-13
SQL Access Advisor
function-based indexes, 1-63
viewing recommendations, 1-63
SQL Apply
and IOT support, 1-18
automatic deletion of archived logs, 1-18

failover, 1-17
tracking runtime performance, 1-21
SQL Profile
using literal values, 1-63
SQL statistics
monitoring long running statements, 1-64
V$SQLSTATS view, 1-63
SQL string comparisons, 1-4
SQL Tuning Sets, 1-62
SQL workloads
transfer of, 1-62
SQL*Plus
and XQuery, 1-35
SQL/XML Standard compliance, 1-13
SSL
and secure HTTP protocol (HTTPS), 1-15
Stacked Bar charts, 1-6
starting a database
with API calls, 1-2
with large memory configurations, 1-20
STARTUP FORCE command, 1-20
statistics
restoring, 1-45
statistics collection
for fixed and dictionary objects, 1-45
stopping a database
with API calls, 1-2
storage policies, 1-44
Streams, 1-48
Advanced Queuing
and Enterprise Manager, 1-50
notification, 1-49
and Enterprise Manager, 1-50
and IOT support, 1-18
application upgrade, 1-50
database migration, 1-50
LOB handling, 1-49
parallel application of changes, 1-49
pool
and shared memory management, 1-61
sizing, 1-62
system statistics cleanup, 1-57

T

tables
online redefinition, 1-22
partitioned
dropping, 1-25
tablespaces
adding, 1-17
creating, 1-17
moving, 1-15
segment details, 1-45
segment display, 1-46
TAF
and OCCI support, 1-2
transferring data
using Automatic Workload Repository (AWR), 1-58
transferring SQL workloads, 1-62
Translated Locale Names mapping function, 1-3
TRANSLITERATE function
and UTL_I18N, 1-4
for Japanese Kana conversions, 1-4
Transparent Application Failover
See TAF
Transparent Application Failover (TAF)
notification of, 1-50
transparent data encryption, 1-35
Transportable Tablespace
adding, 1-17
creating, 1-17
creating from backup, 1-38
moving XMLType tables, 1-15

U

UDDI
and web services, 1-8
Undo Advisor, 1-43
undo retention
enabling and disabling, 1-42
tuning, 1-55
Unicode
character set and Oracle Database, 1-3
data type literals, 1-3
multilingual Unicode data, 1-3
Unicode character set
Oracle Database, 1-3
Universal Description, Discovery, and Integration
See UDDI
unlock statistics, 1-43
upsert ability in MODEL clause, 1-24

V

V$ADVISOR_PROGRESS view, 1-57
V$ASM fixed view, 1-59
V$ASM_CLIENT view
COMPATIBLE_VERSION column, 1-60
SOFTWARE_VERSION column, 1-60
V$SESSION view
performance monitoring, 1-58
V$SQLSTATS view, 1-63
Voting Disk files
and clusterware file redundancy, 1-30

W

Web Services wizard
and UDDI, 1-8
create forms and reports, 1-7
wizards
Create Application, 1-7
Load Data, 1-46
Master Detail, 1-7
Web Services, 1-7
and UDDI, 1-8
WRAP utility
and CREATE command, 1-9

and CREATE_WRAPPED command, 1-9
and REPLACE command, 1-9

X

XAConnection caching
on OracleXADataSource, 1-5
XML
adding nodes to a document, 1-12
audit trail, 1-36
class names and Java objects, 1-14
functions
AppendChildXML(), 1-12
DeleteXML(), 1-12
InsertXML(), 1-12
InsertXMLBefore(), 1-12
management with Enterprise Manager, 1-15
querying, 1-12
removing nodes from a document, 1-12
SQL standard, 1-13
XML DB Repository
customizing metadata, 1-14
XMLDOM
deprecated
now DBMS_XMLDOM, 1-14
XMLPARSER
deprecated
now DBMS_XMLPARSER, 1-14
XMLPROCESSOR
deprecated
now DBMS_XMLPROCESSOR, 1-14
XMLType tables
moving, 1-15
XPath
and Java support, 1-14
XPath Rewrite, 1-13
and Schema Constructs, 1-12
XQuery
and Java support, 1-14
and SQL*Plus, 1-35
XQuery functions
XMLQUERY, 1-12
XMLTABLE, 1-12