Oracle® Database

New Features Guide

10*g* Release 1 (10.1)

Part No. B10750-01

December 2003

ORACLE®

Oracle Database New Features Guide, 10*g* Release 1 (10.1)

Part No. B10750-01

Contributors: Mateo Burtch, Paul Lane, Diana Lorentz, Mark Townsend, Sandra Cheevers, Michael Hartstein, Neal Brenner

Contents

2 Oracle Database 10*g* Documentation

Index

Send Us Your Comments

Oracle Database New Features Guide, 10*g* Release 1 (10.1)

Part No. B10750-01

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

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

If you find any errors or have any other suggestions for improvement, please indicate the 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
  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 book introduces the features, options, and enhancements that are new with this Oracle Database release. This book also provides readers with a list of titles and a brief description of the technical documentation available with this release.

This preface contains these topics:

Audience

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

Organization

This document contains the following chapters:

Chapter 1, "Oracle Database 10g, Release 1 New Features"

This chapter describes the new features of the Oracle Database 10*g*.

Chapter 2, "Oracle Database 10g Documentation"

This chapter lists technical documentation available with this release of the Oracle Database 10*g* and briefly describes the subject matter of each document.

Related Documentation

Printed documentation is available for sale in the Oracle Store at

`http://oraclestore.oracle.com/`

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

`http://otn.oracle.com/membership/`

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

`http://otn.oracle.com/documentation/`

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

This section describes the conventions used in the text and code examples of this documentation set. It describes:

- Conventions in Text
- Conventions in Code Examples
- Conventions for Windows Operating Systems

Conventions in Text

We use various conventions in text to help you more quickly identify special terms. The following table describes those conventions and provides examples of their use.

Convention	Meaning	Example
Bold	Bold typeface indicates terms that are defined in the text or terms that appear in a glossary, or both.	When you specify this clause, you create an **index-organized table**.
Italics	Italic typeface indicates book titles or emphasis.	*Oracle Database Concepts* Ensure that the recovery catalog and target database do *not* reside on the same disk.
`UPPERCASE monospace (fixed-width) font`	Uppercase monospace typeface indicates elements supplied by the system. Such elements include parameters, privileges, datatypes, RMAN keywords, SQL keywords, SQL*Plus or utility commands, packages and methods, as well as system-supplied column names, database objects and structures, usernames, and roles.	You can specify this clause only for a `NUMBER` column. You can back up the database by using the `BACKUP` command. Query the `TABLE_NAME` column in the `USER_TABLES` data dictionary view. Use the `DBMS_STATS.GENERATE_STATS` procedure.
`lowercase monospace (fixed-width) font`	Lowercase monospace typeface indicates executables, filenames, directory names, and sample user-supplied elements. Such elements include computer and database names, net service names, and connect identifiers, as well as user-supplied database objects and structures, column names, packages and classes, usernames and roles, program units, and parameter values. **Note:** Some programmatic elements use a mixture of UPPERCASE and lowercase. Enter these elements as shown.	Enter `sqlplus` to open SQL*Plus. The password is specified in the `orapwd` file. Back up the datafiles and control files in the `/disk1/oracle/dbs` directory. The `department_id`, `department_name`, and `location_id` columns are in the `hr.departments` table. Set the `QUERY_REWRITE_ENABLED` initialization parameter to `true`. Connect as `oe` user. The `JRepUtil` class implements these methods.
`lowercase italic monospace (fixed-width) font`	Lowercase italic monospace font represents placeholders or variables.	You can specify the *`parallel_clause`*. Run `U`*`old_release`*`.SQL` where *`old_release`* refers to the release you installed prior to upgrading.

Conventions in Code Examples

Code examples illustrate SQL, PL/SQL, SQL*Plus, or other command-line statements. They are displayed in a monospace (fixed-width) font and separated from normal text as shown in this example:

```
SELECT username FROM dba_users WHERE username = 'MIGRATE';
```

The following table describes typographic conventions used in code examples and provides examples of their use.

Convention	Meaning	Example
`[ ]`	Brackets enclose one or more optional items. Do not enter the brackets.	`DECIMAL (`*`digits`*` [ , `*`precision`*` ])`
`{ }`	Braces enclose two or more items, one of which is required. Do not enter the braces.	`{ENABLE | DISABLE}`
`|`	A vertical bar represents a choice of two or more options within brackets or braces. Enter one of the options. Do not enter the vertical bar.	`{ENABLE | DISABLE}` `[COMPRESS | NOCOMPRESS]`

`...`	Horizontal ellipsis points indicate either:
	■ That we have omitted parts of the code that are not directly related to the example
	■ That you can repeat a portion of the code
		`CREATE TABLE ... AS `*`subquery`*`;`
		
		`SELECT `*`col1`*`, `*`col2`*`, ... , `*`coln`*` FROM employees;`

.
.
.
	Vertical ellipsis points indicate that we have omitted several lines of code not directly related to the example.	`SQL> SELECT NAME FROM V$DATAFILE;`
		`NAME`
		`----------`
		`/fsl/dbs/tbs_01.dbf`
		`/fs1/dbs/tbs_02.dbf`
		.
		.
		.
		`/fsl/dbs/tbs_09.dbf`
		`9 rows selected.`

Other notation	You must enter symbols other than brackets, braces, vertical bars, and ellipsis points as shown.	`acctbal NUMBER(11,2);` `acct    CONSTANT NUMBER(4) := 3;`
Italics	Italicized text indicates placeholders or variables for which you must supply particular values.	`CONNECT SYSTEM/`*`system_password`* `DB_NAME = `*`database_name`*

Convention	Meaning	Example

UPPERCASE	Uppercase typeface indicates elements supplied by the system. We show these terms in uppercase in order to distinguish them from terms you define. Unless terms appear in brackets, enter them in the order and with the spelling shown. However, because these terms are not case sensitive, you can enter them in lowercase.	`SELECT last_name, employee_id FROM employees;`
		`SELECT * FROM USER_TABLES;`
		`DROP TABLE hr.employees;`

lowercase	Lowercase typeface indicates programmatic elements that you supply. For example, lowercase indicates names of tables, columns, or files.
	Note: Some programmatic elements use a mixture of UPPERCASE and lowercase. Enter these elements as shown.
		`SELECT last_name, employee_id FROM employees;`
		`sqlplus hr/hr`
		`CREATE USER mjones IDENTIFIED BY ty3MU9;`

Conventions for Windows Operating Systems

The following table describes conventions for Windows operating systems and provides examples of their use.

Convention	Meaning	Example
Choose Start >	How to start a program.	To start the Database Configuration Assistant, choose Start > Programs > Oracle - *HOME_NAME* > Configuration and Migration Tools > Database Configuration Assistant.
File and directory names	File and directory names are not case sensitive. The following special characters are not allowed: left angle bracket (<), right angle bracket (>), colon (:), double quotation marks ("), slash (/), pipe (|), and dash (-). The special character backslash (\) is treated as an element separator, even when it appears in quotes. If the file name begins with \\, then Windows assumes it uses the Universal Naming Convention.	`c:\winnt"\"system32 is the same as C:\WINNT\SYSTEM32`

Convention	Meaning	Example
`C:\>`	Represents the Windows command prompt of the current hard disk drive. The escape character in a command prompt is the caret (^). Your prompt reflects the subdirectory in which you are working. Referred to as the *command prompt* in this manual.	`C:\oracle\oradata>`
Special characters	The backslash (\) special character is sometimes required as an escape character for the double quotation mark (") special character at the Windows command prompt. Parentheses and the single quotation mark (') do not require an escape character. Refer to your Windows operating system documentation for more information on escape and special characters.	`C:\>exp scott/tiger TABLES=emp QUERY=\"WHERE job='SALESMAN' and sal<1600\"` `C:\>imp SYSTEM/`*`password`* `FROMUSER=scott TABLES=(emp, dept)`
HOME_NAME	Represents the Oracle home name. The home name can be up to 16 alphanumeric characters. The only special character allowed in the home name is the underscore.	`C:\> net start Oracle`*`HOME_NAME`*`TNSListener`

Convention	Meaning	Example
ORACLE_HOME and *ORACLE_BASE*	In releases prior to Oracle8i release 8.1.3, when you installed Oracle components, all subdirectories were located under a top level *ORACLE_HOME* directory. For Windows NT, the default location was `C:\orant`. This release complies with Optimal Flexible Architecture (OFA) guidelines. All subdirectories are not under a top level *ORACLE_HOME* directory. There is a top level directory called *ORACLE_BASE* that by default is `C:\oracle`. If you install the latest Oracle release on a computer with no other Oracle software installed, then the default setting for the first Oracle home directory is `C:\oracle\ora`*nn*, where *nn* is the latest release number. The Oracle home directory is located directly under *ORACLE_BASE*. All directory path examples in this guide follow OFA conventions. Refer to *Oracle Database Platform Guide for Windows* for additional information about OFA compliances and for information about installing Oracle products in non-OFA compliant directories.	Go to the *ORACLE_BASE*`\`*ORACLE_HOME*`\rdbms\admin` directory.

Documentation Accessibility

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

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

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

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

1

Oracle Database 10*g*, Release 1 New Features

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

Introduction

This version of the Oracle Database has been designed to meet two key goals: reducing the cost of manageability and delivering increased performance for all key workloads. In all areas, manageability and performance have been significantly improved.

In addition, new capabilities that support the provisioning and virtualization of computing resource in a Grid environment are provided, allowing users to take advantage of low cost hardware and storage.

Major new high-availability features are also provided, including new Flashback capabilities that reduce the downtime caused by human errors. In addition, support for rolling upgrades has been provided to reduce the downtime associated with database and application upgrades.

Security improvements in this release focus on improved administration of large numbers of users across an integrated end-to-end environment, and support for key security technologies.

Business intelligence solutions will benefit from improvements in SQL analytical, OLAP, and data mining capabilities. The ability to better support and more efficiently load and refresh larger databases is also improved. This version also provides support for key requirements in the area of bioinformatics.

Support for nonrelational data stored in a database has also been improved in this release, with new formats for multimedia data, addition of text classification capabilities, and improvements in search and retrieval capabilities. Spatial query performance has also been improved, and support for new spatial data models and data types have been provided XML capabilities have also been expanded to focus on the effective management of large amounts of XML data stored natively in the database.

Application development improvements include a new built-in application development environment, support for a high-performance and highly integrated XML capability across the entire technology stack, and a framework provided with the database that enables desktop and middle-tier applications to retrieve and extract data from the database using standard Web Services mechanisms.

Server Manageability

One of the major value propositions for this release of the Oracle Database is a significant reduction in the management cost of deploying and maintaining an Oracle-based solution. A number of major developments in this area incorporate

new techniques and methodologies across the entire database platform. The approach taken is threefold:

- Much of the complexity previously associated with configuring and deploying an Oracle database has been eliminated or simplified.
- Wherever possible, common operating procedures have been automated.
- Many of the activities required to tune and manage different areas of the Oracle Database have been greatly simplified.

Server Configuration

In this release, the overall footprint of the Oracle Database has been significantly reduced. For those upgrading from previous versions, new, easy upgrade capabilities have been provided that greatly reduce the steps required to upgrade an Oracle Database. Optimal database configuration has been made much easier with the initial release. Administrators need to be aware of only a small number of basic initialization parameters that they can use to configure and tune their environment. Many of the other tasks associated with database configuration have also been eliminated or automated. Any future patch sets required during the life of the database can be automatically identified, downloaded, and applied from Oracle*MetaLink*.

Simplified Database Install

The installation process for the database has been enhanced to reduce install time, system resource requirements (CPU, memory, and disk space), and number of installation CDs.

Automatic Storage Management Configuration

You can now use the Database Configuration Assistant (DBCA) to configure Automatic Storage Management for conventional and Oracle Real Application Clusters environments.

Automatic Oracle Real Application Clusters Services Configuration

You can now use the DBCA to automatically configure Oracle Real Application Clusters environments.

Simplified Upgrade for RAC and Oracle Parallel Server Databases

The Database Upgrade Assistant (DBUA) enables you to create services for workload management when you upgrade from a previous Oracle cluster software release.

Automatic Oracle Enterprise Manager Configuration

The Oracle Enterprise Manager repository, job, and event subsystems are now configured automatically, eliminating the need for manual setup.

Automatic Integrated Clusterware Installation

The Oracle Universal Installer (OUI) automatically installs and starts the integrated clusterware and its related components for Oracle Real Application Clusters.

Automatic Configuration of Recovery Area

The Database Configuration Assistant (DBCA) now automatically configures a default recovery area.

Out-of-the-Box LDAP Configuration

The new directory configuration by the DBCA eliminates the need for an `ldap.ora` file by using a network API.

Simplified Initialization Parameters

Initialization parameters are now divided into two groups, basic and advanced. In the vast majority of cases, it is necessary to set and tune only the basic parameters, of which there are 20 to 25, to get reasonable performance from the database. In rare situations, modification of the advanced parameters may be needed to achieve optimal performance.

Default User Tablespace

Database creation now enables specification of a default tablespace for storing permanent objects for all users created. This eliminates the need to use the `SYSTEM` tablespace.

Easy Upgrade

This feature reduces the number of steps needed to upgrade a database and its installed components, which greatly simplifies the database upgrade process.

Upgrade Information Tool

This new tool facilitates successful database upgrade by performing some preliminary checks on the existing database (for example, whether there is enough space, whether there are any obsolete initialization parameters, and so forth), and by providing an estimate of how long the database upgrade will take.

See *Oracle Real Application Clusters Installation and Configuration Guide* for information on Real Application Clusters installation in general as well as information on the Oracle Universal Installer Storage Configuration, the DBCA Service Management features of Real Application Clusters, and Oracle Universal Installer clusterware setup.

Network Management

This release greatly simplifies Oracle network management by eliminating the need to manually configure the Oracle networking environment. It is now possible to connect to a database from a client machine without creating the client configuration files. A number of shared server parameters are now automatically configured, and it is now possible to switch back and forth between dedicated and shared server mode online.

Back Up Directory Naming Entries to Local Naming File

You can now export directory naming entries into a local `tnsnames.ora` file. Clients can use the locally saved file when a directory server is unavailable or when clients have not been upgraded to use directory naming.

Dynamic Connection Manager Configuration

You can now easily change Connection Manager (CMAN) parameters without shutting down the CMAN process.

Easy Connect Naming Method

This feature simplifies network management by enabling a database connection to be completed without the need to first configure a net service name or perform other net configuration on the client. It also enables a connection to different databases or instances running on the same computer.

Improved Network Outage Detection

A new API provides an efficient event-based nonblocking abstraction that supports multiple connections. Oracle clients and servers can now avoid being blocked on a network I/O request.

Automatic LDAP Discovery for Clients

A client Oracle home no longer requires the directory usage configuration file (`ldap.ora`) to use LDAP naming. If an Oracle Internet Directory Server is registered with the Domain Name System (DNS), clients can automatically locate a directory server, and can be mapped to a particular Oracle Context, eliminating the need for local configuration. However, if the `ldap.ora` file is configured, it supersedes the auto-discovery mechanism.

Improved Connection Manager Access Rules

This feature eases configuration and provides more flexibility when specifying access rules for Connection Manager (CMAN). You can filter traffic for CMAN based on timeout, subnet mask, idle timeout, and other rule-level controls.

Automatic Shared Server Configuration

This feature improves flexibility if you are using two server models. You can now easily switch between shared server and dedicated server without having to start a dispatcher.

Simplified Shared Server Configuration Parameters

You no longer need to set as many initialization parameters for shared server environments, as the settings are now handled internally.

Manageability Infrastructure

A number of architectural enhancements have been introduced in this release to facilitate current and future self-management activities. They include a common infrastructure to store all auxiliary metadata and workload information, new monitoring tools and statistics to optimize performance tuning, and a more accurate and efficient wait statistics collection facility.

Database Features Usage Tracking

This release of the database automatically tracks usage (configuration, runtime, or both) of various database features. This enables the user to collect features' usage for future evaluation.

Enhanced Database Time Model

This feature enables the database to keep track of the time spent in performing internal operations such as parse, execute, input/output, and so on. This information is used by the database for making self-tuning decisions and makes it much easier to diagnose performance problems.

End-to-End Application Tracing

This feature simplifies the debugging of performance problems in multitier environments. It also makes debugging and tracing of applications, which use MODULE and ACTION name annotation, easier and more efficient. Further, it supports workload management based on service levels and priorities for applications using services.

JDBC: End-to-End Tracing

Java Database Connectivity (JDBC) enables you to monitor, trace, and correlate resource consumption of Java, Web Applications and EJB components down to the database operation level.

SYSAUX Tablespace

This new system-owned tablespace provides a centralized location for all auxiliary database metadata that does not reside in the SYSTEM tablespace. It reduces the number of tablespaces created by default, both in the seed database and in user-defined databases.

Automatic Routine Administration Tasks

The database Scheduler can schedule routine administration tasks, such as collecting optimizer statistics during a pre-defined maintenance window.

Server-Generated Alerts

This release of the database will proactively send alerts and notifications to administrators when a problem is anticipated or any of the user-selected metrics exceeds a defined threshold.

Automatic Workload Respository

A new, embedded, and completely self-managing repository captures workload information and performance-related statistics, thereby reducing administrative costs. The database uses information contained in the repository for all self-management activities.

Enhanced Wait Model

The enhanced wait model facilitates performance diagnosis. It enables you to determine what sessions are waiting for, keeps a history of waits and wait durations per session, and maintains wait statistics for SQL statements in a dynamic performance view.

Threshold-Based Alerts

New instrumentation for various server-produced, threshold-based alerts is now tightly integrated with the Oracle Enterprise Manager notification framework.

Storage Management

This release adds several features to make the management of database storage more simplified, flexible and automated. The first of these features is Automatic Storage Management, which is a vertically integrated volume manager and file system built for Oracle data files. Two other features that help simplify management of database storage are rename tablespace and multiple default temporary tablespaces.

Multiple Default Temporary Tablespace Support for SQL Operations

You can now define more than one default temporary tablespace, and a single SQL operation can use more than one temporary tablespace for sorting. This feature enables you to create indexes on very large tables without being limited by the size of one tablespace, because the sort operation during index creation can be distributed across multiple tablespaces.

Automatic Storage Management

Automatic Storage Management automates and simplifies the optimal layout of datafiles, control files, and log files. Database files are automatically distributed across all available disks, and database storage is rebalanced whenever the storage configuration changes. This feature also provides redundancy through the mirroring of database files.

Rename Tablespace

You can now rename a tablespace. You no longer have to create a new tablespace, copy the contents from the old tablespace, and drop the old tablespace. This feature makes it easier, for example, to migrate a dictionary-managed tablespace to be locally managed, or to transport a tablespace to a database that already contains a tablespace of the same name.

See Also:

Oracle Database Concepts for more general information

Oracle Database Administrator's Guide for information on multiple default temporary tablespaces, Automatic Storage Management, and renaming tablespaces

Oracle Database SQL Reference for information on creating multiple temporary tablespaces, creating diskgroups, and renaming tablespaces

Space, Object, and Transaction Management

New capabilities have been introduced to simplify object, space and transaction management. Proactive monitoring features have been introduced, capabilities for more efficient space management have been added, manual database administrator tasks have been automated, and new server advisors have been provided to facilitate advice-based intelligent management.

Automatic Undo Retention Tuning

This feature automatically tunes the initialization parameter `UNDO_RETENTION`, that is used to control retention of undo information in rollback segments. Automatic Undo Retention Tuning enables the database server to make a best effort at adjusting to changes in undo requirements of user queries, with changes in system activity given the space previously allocated to the undo tablespace. It thus relieves the administrators from the task of constantly tuning the parameter `UNDO_RETENTION`.

Segment Advisor

The Segment Advisor, based on the level of space fragmentation within an object, gives advice on whether an object is a good candidate for the new online shrink operation. The Advisor also reports on the historical growth trend of segments. This information is very useful or capacity planning purposes.

New Segment Resource Estimation

This feature facilitates space management by giving an estimate of the size of a table or index before it is created, so that an administrator can ensure that sufficient space is available.

Online Segment Shrink

This feature shrinks segments online and in-place (tables, indexes, and materialized views) that have free space in them, thereby enhancing efficiency of space utilization.

Proactive Tablespace Management

The Tablespace Advisor gathers and reports historical data about each tablespace's disk space usage and notifies administrators when tablespaces run low on available space.

Undo Advisor

This feature advises the database administrator in sizing the undo tablespace and in setting `UNDO_RETENTION` appropriately. This helps avoid a 'snapshot too old' error, and, for those using the Flashback features, it helps set the undo tablespace size correctly.

Backup and Recovery Manageability

This release simplifies the implementation and on-going maintenance of your data protection strategy. Recovery Manager (RMAN) now provides backup compression to keep backups on disk when space is a premium and enhanced RMAN scripting capabilities that allow you to modify them using your favorite text editor. Oracle Enterprise Manager has added new backup management functionality to provide easy configuration and monitoring of your backups.

Enhanced RMAN Reporting

Recovery Manager provides enhanced reporting. It can now list the backups required to restore a database and the backups required to restore archive logs.

Backup Compression

If disk space is an issue, or your media-management software does not support compression, Recovery Manager (RMAN) provides the ability to compress RMAN backup sets.

Bounded Backup Window

You can use the new DURATION option to give an RMAN BACKUP command a limit on the amount of time that the backup is allowed to take. You can also control I/O by specifying a window of time in which a backup can run.

RMAN Scripts Manageability

Recovery Manager (RMAN) now enables database administrators to easily view all scripts that are stored in the recovery catalog. Using a text editor, an administrator can create and test an RMAN script before saving it into a recovery catalog. If the script needs to be edited, the administrator can write the script to a file to make necessary changes.

Recovery Area Full Alert

Administrators will now be alerted when an out-of-space condition is imminent in the Recovery Area, potentially preventing new backup and archiving operations.

Instance Tuning

Automated instance tuning capabilities have been provided that greatly simplify an administrator's life. The built-in resource management available in the Oracle Database has been expanded to include CPU usage quotas, allowing administrators to easily set best operating procedures for all types of resource allocation. This in turn makes it very easy to provide predictable response times for key business operations. In addition, new methods of identifying resource consumption groups allow existing applications to take advantage of these capabilities without application changes. This release also supports automated checkpoint tuning, which takes advantage of periods of low I/O usage to advance checkpoints and therefore improve availability.

User-Initiated Buffer Cache Flushing

This feature enables easier testing and performance diagnosis.

Database Resource Manager - Adaptive Consumer Group Mapping

This feature makes it easier to use Database Resource Manager without requiring any application changes.

New Performance Overview Charts in Oracle Enterprise Manager

The enhanced Oracle Enterprise Manager HTML interface provides a central point of access to all database performance-related statistics and facilitates complete monitoring and diagnostics.

Improved SQL Reporting Using Oracle Enterprise Manager

The new Oracle Enterprise Manager HTML interface for analyzing SQL, including Top SQL, helps detect bad SQL and facilitates easy tuning.

Automatic Database Diagnostic Monitor

This feature enables the database to self-analyze its performance. The database can identify potential bottlenecks and fix them automatically or recommend a solution to an administrator. This capability is built inside the database kernel and hence does not require any external tools.

Redo Logfile Sizing Advisor

This feature recommends the optimal size of redo log files in order to avoid excessive disk I/Os due to frequent checkpointing.

Automatic Checkpoint Tuning

The Oracle Database can now self-tune checkpointing to achieve good recovery times with low impact on normal throughput. You no longer have to set any checkpoint-related parameters.

Automatic Shared Memory Tuning

Automatic Shared Memory Tuning automates the configuration of System Global Area (SGA) memory-related parameters (buffer cache, shared pool) through self-tuning algorithms. It simplifies database configuration, ensures most efficient utilization of available memory and improves performance.

Transaction Rollback and Recovery Monitoring

This feature enables you estimate how long it will take to roll back a transaction. You can also monitor the progress of transactions being recovered and estimate the average speed of transaction recovery.

See Also:

PL/SQL Packages and Types Reference for information on the Database Resource Manager

Oracle Database Administrator's Guide for information on the Database Resource Manager

Oracle Database Concepts for more information about the Database Resource Manager, advisors, and other manageability features

Application Tuning

New tools have been introduced in this release to minimize the manual SQL tuning efforts. These tools provide advice to administrators on new indexes or materialized views that can be created to optimize SQL performance and suggest changes to existing indexes and materialized views in order to make them more effective.

SQL Tuning Advisor

SQL Tuning Advisor is a new server tool that eliminates manual tuning of SQL statements as an input and gives advice in the form of precise SQL actions for tuning the SQL along with their expected performance benefit. This feature also introduces a new capability for SQL tuning of packaged applications without requiring any modification to the application code.

SQLAccess Advisor

The SQLAccess Advisor is an expert system that identifies and helps resolve performance problems relating to the execution of SQL statements by recommending which indexes or materialized views to create, drop, or retain.

Materialized View Tuning API

The new `TUNE_MVIEW` API advises what changes you need to make to a materialized view to make it fast refreshable and eligible for advanced query rewrite techniques.

Automatic Optimizer Statistics Collection

This feature automates the collection of optimizer statistics for objects. Oracle gathers statistics on all database objects automatically that have stale or missing statistics and maintains those statistics in a regularly-scheduled maintenance job. Automated statistics collection eliminates many of the manual tasks associated with

managing the query optimizer, and significantly reduces the chances of getting poor execution plans because of out of date statistics.

Database Overall

The utilities that support the database have also been improved in this release, to take advantage of the new capabilities, as well as address some common requirements requested by customers.

Utilities

SQL*Plus and *i*SQL*Plus have been improved in this release to take advantage of new capabilities, and provide improved functionality. In addition, newer versions of Import and Export, based on Oracle Data Pump technology, provide very fast and powerful loading and unloading of data.

SQL*Plus - glogin.sql Changes

`SET PAGESIZE 14` and `SET SQLPLUSCOMPATIBILITY 8.1.7` commands have been removed from the `gloglin.sql` file. The default for `SET PAGESIZE` is 14. The default for `SET SQLPLUSCOMPATIBILITY` is 10.1.

SQL*Plus - Show Recycle Bin

`SHOW RECYCLEBIN [original_name]` enables you to view objects that are available for purging or reverting using the `PURGE` and `FLASHBACK BEFORE DROP` commands.

SQL*Plus - Run glogin.sql and login.sql After Each CONNECT

The `glogin.sql` and `login.sql` scripts are now read after each `CONNECT` command, rather than just when SQL*Plus is started, affording easy per-connection configuration.

SQL*Plus - DBMS_OUTPUT After SELECT

SQL*Plus now displays any `DBMS_OUTPUT` information generated as the result of a procedure or trigger in a `SELECT` statement. This feature improves PL/SQL debugging and reporting capabilities.

SQL*Plus - COMPATIBILITY Command Line Option

SQL*Plus now supports the `-C` command line option. You can use this to set the initial value of the `SQLPLUSCOMPATIBILITY` parameter, giving you more control over SQL*Plus script behavior.

SQL*Plus - SET SQLPROMPT Runtime Variable Substitution

This feature lets you identify the schema and server you are connected to. Substitution variables are now dereferenced in the user prompt. New predefined variables exist for the username, connection privilege, and current date.

SQL*Plus - SPOOL CREATE, REPLACE, and APPEND Options

The `SPOOL` command has been enhanced to include `CREATE`, `REPLACE`, and `APPEND` options. These provide better control over output file creation and bring `SPOOL` syntax into line with the `SAVE` and `STORE` commands.

Data Pump Data Load/Unload

The Data Pump Export and Import utilities provide very high-speed bulk movement of data and metadata from one database to another. These utilities offer several significant advantages over the original Export and Import utilities, including: the ability to completely restart export and import jobs; the ability to detach and re-attach to long-running jobs; the ability to estimate how much disk space an export job would consume; support for export and import operations over the network; and support for fine-grained object selection, based upon object and object types.

iSQL*Plus - Output Over Multiple Pages

Output can now be displayed over multiple pages, or, as previously, on a single page.

iSQL*Plus Input Prompting

User input is now permitted during the running of a SQL script. This improved compatibility for SQL*Plus scripts provides more flexibility for development and deployment of scripts.

Performance and Scalability

As with previous releases, this release also focuses on continuing to deliver the high levels of performance and scalability the Oracle Database is renowned for.

Overall

In this release, many database components have been optimized, resulting in higher performance for transaction processing as well as data warehousing. Performance tuning has also been made simpler and more automated. There are also platform-specific enhancements for the Intel Itanium2 architecture, 64-bit Windows and Infiniband.

Index-Organized Table (IOT) Partitioning Enhancements

IOTs add bitmap index support for (range / hash / list) partitioned IOTs to enhance performance of operations on partitioned IOTs.

Eliminate Duplicated Columns in Index on an Index-Organized Table

IOTs save space when a secondary index defined on an IOT includes columns that are also in the IOT primary key. This improves performance as well because fewer blocks are accessed.

Configurable TCP/IP Send and Receive Buffer Size

This feature adds interfaces to enable a user to specify Send and Receive buffer sizes associated with a transport connection, thus improving network performance, especially in bulk data transfer such as replication, data warehousing, data mining, and so on.

Single-Set Aggregates in DML Returning Clause

This feature enables the use of single-set aggregation functions in the `RETURNING` clause of DML statements. This capability can result in significant performance gains in transactions that process many rows of the same table, a behavior often observed in batch processes.

High-Speed Infiniband Network Support

Oracle protocol support now includes support for the industry-standard Sockets Direct Protocol (SDP) for Infiniband high-speed networks. The SDP protocol is a high-speed communication protocol that speeds up performance of client / server and server / server connections. By using SDP, applications place most of the messaging burden upon the network interface card, freeing the CPU for other tasks.

Sorted Hash Clusters

Sorted hash clusters are new data structures that enable faster retrieval of data for applications where data is consumed in the order in which it was inserted.

Windows Fiber Support

This feature improves performance on Microsoft Windows platforms by enabling Oracle to use Windows fibers.

Optimized 64-bit Database

The Oracle Database has been enhanced for increased performance on 64-bit Windows. This performance improvement is transparent, requiring no database parameter changes.

SQL Tuning Advisor

SQL Tuning Advisor is a new server tool that eliminates manual tuning of SQL statements as an input and gives advice in the form of precise SQL actions for tuning the SQL along with their expected performance benefit. This feature also introduces a new capability for SQL tuning of packaged applications without requiring any modification to the application code.

SQLAccess Advisor

The SQLAccess Advisor is an expert system that identifies and helps resolve performance problems relating to the execution of SQL statements by recommending which indexes or materialized views to create, drop, or retain.

Materialized View Tuning API

The new `TUNE_MVIEW` API advises what changes you need to make to a materialized view to make it fast refreshable and eligible for advanced query rewrite techniques.

Automatic Optimizer Statistics Collection

This feature automates the collection of optimizer statistics for objects. Oracle gathers statistics on all database objects automatically that have stale or missing statistics and maintains those statistics in a regularly-scheduled maintenance job. Automated statistics collection eliminates many of the manual tasks associated with managing the query optimizer, and significantly reduces the chances of getting poor execution plans because of out of date statistics.

Clustering

This release continues to address key requirements in the area of providing reliable, scalable processing power on clusters of machines.

Oracle Real Application Clusters Enhancements

This version of Oracle Real Application Clusters introduces a new service framework that allows administrators to configure, manage, and monitor application workloads as a service, deployed across a number of nodes, in a large-scale cluster deployment. This new framework allows administrators not only to monitor and manage performance levels for a given service but also to manage how to provide these services continuously.

Integrated Clusterware Management

This release offers a complete clusterware management solution as an integral component of Oracle Real Application Clusters, available on all platforms Oracle Database runs on. This clusterware functionality includes mechanisms for cluster connectivity, messaging and locking, cluster control and recovery, and a services provisioning framework. No third party clusterware management software is required, Oracle will, however, continue to support select third party clusterware products on specified platforms.

Automatic Workload Management

With this release, different application workloads can be defined as named services so that they can be individually managed and controlled. Database administrators can then control which processing resources are allocated to each service during both normal operations and in response to failures. CPU resource allocations and resource consumption controls can also be managed for named services using Resource Manager. Performance metrics are also tracked by service and thresholds can be set to automatically generate alerts should these thresholds be crossed. Oracle tools and facilities such as Job Scheduler, Parallel Query, and Oracle Streams Advanced Queuing also use services to manage their workloads.

Single System Image Management

Oracle Enterprise Manager has been significantly enhanced to enable true single system image management of cluster database deployments. Oracle Enterprise Manager's Cluster Database Page provides a single view of system status across multiple nodes. It also enables direct drill down to individual instances as and when needed.

Fast Connection Failover

In this release, Oracle Real Application Clusters enables fast, coordinated recovery between the database and application mid-tier components and products. Oracle Real Application Clusters 10*g* initially supports JDBC Implicit Connection Cache and Oracle Application Server 10*g* with this capability and will support additional mid-tier components and products in the future.

Performance Improvements

Several new optimizations in Oracle Real Application Clusters provide performance improvements for many applications. These include optimizations that reduce message traffic, memory usage and the consumption of other resources. In addition, dynamic file and cache affinity will aid performance when workloads are shifted between instances.

Zero Downtime Patching

In this release, Oracle Real Application Clusters supports the application of patches to the nodes of a system in a rolling fashion, with no downtime. Patches can be applied one node at a time while the other nodes in the RAC system are up and operational. Patches will be labeled as being qualified for installation as rolling upgradeable, or not, depending on the changes being made by the patch.

Cluster Verification and Improved Diagnostic Tools

This release introduces a new cluster configuration verification tool and improvements in the diagnostic tools first introduced in the Oracle9*i* Database release. Together these tools help users both avoid problems and resolve problems more quickly should they occur.

Grid Computing

Oracle Database 10*g* is the database for Grid Computing. A great many of the new features in this release enable users to reduce costs, make more efficient use of resources, and more quickly align their resources to their changing business needs. This release includes features to enable virtualization and dynamic provisioning of resources, and to efficiently manage workloads in a Grid environment.

Resource Virtualization and Provisioning

This release provides features to virtualize and dynamically provision on demand resources such as CPU, storage, and data. Resource virtualization and provisioning

is the key to improving resource utilization and enabling the realignment of resources as business needs change.

Resonance

Reduces costs of operating databases by improving resource utilization and eliminating the need for manually monitoring and allocating resources. Automatically re-provisions resources to meet changing business needs.

Transparent Session Migration

Automatically migrates sessions, including all state, from one instance to another, with no impact on end-users. Rebalances sessions across instances for load balancing, and in preparation for deprovisioning an instance.

Streams - Moving an Operating System File

In this release, the database can move any operating system file. This lets developers of applications move and copy data that is not stored in the database, helping applications keep data inside and outside the database consistent, and providing a mechanism to provision data external to the database in a Grid environment.

Streams - Simple Instantiation of Replica Tablespaces and Databases

You can now create and maintain a replica of one or more tablespaces with a single command, simplifying the provisioning of a database in a Grid environment. Streams will automatically copy the data, ensure the replica is up to date, and optionally set up bi-directional replication.

Integrated Clusterware Management

In this release, Oracle Real Application Clusters offers a complete clusterware management solution as an integral component of Oracle Real Application Clusters, available on all platforms Oracle Database runs on. This clusterware functionality includes mechanisms for cluster connectivity, messaging and locking, cluster control and recovery, and a services provisioning framework. No third party clusterware management software is required, Oracle will, however, continue to support select third party clusterware products on specified platforms.

High-Speed Infiniband Network Support

Oracle protocol support now includes support for the industry-standard Sockets Direct Protocol (SDP) for Infiniband high-speed networks. The SDP protocol is a

high-speed communication protocol that speeds up performance of client/server and server/server connections. By using SDP, applications place most of the messaging burden upon the network interface card, freeing the CPU for other tasks.

Fast Connection Failover

In this release, Oracle Real Application Clusters enables fast, coordinated recovery between the database and application mid-tier components and products. Oracle Real Application Clusters initially supports JDBC Implicit Connection Cache and Oracle Application Server 10*g* with this capability and will support additional mid-tier components and products in the future.

Automatic Storage Management

Automatic Storage Management automates and simplifies the optimal layout of datafiles, control files, and log files. Database files are automatically distributed across all available disks, and database storage is rebalanced whenever the storage configuration changes. This feature also provides redundancy through the mirroring of database files.

Cross-Platform Transportable Tablespaces

The transportable tablespace feature now enables tablespaces to be transported across different platforms.

Data Pump Export and Import Utilities

The Data Pump Export and Import utilities provide very high-speed bulk movement of data and metadata from database to another. These utilities offer several significant advantages over the original Export and Import utilities, including: the ability to completely restart export and import jobs; the ability to detach from and reattach to long-running jobs; the ability to estimate how much space an export job would consume; support for export and import operations over the network; and support for fine-grained object selection, based upon objects and objects types.

Streams - Instantiation Through RMAN

Streams replicas can now be instantiated by way of RMAN, providing faster instantiation of a Streams replica at a remote site or on another system in a Grid environment.

Streams - Instantiation Through Transportable Tablespaces

Streams replicas can now be instantiated through transportable tablespaces, providing faster instantiation of a Streams replica at a remote site or on another system in a Grid environment.

Streams - Instantiation Through Data Pump

Streams replicas can now be instantiated through data pump, providing faster instantiation of a Streams replica at a remote site or on another system in a Grid environment.

Workload Management

New features in this release automate workload management to improve utilization and efficiency. Workloads can be balanced across resources, and new scheduling capabilities distribute and manage workload across time.

Automatic Workload Management

With this release, different application workloads can be defined as named services so that they can be individually managed and controlled. DBAs can then control which processing resources are allocated to each service during both normal operations and in response to failures. CPU resource allocations and resource consumption controls can also be managed for named services using Resource Manager. Performance metrics are also tracked by service and thresholds can be set to automatically generate alerts should these thresholds be crossed. Oracle tools and facilities such as Job Scheduler, Parallel Query, and Oracle Streams Advanced Queuing also use services to manage their workloads.

Database Resource Manager - Adaptive Consumer Group Mapping

This feature makes it easier to use Database Resource Manager without requiring any application changes.

Scheduler - Job Processing

The Scheduler enables job processing in a way that models your business requirements. It lets limited computing resources be allocated appropriately among competing jobs, thus aligning job processing with your business needs. In this release, you can group jobs that share common characteristics and behavior into larger entities called job classes. You can prioritize among the classes by controlling the resources allocated to each class. This ensures that your critical jobs have priority and have enough resources to complete. You can also prioritize jobs within

a job class. The Scheduler lets you change the prioritization based on a schedule. Because the definition of a critical job can change across time, the Scheduler lets you change the prioritization among your jobs over time.

Scheduler - Usability and Manageability

The Scheduler has been designed to be easy to use. It supports user-defined object names, has easy time specification syntax, and has defaults for all the API arguments. Schedules are validated, and both jobs and schedules can be saved in a library for reuse.

Scheduler - Support for Jobs in a Clustered Environment

The Scheduler fully supports execution of jobs in a clustered or Grid environment. To balance the load on your system and for better performance, you can also specify the service where you want a job to run.

Information Integration

This release provides a great many new features that will help you integrate your information within a department, enterprise, or within a Grid environment. New features provided improved performance in heterogeneous environments. There are also numerous improvements to Oracle Streams, many which help improve both performance and manageability. Lastly, a new database scheduler provides a mechanism to automate integration activities, and improve resource utilization over time.

Improved Information Integration Performance

This release provides improved performance for integrating information in heterogeneous environments. A new feature enables reference to remote stored procedures, enabling sophisticated processing in the most efficient manner.

Transparent Gateway - Remote Stored Functions in SELECT Statements

In SELECT statements, you can now refer to remote functions stored in a non-Oracle database. This allows for more efficient local processing of non-Oracle data, reducing the amount of data that must be processed by the gateway.

See *Oracle Database Heterogeneous Connectivity Administrator's Guide* for more details.

Oracle Streams

This release provides functionality, performance, and management improvements for Oracle Streams. New functionality includes downstream capture, which offloads the capture and mining operations from the production database. Oracle Streams has also been expanded to support LONG, LONG RAW, and NCLOB datatypes, index-organized tables, and row subsetting during capture and propagation, expanding the range of applications Streams can be used with. Support for negative rules makes it easier to specify complex subscriptions. Other manageability improvements include better monitoring, additional views, improved diagnostics, and better error handling and cleanup operations. Performance has also been tuned to improve most operations. Oracle Streams also offers two new options for instantiation, using RMAN and transportable tablespaces. Finally, this release extends support for the Messaging Gateway (MGW).

JMS 1.3-Compliant JDBC - Thin Driver Support

This feature provides a J2EE-compatible JMS provider regardless of your choice of Oracle JDBC drivers. You can now use the JDBC thin driver.

OJMS Queue and Topic Unification Support

OJMS (AQ/JMS) now supports Domain Unification as it is specified in the JMS 1.1 Specification. Within one transacted session, both the point-to-point and pub-sub domains are supported. As a result, OJMS applications can now send a message to a queue and receive a message from a topic within one transacted session. OJMS also now supports Streams AQ batch enqueue and dequeue.

Oracle Messaging Gateway - MQSeries JMS Interface

This feature supports integration of Advanced Queuing with the MQSeries JMS interface.

Oracle Messaging Gateway - OJMS (AQ/JMS) Interface

This feature supports integration with the Advanced Queuing OJMS interface.

Oracle Messaging Gateway - Tibco/Rendezvous Java Interface

This feature supports integration of Advanced Queuing with the Tibco/Rv messaging system through its Java interface.

Streams - Batch Enqueue/Dequeue

You can now enqueue or dequeue an array of messages to a single queue. The messages all share the same enqueue or dequeue options, but each message in the array can have different message properties.

Advanced Replication to Streams Migration Tool

This tool enables you to migrate easily from Advanced Replication to Streams.

Streams - Support Delete Cascade Operations

Streams now replicates tables with ON DELETE CASCADE constraints.

Streams - Enqueue Handler

It is no longer necessary to write an apply handler that implicitly dequeues a logical change record (LCR) from the buffered queue and then explicitly enqueues the LCR into the persistent queue. A new name-value pair for the action context of an apply rule lets you specify the queue into which to move an LCR. You can optionally apply the LCR as it is reenqueued.

Streams - Access to Client Information During Processing

You can use new functions to get name and type information about the Streams client processing an LCR, useful for writing more intelligent rules, transformations, apply handlers and error handlers.

Streams - Descending and Function-Based Index Support

You can now use Streams with descending and function-based indexes.

Streams - Precommit Handlers

You can use a new type of apply handler called a precommit handler to record information about commits processed by an apply process.

Streams - Negative Rules

You can now write rules which, if true, will cause the subscription to be false, simplifying the creation of sophisticated rules.

Streams - Rules Engine - Easy Rules Engine Transformation Management

A new function in the DBMS_STREAMS_ADM package makes it easier to set the transformation function name for a rule.

Streams - Row Subsetting During Capture and Propagation

Streams can now automatically subset rows between subset databases without having to send unnecessary data to destination databases.

Streams - Additional Streams Datatype and IOT Support

You can now use Oracle Streams to capture and apply changes to index-organized tables (IOTs) and to LONG and NCLOB datatypes.

Streams - Access to Additional LCR Attributes

You now have access to additional LCR attributes, row_id, serial#, session#, thread#, tx_name, and username, for use in your custom procedures, transformations, and rule conditions.

Streams - Instantiation Through Data Pump

Streams replicas can now be instantiated through data pump, providing faster instantiation of a Streams replica at a remote site or on another system in a Grid environment.

Streams - Clean Up Rules Sets

Some Streams packaged procedures have been enhanced to let you remove unnecessary rule sets, preventing performance degradation and confusion.

Streams - Buffer Queue Monitoring

Streams provides new views to monitor statistics on a Streams buffered queue.

Streams - Message Notification

You can now easily configure email, http, and PL/SQL message notifications for a Streams messaging client. Notifications eliminate the need for clients to block or poll for new messages.

Streams - Allocate Memory from Streams Pool

You can now allocate memory from the new Streams pool, giving database administrators more control over memory usage by Streams.

Streams - High-Level API for Message Queuing

Streams now includes a simplified API for enqueue and dequeue operations into and from an `AnyData` queue.

Streams - Easier Transformation Monitoring

A new data dictionary view makes it easier to monitor transformation functions.

Streams - Improved Capture and Apply Error Handling

Streams now writes warnings during some error conditions instead of stopping the capture and apply.

Streams - Instantiation Through RMAN

Streams replicas can now be instantiated by way of RMAN, providing faster instantiation of a Streams replica at a remote site or on another system in a Grid environment.

Streams - Expose Commit Order of Changes

You can now implement record keeping and order transactions when LCRs are applied by external applications rather than by the Streams apply process.

Streams - Clean Up LogMiner Information

A new `DBMS_CAPTURE` procedure enables you to remove old LogMiner information. This cleanup capability can prevent performance degradation and improve manageability by eliminating clutter.

Streams - Set Instantiation SCN at Apply Site

Streams packaged procedures have been enhanced to set the instantiation system change number (SCN) for all table objects owned by a schema, eliminating the need to set the instantiation SCN for each object individually.

Streams - Views to Monitor Instantiation Data

New data dictionary views make it easier for you to ensure that Streams is operating and to discover and resolve any issues that may arise.

Streams - New APIs to Grant Privileges

You can now grant privileges through an API, which makes it easier in a Streams environment to grant only those privileges that are necessary.

Improved Streams Oracle Real Application Clusters Support

When used in an Oracle Real Application Clusters environment, Streams now supports hot mining and can automatically restart processes after failover.

Streams - Downstream Capture

Streams can now perform capture from log files on a system other than the source database. This can reduce the overhead on the source system and provide better disaster protection. In a data warehousing environment, it eliminates the need to access the production OLTP system.

Streams - Instantiation Through Transportable Tablespaces

Streams replicas can now be instantiated through transportable tablespaces, providing faster instantiation of a Streams replica at a remote site or on another system in a Grid environment.

OCCI Support for Streams AnyData Queues

C++ developers can now use the OCCI API to enqueue and dequeue to and from Streams and `AnyData` queues.

Streams - Enhanced Capture and Apply Performance

Streams change capture and apply processes have been improved to handle greater workloads.

Streams - Performance Enhancements for AnyData

Streams `AnyData` queue operations, rules evaluation, and propagation have been improved to handle greater workloads.

Rules Engine - Rules Enhancements

This feature reduces the time you need to spend creating and managing rules used by Streams and other applications. Additional views and statistics are provided for tuning purposes.

See Also:

Oracle Streams Replication Administrator's Guide for information on migration from Advanced Replication to Streams, supplemental logging, and the following instantiation enhancements: RMAN, SCN at the apply site, and transportable tablespaces

Oracle Streams Concepts and Administration for information on the enqueue handler, rules enhancements, negative rules, the rules engine, row migration, monitoring enhancements for transformations, downstream capture, and new APIs for granting privileges

PL/SQL Packages and Types Reference for information on message queuing and LogMiner information

The Scheduler

This release includes a new database scheduler to provide enterprise scheduling functionality. The Scheduler gives administrators the ability to schedule a job to run at a particular date and time. It also provides the ability to create libraries of the Scheduler objects, thus allowing existing objects to be shared by other users. It also enables scarce computing resources to be allocated appropriately among competing jobs, thus aligning job processing with the service-level needs of the business. Jobs that share common characteristics and behavior can be grouped into larger entities, called job classes, which can be prioritized by controlling the system resources allocated to each. For finer control, the prioritization among the job classes can also be based on a schedule.

Scheduler - Core Scheduler Features

The Scheduler integrates many job scheduling capabilities, such as time-based job execution, in one tool. It supports PL/SQL stored procedures and anonymous blocks, C functions, and Java stored procedures by way of callouts and operating system scripts. It also supports distributed database scheduling.

Scheduler - Usability and Manageability

The Scheduler has been designed to be easy to use. It supports user-defined object names, has easy time specification syntax, and has defaults for all the API arguments. Schedules are validated, and both jobs and schedules can be saved in a library for reuse.

Scheduler - Monitoring

A job undergoes multiple states from its creation to its completion. All Scheduler activity is logged, and you can easily track information such as the status of the job and the time to completion of the job. This information is stored in views and can be easily queried using Oracle Enterprise Manager or a SQL query.

Scheduler - Job Processing

The Scheduler enables job processing in a way that models your business requirements. It lets limited computing resources be allocated appropriately among competing jobs, thus aligning job processing with your business needs. In Oracle Database, you can group jobs that share common characteristics and behavior into larger entities called job classes. You can prioritize among the classes by controlling the resources allocated to each class. This ensures that your critical jobs have priority and have enough resources to complete. You can also prioritize jobs within a job class. The Scheduler lets you change the prioritization based on a schedule. Because the definition of a critical job can change across time, the Scheduler lets you change the prioritization among your jobs over time.

Scheduler - Recovery

The Scheduler recovers back to a transactionally consistent state from a system or slave process crash. Jobs and windows that would have started had the system not crashed will be started automatically, so no manual intervention is required.

Scheduler - Support for Jobs in a Clustered Environment

The Scheduler fully supports execution of jobs in a clustered or Grid environment. To balance the load on your system and for better performance, you can also specify the service where you want a job to run.

Scheduler - ILMS Support

This feature improves the performance of jobs involving PL/SQL stored procedures, Java stored procedures, and C functions by skipping the parse phase.

Scheduler - Increased Throughput

Improved job coordinator algorithms have increased the number of Scheduler jobs that can be executed per hour.

Scheduler - Dynamic Slave Pool

The slave pool grows and shrinks dynamically depending on demand, thus reducing the resource use and increasing the number of jobs that can be executed.

Scheduler - Redo Minimization

This feature reduces the amount of Scheduler information that is written to disk. It is very beneficial for polling type jobs, which would otherwise require a huge amount of information to be written to disk.

Availability

Availability of data is a major requirement for organizations wishing to deliver world-class data management capabilities. With this release, Oracle extends the ability of the database to deal with any type of human error, and also provides support for reducing the time it takes to implement database and application upgrades.

Backup and Recovery

In line with the management improvements in other areas, this release also greatly simplifies the management of backup and recovery in an Oracle environment. New in this release is a disk-based recovery area that can be used as an online disk cache for backup and recovery operations for one or more Oracle databases. Backups themselves can be scheduled to happen automatically, and automated backup tuning, in conjunction with fast incremental backups, ensures that backups complete within the window provided for these operations. In recovery scenarios, the relevant database is also able to identify the correct backup to use for restore operations, preventing the administrator from mistakenly using the wrong backup. A new offline synthetic recovery is also supported, ensuring that an up-to-date image is always available for fast media restore and recovery.

Flash Backup and Recovery

This release supports automated, disk-based backup and recovery. The benefits include simplified and unified storage location for backups, archive logs, and any other files needed for Oracle recovery; automatic deletion of the files after they have been successfully backed up by the Recovery Manager (RMAN); the equivalent of a disk cache for tape, which reduces the time needed to restore a file from tape; and reduced risk of an out-of-space condition on disk, by deleting files that are no longer required for database recovery.

Drop Database

The new DROP DATABASE command deletes all database files, all online logs, control files, and server parameter files (spfiles).

RMAN Database Deregistration

The new UNREGISTER DATABASE RMAN command removes all metadata for one database from the recovery catalog.

Backup and Restore of Standby Control File

This feature enables you to quickly restore the standby control file in case of media recovery, providing Oracle-managed files and automatic storage management support for standby databases.

Automatic TSPITR

This feature automatically creates the auxiliary instance needed to perform tablespace point-in-time recovery (TSPITR) and incorporate the RMAN TSPITR operations.

Simplified Recovery Manager Cataloging of Backup Files

You can now catalog RMAN proprietary backup metadata into a backup repository. If a backup is overwritten in the control file or a backup file is moved to a new location on disk, then you can easily uncatalog the backup metadata from the repository.

Automatic Channel Failover for Backup and Restore

Recovery Manager (RMAN) now automatically retries a failed backup or restore operation, reducing the risk of leaving you with no backup of the database because of an error.

Automatic File Creation During Recovery

This feature enhances RMAN recovery by automatically creating and recovering datafiles that have never been backed up.

Simplified Backups to Disk

Image backups provide fast recovery by being readily usable. The Recovery Manager (RMAN) BACKUP command has been enhanced to perform image copy backups at the database, tablespace, and datafile level.

Proxy Copy Backup of Archivelogs

You can now back up archive logs by way of the Recovery Manager (RMAN) Proxy Copy.

Incrementally Updated Backups

You can now apply a Recovery Manager (RMAN) incremental backup to a datafile image backup. This results in reduced recovery time, because fewer logs need to be applied, and reduced time to back up the database, because you do not always have to back up the whole database.

Simplified Recovery Through Resetlogs

You no longer have to back up your database following an incomplete recovery and `OPEN RESETLOGS` operations.

Restore Tolerates Corrupt or Missing Backups

When the latest backup is not available, Recovery Manager (RMAN) now automatically uses an older backup for restore operations.

Full Database Begin Backup Command

It is no longer necessary to issue a separate command to place each tablespace in hot backup mode. You can now use the `ALTER DATABASE` statement to place all tablespaces in backup mode. Also, the `BEGIN BACKUP` command now runs faster than before.

Change-Aware Incremental Backups

By using a new type of log file to track blocks that have changed in the database, Recovery Manager (RMAN) can avoid scanning the entire datafile during an incremental backup. Instead, the amount of data scanned is proportional to the amount of data changed.

See Also:

Oracle Database SQL Reference for information on the DROP DATABASE statement

Oracle Database Backup and Recovery Basics for information on automated disk-based backup and recovery

Oracle Database Administrator's Guide for information on dropping a database and on simplified recovery through resetlogs

Oracle Data Guard Concepts and Administration for information on automated disk-based backup and recovery and on simplified recovery of resetlogs

Oracle Database Concepts for more information about flash forward and an overview of backup and recovery features

Enhanced Oracle Data Guard Infrastructure

The Oracle Data Guard infrastructure introduced to support standby databases has been improved. Support for log mining of additional datatypes is provided, as well as hot log mining capabilities, allowing Oracle Data Guard to be used with a wider variety of existing legacy applications. A new real-time standby apply capability is provided, which ensures that the standby database is in close synchronization with the production system. Zero downtime instantiation and easier zero data loss operations are also supported for standby databases using SQL Apply. Management of an Oracle Data Guard environment has also been improved in this release, with support for more fine-grained supplemental logging at the database, schema, and table level, and improved monitoring capabilities. Performance and security of data transmission have also been improved.

See *Oracle Data Guard Concepts and Administration* for information on all Data Guard enhancements.

Data Guard Broker Support for RAC

You can now configure and support Real Application Clusters instances in a Data Guard configuration using Data Guard Broker interfaces, both the GUI and the command line.

Automatic LogMiner Configuration

For users who conduct LogMiner mining in the same database that generated redo logs, LogMiner can scan the control file itself and determine the redo logs that satisfy a requested time or SCN range. Through this feature, LogMiner by default adds redo logs from the mining database.

Log Miner Support for Index-Organized Tables

LogMiner support for index-organized tables increases transaction recovery options for DBAs using LogMiner and makes logical standby more complete.

LogMiner Support for More Types: LONG, Multibyte CLOB and NCLOB

LogMiner and SQL Apply now support multibyte CLOB and NCLOB data. SQL Apply now also supports LONG data. Support of additional datatypes means that you can now mine a greater variety of data.

Fine-Grained Supplemental Logging

This feature adds a new set of data definition language (DDL) statements that enable you to add or drop supplemental logging at the database, schema, and table levels.

Secured Redo Transmission

This feature increases the security of a Data Guard environment by preventing possible tampering of redo data as it is being transferred to the standby database.

Uniquely Named Databases with DB_UNIQUE_NAME

With `DB_UNIQUE_NAME`, it is now possible to dynamically add a standby database to a Data Guard configuration that contains a Real Applications Clusters primary database, when that primary database is operating in either the maximum protection or maximum availability protection modes, without shutting down the primary database. `DB_UNIQUE_NAME` also enhances the usability aspects of a Data Guard configuration.

Simplified Zero Data Loss for Data Guard SQL Apply

SQL Apply now supports standby redo logs (SRLs) that enable complete zero-data-loss support for logical standby databases.

Zero Downtime Instantiation for SQL Apply

Shutdown or quiesce of the primary database is no longer required when you create a logical standby database from a primary database.

Real Time Apply

The managed recovery process (MRP) can now recover redo from standby online redo logs as the logs are being filled, without requiring them to be archived at the standby database, resulting in faster recovery, switchover, and failover times.

Automating Recovery Through Open Resetlogs in Standby Databases

Data Guard supports the new in this release recovery through resetlogs feature that simplifies recovery with backups taken from an earlier incarnation so that it is as easy as recovering a backup from the same incarnation. Hence, database administrators no longer need to make new backups of a database after a RESETLOGS operation. Data Guard supports the recovery through resetlogs feature by automating operations on a standby database when an ALTER DATABASE OPEN RESETLOGS statement is being performed on a corresponding primary database. By providing this automation, Data Guard eliminates the possibility of user error in the event the standby database has not applied beyond the primary database OPEN RESETLOGS point-in-time.

See *Oracle Database Utilities* and *PL/SQL Packages and Types Reference* for information on LogMiner configuration and *Oracle Database Utilities* for information on LogMiner support for index-organized tables and for additional datatypes.

Recovering from Human Error with Flashback Technology

In this release, Oracle introduces expanded database Flashback capabilities. If a major error occurs, such as a batch job being run twice in succession, the database administrator can request a Flashback operation that quickly recovers the entire database to a previous point in time, eliminating the need to restore backups and do a point-in-time recovery. In addition to Flashback operations at the database level, it is also possible to flash back an entire table in this release of the Oracle Database. Similarly, a new capability allows the database to recover tables that have been inadvertently dropped by a user. The existing Oracle Flashback Query capabilities have also been improved.

Oracle Flashback Database

This feature introduces the FLASHBACK DATABASE statement in SQL. It let you quickly bring your database to a prior point in time by undoing all the changes that

have taken place since that time. This operation is fast, because you do not need to restore the backups. This in turn results in much less downtime following data corruption or human error.

Oracle Flashback Table

This feature introduces the FLASHBACK TABLE statement in SQL, which enables you to quickly recover a table to a point in time in the past without restoring a backup.

Oracle Flashback Version Query

Using undo data stored in the database, you can now view the changes to one or more rows along with all the metadata of the changes.

Oracle Flashback Drop

Oracle now provides a way to restore accidentally dropped tables.

Oracle Flashback Transaction Query

This feature introduces Oracle Flashback Transaction Query, which enables you to examine changes to the database at the transaction level. As a result, you can diagnose problems, perform analysis, and audit transactions.

See *Oracle Database SQL Reference* for information on the FLASHBACK DATABASE, FLASHBACK TABLE, and UNDROP statements and on flashback queries and *Oracle Database Backup and Recovery Advanced User's Guide* for information on FLASHBACK DATABASE.

Reduced Downtime for Application and Database Upgrades

This release further reduces the downtime associated with application and database upgrades by supporting rolling window upgrades of hardware, operating system, or database versions. This is achieved by using standby databases, with new support allowing an organization to switch between different versions of standby and production databases. Different patch releases of database software running in an Oracle Real Application Clusters environment are also supported. For database changes and application upgrades, the current online redefinition capabilities have been expanded to support one-step cloning of all the associated database objects. In addition, changes to the underlying objects no longer invalidate the PL/SQL packages built on these objects, allowing new redefinition operations to be performed without the need to recompile the corresponding stored procedures.

See *Oracle High Availability Architecture and Best Practices* for general information on application and database upgrades.

Enhanced Online Redefinition

This feature increases application availability by minimizing downtime, such as the downtime for application upgrades.

Rolling Upgrades with SQL Apply

In a future patchset release of this release, it will be possible to do a rolling upgrade using logical standby databases. The foundation for rolling upgrades is now implemented into the SQL apply technology so that the primary database incurs minimal downtime when you upgrade the Oracle Database software on each database in the Data Guard configuration. For example, using SQL apply and logical standby databases, you will be able to upgrade the oracle database software from patchset release 10.1.0.n to the next database 10.1.0.(n+1) patchset release. See the README file for the applicable patchset release.

Signature-Based Dependency Tracking Using Synonyms

When a PL/SQL subprogram or a view references a table using a synonym, then repointing the synonym to a table with the same signature and grants no longer invalidates its dependents. Thus, the need for time-consuming recompilation is avoided.

See Also:

PL/SQL Packages and Types Reference for information on online redefinition using the `RDBMS_REDEFINITION` package

Oracle Database Administrator's Guide for information on online redefinition enhancements

Oracle Data Guard Concepts and Administration for information on rolling upgrades with Data Guard SQL Apply

Security and Directory

Security improvements in this release focus on improved administration of large numbers of users across an integrated end-to-end environment, and support for key security technologies.

Improved Administration

Administration enhancements lower the costs associated with development, installation, deployment configuration, and management of security for realistic applications. Security procedures are simplified so that complexity is not the primary barrier to secure operational deployments.

DML Support in FGA

SQL support of fine-grained auditing (FGA) has been enhanced to support granular auditing of queries as well as `UPDATE`, `INSERT`, and `DELETE` operations.

Communication over SASL

This feature provides an alternate secure channel for communication between Oracle Internet Directory (OID) and the database or between two databases. If your enterprise deploys password-authenticated users, then you no longer have to invest in and administer public key infrastructure (PKI). Simple Authentication and Security Layer (SASL) communication provides an equally secure channel.

Unified User Model

A single sign on user defined in the directory can now use the features offered by Enterprise User Security. No additional administration tasks are required for provisioning and credential management. In addition, the administrative groups for Enterprise User Security can now designate an owner, which facilitates stronger security overall.

Easy Database Registration

This feature eliminates the need for the `RDBMS_SERVER_DN` parameter, making configuration of enterprise users easier.

Extended and Uniform Audit Trail

This feature improves security administration by providing a uniform audit trail for RDBMS auditing in the form of standard and fine-grained audit tables. Transactions and SQL information have been added to the audit tables to further improve accountability of all users.

Oracle Label Security Directory Integration

You can now manage Oracle Label Security policies and user label authorizations in a central OID/LDAP repository. This reduces administration costs and increases security by eliminating multiple management points.

Integration and Interoperability

Database integration with the Oracle Internet Directory has been enhanced, with support added for directory management of Oracle Label Security policies, and Certificate Revocation Lists (CRLs). The Oracle Internet Directory can also be used to store the passwords as SASL 'secrets'. In addition, this release provides full support for Kerberos based security frameworks, supporting both Kerberos based user authentication, and database-to-database communications based on Kerberos credentials. Fine grained auditing, has been expanded to apply not only to query operations to but DML operations as well, providing better accountability of all user operations.

Certificate Validation with Certificate Revocation Lists (CRLs)

Certificate validation is an important element of enabling public key infrastructure (PKI) in an enterprise. If you use SSL in an Oracle Database environment, you can now validate the certificates presented by servers and clients for authentication.

Centralized CRL Management

This feature reduces administrative costs by enabling certificate revocation lists (CRLs) to be centrally stored in the Oracle Internet Directory as an alternative to local CRL stores.

Centralized User Management for Kerberos Users

Enterprise users can be authenticated using their Kerberos credentials. This feature extends the support for Kerberos users from external users to global users and enables users to be mapped to an exclusive or shared schema.

Kerberos-Enabled Database Links

You can now enable Kerberos-based authentication across current-user and connected-user database links.

Operating System Credential Cache

This feature supports interoperability with commercially available key distribution center (KDC) vendors. For example, customers on Microsoft Windows platforms (including Windows 2000 and Windows XP) using Kerberos as the security infrastructure can now integrate with an Oracle environment with the same Kerberos credentials, achieving secure single sign on.

Access to Single Sign On Wallet

Java applications, including those running within OC4J, can now access the single sign on wallet through the key store interface without user input or programmatically provided password.

Key Store Interface

J2EE applications that must run over SSL can now use standard key store interfaces provided in the Java Development Kit to open Oracle Wallets.

Database Authentication with Standard LDAP Password Verifiers

This feature lets an application use Internet Engineering Task Force (IETF) standard Digest-MD5 verifiers for authentication of users and of the application itself, providing easier integration of users into the Oracle environment.

Single Station Administration for Password Authentication to Oracle Database

Oracle Database users and Oracle *i*AS users managed in the directory now use the same attribute in the directory for authentication. This feature provides the same level of security for safeguarding passwords (verifiers) as is available for public key infrastructure (PKI) credentials.

Smart Card Support for X509v3 Certificates

Oracle SSL users can now store their private key material on smart cards. This feature provides complete machine independence for performing SSL-based transactions.

Transport Layer Security (TLS) Support

Oracle Database and Oracle *i*AS now support transport layer security (TLS), which reduces network activity, improves caching, and provides a framework to extend public key and bulk encryption algorithms.

SSL Session Renegotiation

A client or server can now renegotiate the security parameters of an existing SSL session. Because SSL handshake is highly CPU intensive, this feature improves performance of SSL connection requests.

4096-Bit Key Size Support

Users and servers can now be provisioned with certificates with up to 4096-bit keys as well as the ability to honor the keys at run time.

Performance Improvements

This feature incorporates the performance enhancements RSA Data Security, Inc., has made for SSL Crypto, which will improve all Oracle products that are configured to run over SSL.

Security and Privacy

Virtual Private Database (VPD) security policies can be defined to trigger on relevant column access, providing both better accountability and more fine-grained data security. VPD also introduces static security policies for security rules that are always enforced, that is, not based on changing criteria, such as time of access. This provides a performance advantage in large-scale hosted environments. In addition, VPD support has been enhanced to allow parallel execution of application contexts, improving performance and scalability in data warehouses where parallel queries are routinely used.

VPD Support for Parallel Query

Virtual Private Database now supports parallel query, resulting in performance and scalability improvements.

VPD Static and Dynamic Policies

Virtual Private Database now lets you distinguish between static policies, which are suitable for hosting environments that always need to enforce an unchanging policy, and dynamic policies, which are suited for time-dependent enforcement, such as time of day, where rows returned must vary at a particular time.

Column-Level VPD

Virtual Private Database (VPD) is now more fine grained. You are now able to enforce VPD rewrite when a query references a particular column.

Business Intelligence

This release continues to focus on key requirements in the business intelligence and data warehousing areas. In addition, it introduces new capabilities to better support the needs of bioinformatics and life sciences customers.

Bioinformatics

This release contains specific improvements for bioinformatics customers. These include the native support of new DOUBLE and FLOAT datatypes in the database, resulting in improved performance for applications that require large amounts of numeric computations. New built-in statistical functions are also provided to support common statistical analysis. For genetic data, text mining capabilities are provided that support document clustering and classification using Support-Vector Machine algorithms. In addition, new BLAST similarity searches provide specific support for analysis of genetic data.

BLAST Queries

Oracle Data Mining now supports specialized nucleotide and amino acid sequence matching and annotation algorithms. A version of BLAST, like NCBI BLAST 2.0, exists in the database using table functions. As the algorithms are implemented as table functions, parallel computation is intrinsically supported. ODM supports the 5 core variants of BLAST (BLASTN, BLASTP, BLASTX, TBLASTN, TBLASTX. BLAST can be used to just return the `seq_id`, expect value and score, or to return that information along with full alignment details. The value of implementing BLAST in the database comes from: 1. The ability to craft complex queries, enabling complex analytical pipelines including BLAST searches 2. The ability to subselect portions of the database using SQL 3. Not having to export the sequence data and preprocess them to create BLAST datasets and import the results back into the database.

Statistical Functions

This feature provides the most commonly used statistical functions as part of the database, including summary statistics, hypothesis testing, ANOVA analysis, distribution fitting, and enhancements to cross-tabulations. As a result, you no longer need to extract data to external statistical engines for these analyses.

Document Clustering

This feature provides a statistical clustering package for automatic organization of documents.

Support Vector Machines

This feature provides a text mining classification algorithm.

Document Named Entity Extraction

This feature provides document metadata extraction by providing statistically assisted discovery and extraction of patterns from text.

See Also:

Oracle Database SQL Reference for information about statistical functions and floating-point numbers

Oracle Database Application Developer's Guide - Fundamentals for information on floating-point numbers

PL/SQL Packages and Types Reference for information on PL/SQL packages relating to floating-point numbers

Data Mining Support for Analytic Applications

Data mining operations are enhanced in this release. There are improvements in the data preprocessing and automated binning capabilities (which support outlier and missing value accounting and dynamic discretization), new attribute importance algorithms (Mutual Information and Pearson Correlation), Adaptive Bayes Network enhancements, and better Model Seeker capabilities. Development and deployment of data mining solutions are also improved, with provision of new JSR-73 standard-based Java components that allow the easy specification of data mining operations and the automatic generation of code to perform these operations. In addition, Oracle Enterprise Manager can be used to monitor and manage the Oracle Data Mining environment.

Frequent Itemsets

This feature provides an efficient mechanism for computing frequent itemsets, a common computation in data mining applications. Frequent itemsets are typically used in market-basket analysis, to find which items are most frequently purchased together.

DBMS_DATA_MINING PL/SQL Interface for Oracle Data Mining

This feature expands customer base for data mining by providing a PL/SQL interface in addition to the existing Java interface.

Features Extraction using Non-Negative Matrix Factorization Algorithm

Support feature extraction for life science and text mining.

Enhanced Data Preprocessing

This feature improves the performance and ease of use of essential data preprocessing tasks (transformations) required by algorithms.

Enhanced Adaptive Bayes Network

This feature enhances the model building and scoring functionality of Oracle Data Mining, and improve accuracy of the resulting models.

Multi-User Access Control

Oracle standard database security is now extended to provide added security to Oracle Data Mining user data and data mining results.

Oracle Data Mining Client (DM4J)

New Oracle Data Mining Client (DM4J) based on Oracle JDeveloper components enable graphical specification of Oracle Data Mining (ODM) objects and a graphic user interface for interacting with key Java objects and processes in the ODM server.

Support Vector Machines (SVM)

Support Vector Machines (SVMs) introduce a state-of-the-art classification and regression algorithm. SVMs are known from the literature to produce highly accurate models in difficult domains such as bioinformatics, image and text classification.

See *Oracle Database SQL Reference* for information on grouped table outer joins and SQL model calculations and *Oracle Data Warehousing Guide* for information on summary management as well as information on grouped table outer joins, SQL model calculations, and upsert functionality.

Improved Very Large Database (VLDB) Support

Several new capabilities and improvements have been provided to support customers who are planning to build data warehouses and transaction systems that will grow to hundreds or even thousands of terabytes. First of these is support for ultra large datafiles in an Oracle environment, which raises the limit of data addressable by an Oracle database to 8 exabytes (8 million terabytes). Partitioning improvements are also provided, with support for hash partitioning of global

indexes, providing increased throughput for applications performing very high numbers of concurrent inserts. In addition, partitioning capabilities have been expanded to include support for index-organized tables (IOTs), with support for list partitioning, partitioning of IOTs containing large object binaries (LOBs), and automatic global index management. Oracle Enterprise Manager also provides complete management of partitioned tables, materialized views, and indexes.

Bigfile Tablespace

A database can now contain up to 8 exabytes (8 million terabytes) of data. You can also store data in much larger files, thus decreasing the number of files in large databases.

Enhanced Partition Management in Oracle Enterprise Manager

This feature provides a user-friendly, intuitive graphical user interface for complete management of the wide range of partitioning options provided by Oracle.

Global Partitioned Indexes - Hash Partitioning

You can now hash-partition indexes on tables, partitioned tables, and index-organized tables. This feature provides increased throughput for applications with large numbers of concurrent inserts.

Partitioned Index-Organized Tables - List Partitioning

You can now list-partition index-organized tables, which can improve performance and manageability.

Partitioned Index-Organized Tables - Global Index Maintenance

Oracle now automatically maintains global indexes when data definitions language (DDL) operations are executed against partitioned index-organized tables. This feature greatly reduces the complexity of partition maintenance operations.

Partitioned Index-Organized Tables - LOB Support

LOB columns are now supported in all types of partitioned index-organized tables.

Enhanced Bitmap Index Performance and Space Management

Bitmap indexes now perform better and are less likely to be fragmented when subjected to large volumes of single-row data manipulation language (DML) operations.

Local Partitioned Indexes Manageability Improvements

You can now specify the location of new local index partitions during partition maintenance operations, enabling automatic maintenance of local indexes for any partition maintenance operation.

See Also:

Oracle Database SQL Reference for information on partitioning enhancements and index enhancements

Oracle Data Warehousing Guide and *Oracle Database Administrator's Guide* for information on ultra-large datafiles

PL/SQL Packages and Types Reference or information on using the `DBMS_ROWID` package for ultra-large datafiles

Oracle Data Warehousing Guide and *Oracle Database Concepts* for information on partitioning improvements

OLAP Support for Analytic Applications

This release contains support for new OLAP capabilities using the built-in analytical workspaces of the Oracle Database. New PL/SQL and XML-based interfaces are provided for the creation of workspaces based on the cubes and dimensions defined in the OLAP catalog in the database. Measures and calculations can also be similarly defined and created. These new interfaces are used directly or by way of Oracle Enterprise Manager to define and build analytical workspaces, removing the need for the user to learn OLAP DML commands. New cross-tabular analysis capabilities are also provided, supporting the aggregate of attributes within a dimension, such as color by size within a product dimensions. The release contains significant focus on performance enhancements. New parallel capabilities are provided for `AGGREGATE` and SQL `IMPORT` operations, making it much faster to load and materialize the analytical workspaces from relational information.

Parallel AGGREGATE Command

The Oracle OLAP `AGGREGATE` command is now run in parallel, which reduces the amount of time required to materialize summary data in analytic workspaces.

Parallel SQL IMPORT Command

The amount of time required to load data from relational tables and views into multidimensional datatypes in analytic workspaces is reduced, because the process is now run in parallel.

PL/SQL Creation of Analytic Workspaces

A new PL/SQL interface uses Oracle Enterprise Manager to build analytic workspaces based on Cubes and Dimensions as defined in the OLAP catalog. SQL-oriented developers no longer need to learn and use OLAP DML commands to build analytic workspaces.

Intra-Dimensional (Attribute) Aggregation

The OLAP AGGREGATE command now provides support for cross-tabular analysis.

Enhanced Composite Dimensions Indexing

New indexing techniques for COMPOSITE dimensions in analytic workspaces result in better query performance in cases with many dimensions and sparse data.

PL/SQL OLAP Measure Calculation Definition

A new PL/SQL interface defines multidimensional calculations in analytic workspaces. SQL-oriented developers no longer need to learn OLAP DML commands to define such calculations.

Complete Analytical Workspace Management

Analytic workspace management tools in Oracle Enterprise Manager provide a graphical user interface for creating and refreshing multidimensional datatypes in analytic workspaces without the need for programming.

Multi-Writer Support

Different users or sessions can now attach different multidimensional datatypes within the same analytic workspace. This significantly simplifies application development in cases where multiple users must write to the same analytic workspaces.

XML Interface for Analytic Workspace

The XML interface to analytic workspaces eliminates the need for application developers to learn OLAP DML syntax in order to build and add analytic content to analytic workspaces.

Reduced Information Cycle Time

This release of the Oracle Database provides significant new capabilities to extract, load, and transform data. The first of these is a new simple, scalable, and nonintrusive change data capture framework that allows administrators to asynchronously capture and publish changed data to data warehouses, data marts, and even applications. For large-scale bulk movement of data, cross-platform transportable tablespaces are provided, allowing large amounts of data to be very quickly moved from one database platform to another. New external table capabilities are also provided that support the unloading of data to flat files for data propagation or storage. Newer versions of SQL*Loader and the Import and Export utilities that take advantage of new high-performance data pump capabilities are also provided.

Asynchronous Change Data Capture

This feature provides a framework for capturing change data, publishing it, and enabling applications to subscribe to the change data in a controlled fashion. Change capture occurs asynchronously based on the information in the Oracle redo logs.

Cross-Platform Transportable Tablespaces

The transportable tablespace feature now enables tablespaces to be transported across different platforms.

Enhanced Table Functions

Parallel pipelined table functions have been enhanced to improve performance by returning only required rows. In addition, anonymous return types are supported for `AnyDataSet` table functions.

External Tables Unload

You can now load and transform large volumes of data into a platform-independent, Oracle proprietary flat file for data propagation or storage, either serially or in parallel.

Enhanced MERGE Functionality

The MERGE statement has been extended to cover a broader variety of complex and conditional data transformations, leading to faster loading of large volumes of data.

Data Pump Export and Import Utilities

The Data Pump Export and Import utilities provide very high-speed bulk movement of data and metadata from database to another. These utilities offer several significant advantages over the original Export and Import utilities, including: the ability to completely restart export and import jobs; the ability to detach from and reattach to long-running jobs; the ability to estimate how much space an export job would consume; support for export and import operations over the network; and support for fine-grained object selection, based upon objects and objects types.

Parallel Data Pump Export and Import

The new Data Pump Export and Import utilities can each be run in parallel, resulting in better performance in loading and unloading data and metadata.

SQL*Loader Direct-Path Load Support for ROWID Datatype

This feature produces faster SQL*Loader load times for tables containing ROWID columns.

SQL*Loader Direct-Path Load Support for VARRAY Datatype

This feature produces faster SQL*Loader load times for tables containing VARRAY columns.

SQL*Loader Direct-Path Load Support for XMLType Tables

This feature results in faster SQL*Loader load times for XMLType data in schema-based XMLType tables.

See Also:

Oracle Database Utilities for information on the new Data Pump Export and Import utilities as well as information on external tables and SQL*Loader direct-path support for ROWID, VARRAY, and `XMLType` datatypes

Oracle Data Warehousing Guide and *PL/SQL Packages and Types Reference* for information on asynchronous data change capture

Oracle Database SQL Reference for information on `MERGE` enhancements

Oracle Data Warehousing Guide for information on asynchronous change data capture

Oracle Database Administrator's Guide for information on cross-platform transportable tablespaces

Oracle Database Concepts for more information about external tables, data pump export and import, and other utilities

SQL Support for Analytic Applications

This release introduces interrow calculations by way of new SQL clauses, which add support for symbolic cell addressing and automatic and custom formulas. These new capabilities make it easy to build models and perform complex calculations without needing to code multiple joins and union clauses, which can soon become unwieldy. In addition to query operations, `INSERT`, `UPDATE`, and `MERGE` operations are also supported by the new capabilities, allowing users to build persistent models and forecasts. In addition to these new, powerful SQL capabilities, there are improvements in other areas of SQL analytics capabilities as well.

Partition Outer Join

This extension to the ANSI join syntax improves performance and simplifies SQL queries for time-based calculations.

Increased Number of Aggregates per Query

There is no longer a limitation on the number or size of aggregations in a single SQL statement.

SQL MODEL Clause

Queries and subqueries can include new syntax that enables highly expressive computations using sets of interrelated formulas. The feature provides a building block for complex calculations such as forecasts and budgets. With the `MODEL` clause, relational tables are treated as n-dimensional arrays and inter-row references can be specified without SQL joins and unions.

Upsert Through SQL Model Calculations

The SQL `MODEL` clause supports upsert operations, enabling easy `INSERT` and `UPDATE` of calculated values for building business models such as forecasts and budgets.

Summary Management OLAP Query Performance Improvements

Query rewrite enhancements for OLAP queries enable greater flexibility in creating materialized views in two ways. First, users have more choice about the fraction of the data that is preaggregated. Second, users have more choice about the number of materialized views in which the preaggregated data is stored.

Summary Management - Enhanced Partition-Aware Materialized View Refresh

This feature extends the materialized view refresh functionality by optimizing refresh of materialized views that are partitioned on a column with a functional dependency on the partitioning columns of the underlying table(s).

Summary Management - EXPLAIN PLAN Shows Materialized View Access

The `EXPLAIN PLAN` statement now indicates that a materialized view was accessed or used by query rewrite, instead of showing only table access.

Summary Management - Enhanced Dimensions

You can now specify a name for an attribute of a dimension, and you can display the structure of a dimension using the `DESCRIBE_DIMENSION` API.

Summary Management - Nested Materialized View Refresh

Materialized view refresh has been extended so that you can now refresh all the materialized views in a nested materialized view.

Summary Management - Enhanced PCT Refresh

Partition change tracking (PCT) refresh now supports list partitioning and ROWID, when it is used as a partition marker. It also performs TRUNCATE operations when appropriate.

See *Oracle Database SQL Reference* for information on grouped table outer joins and SQL model calculations and *Oracle Data Warehousing Guide* for information on summary management as well as information on grouped table outer joins, SQL model calculations, and upsert functionality.

Content Management

Support for nonrelational data stored in a database has also been improved in this release, with new formats for multimedia data, addition of text classification capabilities, and improvements in search and retrieval capabilities.

Improved Multimedia Support

As larger and larger documents are stored and managed in content management applications, the 4 GB restriction on LOBs stored in the database is removed in this release, raising the limit to 8 to 128 terabytes. Standards-based access to image data is also supported through the SQL Multimedia Still Image Standard (ISO / IEC 13249-5). The Java Advanced Imaging package used by Oracle *inter*Media for image support has been upgraded to JAI 1.1.1_01, and additional image processing is now supported. In addition, new audio and video media formats are supported, including MPEG4 and MPEG2, with the associated automatic recognition and extraction of metadata supported by these formats.

ISO/IEC 13249-5 SQL Multimedia Still Image Support

This feature implements Oracle *inter*Media compliance with ISO / IEC Still Image standards, enabling application portability.

Terabyte-Size LOBs

As large documents proliferate in content management applications, the Oracle Database has been enhanced to store large documents from 8 to 128 terabytes in size.

Java Advanced Imaging 1.1.1_01 Support

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

Native Support for Additional Audio/Video Media Formats

You no longer have to write special code to parse MPEG2 and MPEG4 multimedia content. *inter*Media can now extract metadata from these formats and make it available for indexing and querying.

See Also:

Oracle interMedia Reference for information on ISO / IEC 13249-5 SQL / MM Still Image support, the new operators for storing and processing image content, audio media formats, and video media formats

Oracle Database Application Developer's Guide - Large Objects for information on unlimited-sized LOBs

Improved Text Globalization Capabilities

For better handling of documents in a global environment, automatic discovery of the language and character set of unknown documents is provided, as well as support for new German spelling rules, Japanese adverb and verb stem indexing, and additional Japanese and Chinese character sets.

Document Character Set and Language Detection

This feature enables you to determine the character set and language, including Asian languages, of unknown documents.

New German Spelling Rules Support

Transliteration spelling rules, as mandated by the German government, are now supported, enabling queries against both the old and new spellings of a word.

Chinese GB18030 Character Set Support

This feature provides support for this increasingly common Chinese character set.

Enhanced Japanese Language Support

Japanese language support has been enhanced to include adverb and verb stem indexing, a Japanese-specific "fuzzy machine," a customizable user dictionary, and full support for AL32UTF8.

Improved Text Manageability and Usability

To make it easier to develop text-enabled applications, a series of JDeveloper Wizards are provided that enable the generation of catalogs and text-driven applications. In addition, an HTML-based Thesaurus Manager is also provided. Manageability of text documents in the database has also been improved: locally partitioned text indexes can be created online, and documents can be inserted during index creation and rebuild. Document services, such as highlighting, themes, and gists, no longer require the building of a text index, and a text framework allows user-defined components to be plugged in while invoking these document services.

Text Classification and Routing

Oracle Text now supports content-based as well as text-only routing of word processor format documents.

Classification Training Set Wizard for JDeveloper

JDeveloper now includes a classification training set application.

Thesaurus Manager

The Thesaurus Manager is an easy-to-use HTML-based thesaurus management tool.

Catalog Wizard for JDeveloper

JDeveloper now includes a catalog search application generator.

Text Wizard for JDeveloper

JDeveloper now includes a Text application generator.

Indexless Document Services

This feature lets you call document services without a Text index.

Progressive Text Query Relaxation

This feature provides an easier way to query in several dimensions.

Query Log Analysis

This feature enables you to customize search sites according to actual end-user queries.

Multipart MIME Filter

This new configurable filter can understand multipart MIME-encoded documents and mail messages.

Highlighting for INPATH and HASPATH Operators

This feature provides more complete handling of XML paths.

Positional Operator

An order position identifier for XML sections supports more complete handling of XML path expressions.

Improved Text Quality of Retrieval

The quality of retrieval of text searches has also been improved. Query log analysis is supported, allowing searches to be customized based on the success of past user queries, and a new progressive text query relaxation template allows multiple queries, with relaxed restrictions if earlier queries fail.

Link Analysis

A new algorithm boosts scores based upon popularity, and affords better search ranking for Web pages and content.

NEAR-ACCUM Text Operator

This Text operator combines the effects of proximity (NEAR) and progressive relaxation (ACCUM) operations, increasing the quality of search result sets, especially for passage-based queries.

Theme Proximity Searching

This feature enables you to use proximity (NEAR) with theme queries for better relevance ranking.

Location Services

New in this release, Oracle Locator supports parallel spatial queries; improves performance for spatial queries and joins, index updates and inserts; and conforms to the OpenGIS Simple Features Specification. New capabilities for the Oracle Spatial option include topology and network data models, a GeoRaster data type, spatial analytic functions, and a geocoder. Workspace Manager adds event management support, enhanced usability, and more support for database features.

Overall

This release continues to address key requirements of our location services and spatial customers. It includes new capabilities and performance enhancements for Geographic Information Systems customers (such as land management, energy, defense / homeland security) and location services providers and users. New in this release, Oracle Locator supports parallel spatial queries and improves performance for spatial queries, index updates, and index inserts. New Oracle Spatial capabilities include topology and network data models, GeoRaster data type support, and spatial analytic functions. Workspace Manager now provides event management support, enhanced usability, and more support for database features.

GeoRaster Support

This feature provides an open data format for database storage, management, and retrieval of raster data with a location reference or geo-reference (such as satellite imagery, remotely sensed data, gridded data). It supports the specific image processing requirements of GIS and remote sensing applications such as homeland security, energy, defense, and land management.

Spatial Analytics and Geocoder

New server-based spatial analysis capabilities include classification, binning, association, and spatial correlation - essential for business intelligence applications. In addition, a geocoding engine provides international address standardization, geocoding, and POI matching by querying geocoded data stored in an Oracle Database.

Network Data Model

This feature provides an open data model to store and manage networks (graphs) in support of utility, transportation, and life sciences applications. Path computations and network-traversal queries are supported, as well as segment-level updates for network data by data providers. In addition, a scalable routing engine provides

fastest or shortest routes, summary or detailed driving directions, and time and distance along a street network from a single location to multiple destinations.

Topology Data Model

This feature provides an open data model to store and manage topology, which is required for managing data in land information systems. It provides APIs that support edit and update capabilities for managing parcel and linear feature data as well as a query model to efficiently relate spatial objects by feature.

Spatial Index Updates and Queries - Performance Improvements

Spatial R-tree index update time has been reduced by 40% or more, which is especially helpful for location-based services and enterprise geographic information systems, and index inserts run 5 to 10 times faster. Spatial distance and "relate" queries (which check for specific location relationships) now run 20-40% faster than before, and spatial joins run 2-6 times faster.

Parallel Spatial Queries

Spatial queries can now run in parallel on partitioned spatial indexes, improving the performance of within distance, nearest neighbor, and relate queries. Performance scales with the number of CPUs used to execute a query. This helps location service and land management applications, which need to execute high volumes of spatial queries quickly.

OpenGIS Simple Features Specification Conformance

This feature ensures that the new release continues to enable third-party tools as well as GIS and location service applications to conform to the OpenGIS interface standards for location-based services and geographic information systems, when using Oracle Locator or the Oracle Spatial option.

Workspace Manager Event Management

Workspace Manager adds the ability to define workspace event callbacks that associate application logic, such as business practices, with Workspace Manager operations. Events exist for the following workspace operations: create and delete savepoint; create, merge, remove, refresh, and rollback workspace. User-defined event handlers can be written in PL/SQL or any other language that can be wrapped in a PL/SQL procedure.

Enhanced Workspace Manager Usability

Workspace Manager adds support for multiparent workspaces that let a child workspace merge with and refresh from multiple parent workspaces. It also supports continually refreshed workspaces anywhere in a workspace tree with optimistic locking, and enables other users in the workspace tree to access versions locked in a workspace.

Enhanced Workspace Manager Database Support

Workspace Manager adds support for `UNIQUE` constraints, SQL*Loader, materialized views, Virtual Private Database, table statistics, and nested tables. Support has been enhanced for many DDL operations and for finer grained import and export.

Application Development

Improvements have been provided to simplify development of high-performance, global applications in standard languages.

Database Centric Application Development Environment

Oracle HTML DB is a declarative development tool and a framework for the development and deployment of database-centric web applications. Oracle HTML DB accelerates application development through built in features such as design themes, navigational controls, form handlers and flexible reports. Using only a web browser, users can quickly assemble a sophisticated database driven web application.

Service Administration

With Oracle HTML DB service administration, you can maintain a hosted development service consolidating many independent development projects in a single database.

SQL Workshop

SQL Workshop enables you to build database objects and run SQL statements and SQL scripts from a web browser. In addition, it enables you to store and retrieve data, execute SQL commands, and perform many additional tasks to speed and simplify development.

Data Workshop

With the Data Workshop, you can import data into, and export data from the database using a web browser. Supported formats include: structured text, such as comma separated and tab delimited data, data from spreadsheets and certain XML documents.

Application Builder

Application Builder in HTML DB enables you to build database-centric interactive web applications. Application Builder helps you assemble an HTML user interface on top of database objects such as tables and procedures. Once your application is assembled, the engine takes care of rendering your application, using templates and UI elements that you specify. It also handles all the processing, validation, and branching that your application needs.

Globalization and Unicode

To aid development of global applications, this release provides a Globalization Development Kit (GDK) that includes comprehensive programming APIs, tools, and documentation that address many of the design, development, and deployment issues encountered while creating global applications. The Oracle NLS definition files (language, territory, linguistic sort, and character set) also become platform independent, reducing the need to regenerate new binary files on each platform of developers using these capabilities. In addition, the database provides expanded locale coverage with support for new languages and territories, and support for Unicode 3.2.

Globalization Development Kit

The Oracle Globalization Development Kit (GDK) is a toolkit that simplifies the development process and reduces the cost of developing Internet applications that will be used to support a global environment. This release of the GDK includes comprehensive programming APIs (Java and PL/SQL), code samples, and documentation that address many of the design, development, and deployment issues encountered while creating global applications. The GDK is a set of Java and PL/SQL APIs that provide application developers with the framework to develop globalized Internet applications using the best globalization practices and features designed by Oracle.

Character Set Scanner Utilities

The Database Character Set Scanner has been enhanced to support the scanning of nested tables and character semantics objects. The database scan summary report now provides additional information on the source database along with statistics on possible size expansion. A new utility, the Language and Character Set File Scanner, provides automatic identification of language and character set pairs for plain text files. this is a statistically based utility that is capable of detecting language and character sets supported throughout the world.

Database Character Set Scanner CSALTER Script

The CSALTER script is a DBS tool for special character set migration.

Expanded Locale Coverage

This feature adds new territories and languages, and augments existing definition files with additional information. In previous releases, Oracle defined language and territory definitions separately. This resulted in the definition of a territory being independent of the language setting of the user. In this release, some territories can have different date, time, number, and monetary formats based on the language setting of a user. This type of language-dependent territory definition is called a locale variant. Also, NLB files that are generated on one platform can be transported to another platform by, for example, FTP. The transported NLB files can be used the same way as the NLB files that were generated on the original platform. This is convenient because locale data can be modified on one platform and copied to other platforms.

Unicode 3.2 Support

This feature provides support for the Unicode standard, Unicode 3.2, by adding new Unicode code points, character classifications, and mapping information to existing Unicode character sets.

CLOB and NCLOB Implicit Conversions

This feature provides implicit conversion between CLOB and NCLOB datatypes. Global internet applications that support multiple national language character sets no longer require development and deployment of explicit function calls to achieve this conversion.

See Also:

Oracle Database Globalization Support Guide and *PL/SQL Packages and Types Reference* for information on the `UTL_ENCODE`, `UTL_I18N`, and `UTL_LMS` packages, which are part of the Globalization Development Kit

Java, JDBC, and Web Services

In this release, the Java Virtual Machine (JVM) built into the Oracle Database is brought up to J2SE 1.4 compatibility. Overall improvements are provided in the supplied JDBC drivers, including support for varray enhancements, LONG-to-LOB conversions, support for the INTERVAL DAY TO SECOND datatype, autogenerated keys retrieval, and Web RowSet support. New JDBC 3.0 J2EE support is included for DATALINK datatype and REF interface, Connector Architecture Resource Adapters, and connection caching. The thin JDBC driver is improved to support key capabilities provided by the similar JDBC OCI driver, including support for PL/SQL index tables, and passing, retrieving, and registering parameters by name. Loadjava performance has also been dramatically improved.

JDBC Thin Driver PL/SQL Index Table

This feature enables you to send and receive PL/SQL tables in the thin driver. For example, you can exchange Java collections with PL/SQL collections.

JDBC 3.0 Named Parameter

This feature enables JDBC applications to pass parameters by name with `CallableStatement` and to register and retrieve output parameters by name.

SOAP Client for Database Web Services

This feature lets the database consume external Web Services. As a result, Java classes in the database as well as SQL statements, PL/SQL packages, triggers, table functions, and so on can call out to the external Web Services.

Java-in-Database Web Services

This feature renders Java classes in the database that implement data-bound logic (stored procedures) as Web Services. Java portability allows the partitioning of these classes between the middle tier and the database.

SQL Query Web Services

Middle-tier applications, desktop applications, and tools can search, retrieve, and extract business data (relational, XML, text, or spatial) from the database using Web Services mechanisms. As a result, you can implement data-bound services, such as predefined warehousing queries, catalog search queries, and Map / GIS services, and expose them to other applications by reusing predefined SQL queries.

Web Services Data Source

This feature renders the output resulting from external Web Services calls as regular SQL output, which can then be used in the `FROM` clause of a SQL query. SQL functions can be applied to `WHERE` clauses before returning the query results. This conversion of external Web Services into SQL data sources is a building block of enterprise data integration.

Enhanced PL/SQL Web Services

This feature improves PL / SQL Web Services by supporting BOOLEAN, CLOB, BLOB, and PL / SQL types, enabling you to use most of your existing PL / SQL packages as Web Services.

Consuming External Web Services from SQL and PL/SQL

Any SQL-enabled tool or application can transparently and easily consume dynamic data from external web services.

Consuming External Web Services from Database Using Java

This feature offers an easy-to-use interface for calling-out web services. Java classes running in the database can simply and directly invoke external web services, using their Java proxy, insulating Developers from low-level SOAP programming.

SQL DML Web Services

This feature makes possible the implementation of a Web Services as a single or a group of database `INSERT`, `UPDATE`, and `DELETE` operations.

OracleJVM: Native Java Interface

This feature provides simplified application integration. It enables client-side and middle-tier Java applications to invoke Java directly in the database without the need for a PL / SQL wrapper.

Enhanced JDBC VARRAY Support

Aggregation and set operations are now available to data mining and JDBC applications.

JDBC IEEE DOUBLE and FLOAT Datatypes

The new SQL floating-point number datatypes are supported in JDBC, enabling Java/J2EE applications using JDBC to perform faster arithmetic calculations without loss of information and with reduced storage.

JDBC LONG to LOB Conversion

This feature improves JDBC application portability by enabling the conversion of CLOB and BLOB data so that it is compatible with LONG, RAW, and LONG RAW, and by simplifying CLOB and BLOB manipulation in JDBC applications.

JDBC INTERVAL DAY TO SECOND

JDBC applications can now utilize the database INTERVAL DAY TO SECOND datatype for improved time management.

JDBC RAC/HA Fast Connection Failover

This feature ensures High Availability of Data-sources in RAC environment. On instance up, JDBC proactively balances data-source connections across all available instances; conversely on instance or node down, it removes invalid connections from the cache, transparently to the application.

JDBC 3.0 J2EE Connector Architecture Resource Adapter

The Oracle JDBC driver can now function as a JCA-compliant resource adapter for Oracle databases, enabling improved pluggability, packaging, and deployment.

JDBC 3.0 DATALINK Datatype and Ref Interface

This feature enables distributed JDBC applications to retrieve and store references to external data, obtain DATALINK-related metadata information, and map into Java a reference to a SQL structured type value in the database.

JDBC Web RowSet

This feature enables disconnected applications such as Web Services clients or J2EE components to fetch a collection of rows from database tables (or other data

sources) in XML format and to scroll through, update locally, and synchronize the rows back to the data source.

OracleJVM J2SE 1.4.x Compatibility

Java-client applications that are J2SE 1.4.x compatible can now be moved seamlessly into OracleJVM as Java Stored Procedures. New Java libraries can be loaded by using new data-driven functionalities in the database.

Implicit JDBC 3.0 Connection Caching

This feature provides JDBC applications with new, faster, simplified connection caching for DataSource. Java / J2EE applications will benefit from transparent access to the cache, support for multiple users, and the ability to request connections based on user-defined profiles.

JDBC Connection Cache Manager

The Connection Cache Manager APIs let the middle-tier (application servers and packaged applications frameworks) create, manage, and maintain caches. They can also bind connections in caches to DataSources, configure and search connections based on specified profile.

Other APIs and Precompilers

To facilitate C++ development, the Oracle C++ Call Interface (OCCI) has also been improved in this release, with support for UTF16 character data and other globalization capabilities. In addition, performance has been improved with internal object management algorithms, SQL statement caching, and stateless connection pooling. Distributed transaction capabilities are supported in OCCI with XA support in the interface.

Instant Client

Applications that use OCI and OCCI, including the type-2 JDBC driver, can now easily install only the required files instead of installing and configuring the entire Oracle client; resulting in simplified deployment and smaller memory footprint.

OCCI Globalization Support

Globalization support for OCCI users enables them to bind and define multibyte and UTF-16 character data and to use strings in any character set.

OCCI Stateless Connection Pooling

Formerly, every time a connection was chosen, it was authenticated with the server. In stateless connection pooling, a connection is picked up and given to the user without having to do a server round-trip for authentication.

OCCI Performance Improvements

OCCI applications now experience improved performance in the creation and update of object instances as well as fetch, delete, and pickle and unpickle operations.

OCCI Statement Caching

By caching the cursors containing parsed SQL and associated data structures, this feature provides improved performance for statements that are reused. You can cache statements for a particular connection and enable or disable caching for a connection. You can also tag an individual statement so it can subsequently be retrieved from the cache.

OCCI XA Support

C++ developers using OCCI are now able to use distributed transactions options as provided by an XA-compliant database.

Preserved User Code in OTT-Generated Classes

Applications using Oracle types need to use the Oracle Type Translator (OTT) to generate classes for the types stored in the database. Application developers often must modify OTT-generated code to handle application-specific requirements. OTT now preserves the user-added code for use in subsequent calls to OTT.

See *Oracle Call Interface Programmer's Guide* for information on OCI installation.

PL/SQL Improvements

This release contains significant PL/SQL enhancements, both for the application end-user and for the application developer. It introduces a brand-new optimizing compiler that generates faster code, and it exposes improvements in Native PL/SQL compilation, in particular, adding full support for RAC. A number of powerful new language features are added (for example, support for ANSI syntax for multiset operations on nested tables, for regular expressions and for IEEE floating point arithmetic). New utility packages are added for data compression and to provide a simple API for sending email.

Sparse Table Bulk Binds

This feature extends the usefulness and efficiency of the FORALL bulk DML syntax by allowing the bound array to be sparse, or by allowing a second array to denote a sparse subset in the bound array.

Collection Set Operations

SQL introduces support for ANSI compliant multiset operations (for example, equality, union, and intersect). Oracle's implementation of the multiset is the nested table, which, since its introduction, has been available also as a PL/SQL datatype. The new, compact syntax allows efficient implementation of common operations on collections which would otherwise require tedious coding and would result in a less efficient implementation, and these benefits are also brought to PL/SQL programs. In addition, support is now provided for using a collection in the semantic role of an IN list.

Database Storage of PL/SQL Native Compilation Units

The output of PL/SQL native compilation is now stored in the database as BLOB data. This simplifies backup procedures in a native PL/SQL environment. The performance benefits of PL/SQL native compilation are now available in an Oracle Real Application Clusters configuration. The usability of the configuration steps has been improved.

Fine-Grained Debug Privileges

Debugging privileges can now be granted on a per-program-unit basis. As a result, application developers can debug appropriate code without having access to data they should not see or change.

Regular Expressions

The new SQL builtins for POSIX-compliant regular expressions are also available in a pure PL/SQL setting. This allows the compact and efficient implementation of a new class of algorithms as stored database procedures.

User-Specified Quoting Character

You can now choose any convenient delimiter and define it dynamically as the quoting character, resulting in increased usability for application developers and database administrators.

UTL_COMPRESS Package

This new package delivers the familiar functionality of zip and unzip utilities in a PL/SQL environment. It lets you compress and uncompress a RAW bytestream, and guarantees the exact original bytestream after the round trip. This package also handles lobs, and has features to add and later extract several pieces into a single archive.

UTL_MAIL Package

This package makes it possible for a PL/SQL programmer to send programmatically composed e-mails by way of a PL/SQL API. It requires only the normal mental model of a user of the GUI email client, rather than an understanding of the underlying protocol (SMTP) features.

Compile-Time Warnings

This feature, which is well-known in other programming languages, is now introduced for PL/SQL. Constructs that are legal, but are in some sense dubious, will now optionally raise a compilation warning. Warnings can be turned on or off individually or by category. For example, a particular warning in the performance category will be raised when an opportunity is detected to improve the efficiency of parameter passing by using different syntax.

DBMS_PROFILER Improvements

This package now helps you better measure the performance characteristics of your PL/SQL programs. It enables profiling of natively compiled PL/SQL program units, removes correspondence problems between reported and actual source code line numbers, and flushes profile data to a table faster at the end of the profiling run.

PL/SQL Optimizing Compiler

This release replaces the PL/SQL compiler with a completely new version that provides a framework for and support of many optimizations. The result is improved performance, especially for computationally intensive PL/SQL programs.

See *PL/SQL Packages and Types Reference* for more information.

SQL Language Improvements

To aid migration from other database systems, this release provides support for case-insensitive queries and sorts, allowing users to search and sort their data

regardless of the casing and accent of the characters. Regular expressions are also supported in SQL and PL/SQL, allowing developers to write one-line queries that previously would have taken multiple lines of SQL code. This POSIX-compliant implementation also supports multilingual queries and is locale sensitive. A new capability, Expression Filtering, is also supported in this release, that allows application developers to manage and evaluate conditional expressions that describe users' interests in data. Other SQL improvements include new `CONNECT BY` processing that supports ancestor-descendant pairs, and new collection performance and type evolution enhancements.

Case-Insensitive and Accent-Insensitive Query and Sort

Oracle provides linguistic sorts and queries that use information about base letter, accents, and case to sort character strings. Now you can also specify a sort or query on the base letters only (accent insensitive) or on the base letter and the accents (case insensitive).

BINARY_DOUBLE and BINARY_FLOAT Datatypes

This release introduces single- and double-precision floating point number datatypes. These new types enable faster arithmetic calculations and reduce storage requirements. These datatypes also enable you to implement published algorithms that specify IEEE behavior for arithmetic operations. Intense number crunching computations (whose scale and precision requests can be accommodated by the IEEE types) will run very substantially faster using these types than using NUMBER. Additionally, this feature enables cleaner integration with XML and Java environments because it provides similar numeric datatypes for those environments.

Enhanced Collections

Collection types have been enhanced to provide better performance and greater functionality. You can now alter the size of a VARRAY type object, use varrays in temporary tables, and specify different tablespaces for different columns of nested table type. Nested Table type now also supports comparison conditions (for example, Equal, Not Equal, Member Of) and ANSI SQL 2003 Multiset operations (for example, Multiset Intersect, Multiset Union).

Enhanced CONNECT BY Support

Additions to the `CONNECT BY` clause enhance queries of hierarchical data in the following ways: All ancestor-dependent pairs can be returned (not just parent-child pairs); a new pseudocolumn specifies whether or not a given node is a leaf of a

hierarchy; a cycle in a hierarchy triggers an error message with information about the rows involved in the cycle; and the CONNECT BY clause is now supported for simple subqueries.

Expression Filter

Some classes of applications use conditional expressions to describe interests in expected data and business rules. Expression Filter provides an Expression datatype, SQL EVALUATE operator, and indexing. These allow conditional expressions to be stored in a column of a table and searched. The EVALUATE operator matches incoming data with the expressions to identify rows of interest. The incoming data can also be stored in a table and matched with expressions in another table using a join to derive complex relationships between the tables.

Extensible Indexing

Extensible Indexing adds the ability to alter user-defined domain operators, maintain global and local domain indexes during partition maintenance operations and enables parallel creation of local domain indexes.

SQL Regular Expressions

This release supports POSIX-compliant regular expressions to enhance search and replace capability in programming environments such as Unix and Java. In SQL, this new functionality is implemented through new functions that are regular expression extensions to existing functions such as LIKE, REPLACE, and INSTR. This implementation supports multilingual queries and is locale sensitive.

Row Timestamp

A new pseudocolumn consisting of the committed timestamp or SCN provides applications and users the ability to efficiently implement optimistic locking. Until now, when posting updates to the database, applications had to read in all column values or user-specified indicator columns, compare them with those previously fetched, and update those with identical values. With this feature, only the row SCN needs to be retrieved and compared to verify that the row has not changed from the time of the select to the update.

See Also:

Oracle Database SQL Reference for information on Oracle regular expression support, case- and accent-insensitive sorting, and hierarchical query enhancements

Oracle Database Globalization Support Guide and *Oracle Database Application Developer's Guide - Fundamentals* for information on Oracle regular expression support

Oracle Database Globalization Support Guide for information on case- and accent-insensitive sorting

Oracle Database Application Developer's Guide - Expression Filter and *Oracle Database Application Developer's Guide - Fundamentals* for information on Oracle Expression Filter

XML Application Development Support

To provide better support for standards-based access to XML, this release provides SQL / XML Standard compliance. Applications can use standard SQL / XML operators to generate complex XML documents from SQL queries and to store XML documents. The XML parser is also extended to support the updated and new W3C XML standards. There have also been major improvements in XML processing performance. The XML Developer's Kit (XDK) libraries and interfaces in Java, C, and C++ all transparently support the database XMLType, increasing throughput and scalability without high resource and processing costs. Additionally, the architecture has been redesigned using a pipeline process model and SAX to increase performance while reducing resources.

SQL/XML Standard Compliance

SQL / XML is a new part (Part 14) of the SQL standard that provides interoperability between SQL and XML. Applications can now use SQL / XML operators to generate complex XML documents from SQL queries and to store XML documents in relational tables.

Oracle XML DB LDAP Access Performance

This release provides improved performance for large-scale enterprises using Oracle Internet Directory and Lightweight Directory Access Protocol for access control.

XML DB Schema Evolution

This feature supports the evolution of XML schemas as your business requirements change.

XML Processing Performance Improvements

XDK libraries and interfaces in Java, C, and C++ now support XMLType data in the database, thus increasing throughput and scalability without high resource and processing costs.

XML Transformation Performance Improvements in C/C++

The C XSLT Processor is now over 100% more efficient in the XSL transformation of XML, yielding huge benefits in performance and scalability of XML-enabled applications and the XML database.

XML DB Internationalization

The Oracle XML DB and the Oracle XML database repository now support multibyte character sets and the use of multiple client characters sets. You can now set your client character set different from the database character set. Appropriate conversion will take place to present the XML data in the character set of the client. In addition, using FTP or HTTP, you can specify multibyte characters in the directory, filename, or URL, and you can transfer or receive data encoded in a different character set from the database. Oracle XML DB can handle all popular XML character encodings as long as the database character set supports characters in use. For full support of all valid XML characters, use UTF-8 as your database character set.

Extended XML Standard Support in C/C++

You can use the Oracle XDK as an XML platform because of its extensive and conforming support of the Internet XML Standards. This feature enables you to develop and deploy on a platform that does not lock you into proprietary interfaces. It also implements DOM Level 2 Range and Traversal to improve support for DOM.

Extended XML Standard Support in Java

You can use the Oracle XDK as an XML platform because of its extensive and conforming support of the Internet XML Standards. This feature enables you to develop and deploy on a platform that does not lock you into proprietary

interfaces. It also implements JSR-031 Java Bindings (enhanced class generator), which provide easy XML access for Java development.

XML Transformation and Querying Performance Improvements in Java

To help speed searches on XML documents, this release adds support for indexing XML documents based on the XMLSchema. It also adds support for transforming XML files using SAX-based streaming of XML or partially built DOM trees. This results in dramatic increases in scalability and performance for all of the Oracle XML infrastructure and XML-enabled applications, including Discoverer, JDeveloper, EJBs, and JSPs.

High Performance XML - Java

Companies and developers can now leverage new optimizations and APIs to improve XML processing performance.

XML DB XMLType View Performance

Query rewrite is now supported for `XMLType` views. This makes `XMLType` views easier to use and more like traditional relational views, and reduces the complexity of applications that use them.

See *Oracle XML Developer's Kit Programmer's Guide* for information on the XSLT compiler and the XSLT virtual machine, *Oracle XML DB Developer's Guide*, and *Oracle XML API Reference*.

Windows

Oracle Database 10*g* is closely integrated with the Windows operating system and Microsoft data access technologies. Its new database features improve data access and database performance on Windows. Many of these performance improvements are transparent, requiring no Oracle parameter or application code changes to be made. Oracle Database 10*g* is easier to use and exposes more advanced database features for Microsoft COM and .NET data access users, thereby improving developer productivity.

Data Access Support

This release provides further integration with Microsoft's COM and .NET Framework by exposing more Oracle advanced database functionality in an easy-to-use manner. These improvements include native XML DB support and additional Oracle datatype and PL/SQL support.

Oracle Objects for OLE Support for Oracle Datatypes and Grid

Oracle Objects for OLE (OO4O) now supports all of Oracle's timestamp (TIMESTAMP, TIMESTAMP WITH LOCAL TIME ZONE, TIMESTAMP WITH TIME ZONE) and interval (INTERVAL DAY TO SECOND, INTERVAL YEAR TO MONTH) datatypes. It has also added support for the new optimized float and double datatypes in this release. Additionally, OO4O is Grid-enabled, allowing developers to take advantage of Oracle's database Grid support without having to make changes to their application code.

ODP.NET Support for Advanced Oracle Features

The Oracle Data Provider for .NET (ODP.NET) has been enhanced to support advanced Oracle datatypes and features. .NET now supports PL/SQL Associative Array parameter binding; the new optimized float and double datatypes introduced in Oracle Database; and improved overall data access performance. Moreover, ODP.NET has introduced a new LOB property, InitialLOBFetchSize, to allow developers to tune their LOB applications more easily. ODP.NET also enables Oracle's Grid technology, allowing developers to take advantage of it without having to modify their code.

ODP.NET Support for XML DB

ODP.NET has been enhanced to include native support for Oracle XML DB, making XML easy to use and manage between the database and .NET. This new functionality introduces a native OracleXMLType data type for flexible and efficient XML manipulation. Developers will be able to query and save relational and object-relational data as XML using ODP.NET. When data is stored within XML DB, developers can query and save both non-schema (CLOB) and schema-based (object) XML.

Operating System Integration

This release provides new 32-bit and 64-bit Oracle Databases for Windows performance and scalability optimizations for enterprise data centers.

Optimized 64-bit Database

The Oracle Database has been enhanced for increased performance on 64-bit Windows. This performance improvement is transparent, requiring no database parameter changes.

Oracle Services for Microsoft Transaction Server Support

Oracle Services for Microsoft Transaction Server now supports distributed transactions set to a serializable isolation level.

2

Oracle Database 10*g* Documentation

This chapter lists the books available online that document Oracle Database 10*g*. It contains the following section:

- Documentation Titles and Online Documentation Structure

Documentation Titles and Online Documentation Structure

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

- Oracle Database 10g Server and SQL*Plus Documentation
- Oracle Data Mining Documentation
- Oracle Spatial Documentation
- Oracle OLAP Documentation
- Oracle Database10g Documentation for Windows
- Java Documentation
- Oracle Enterprise Manager Documentation

Oracle Database 10*g* Server and SQL*Plus Documentation

Table 2–1 Oracle Database 10g Server and SQL*Plus Documentation

Documentation	Part Number	Description
Administration, Concepts, and Reference		
Oracle Database New Features	B10750-01	Describes the new features of Oracle Database 10*g* and lists the documentation for this release.
Oracle Database Administrator's Guide	B10739-01	Explains how to create and to manage an Oracle database. Includes information about distributed processing.
Oracle 2 Day DBA	B10742-01	Provides a single source reference for administering an Oracle database in a small to midsize environment. If you are familiar with computers, but unfamiliar with administering an 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 Concepts	B10743-01	Describes how the Oracle Database functions and lays a conceptual foundation for more detailed information contained in other Oracle Database books.
Oracle Database Upgrade Guide	B10763-01	Explains the process of planning and executing database upgrades on the 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 Database Performance Tuning Guide	B10752-01	Explains in detail how to enhance performance by writing and tuning SQL properly, by using performance tools, and by optimizing instance performance.
Oracle Database Reference	B10755-01	Provides reference information about database initialization parameters, static data dictionary views, dynamic performance views, database limits, and SQL scripts that are part of the Oracle database system.
Oracle Database Utilities	B10825-01	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.

Table 2–1 (Cont.) Oracle Database 10g Server and SQL*Plus Documentation

Documentation	Part Number	Description
Oracle Database Globalization Support Guide	B10749-01	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.
Oracle Globalization Development Kit Java API Reference	B10971-01	Describes APIs for the Oracle Globalization Development Kit.
Oracle High Availability Architecture and Best Practices	B10726-01	Describes tested architectures and recommended practices that can be used to achieve high availability with Oracle Database products and features.
Oracle Database Sample Schemas	B10771-01	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 Quick Reference	B10758-01	Contains a high-level description of the SQL used to manage information in an Oracle database.
Oracle Database SQL Reference	B10759-01	Contains a complete description of the SQL used to manage information in an Oracle database. The first volume describes all of the building blocks of SQL statements. The remaining volumes describe the SQL statements alphabetically.
Oracle Database Heterogeneous Connectivity Administrator's Guide	B10764-01	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 Error Messages	B10744-01	Online only. Includes all database error messages.
Oracle Database Readme	B12304-01	Provides important last minute information not included in the Oracle Database Documentation Library.
Oracle HTTP Server Administrator's Guide	B12255-01	Describes how to administer Oracle HTTP Server.
Oracle HTTP Server mod_plsql User's Guide	B12303-01	Describes how to use `mod_plsql` with HTTP Server.
Oracle HTTP Server Installation Guide for Windows	B12144-01	Describes installation instructions for HTTP Server.

Table 2–1 (Cont.) Oracle Database 10g Server and SQL*Plus Documentation

Documentation	Part Number	Description
Backup and Recovery		
Oracle Database Recovery Manager Quick Start Guide	B10769-01	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 Basics	B10735-01	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 Advanced User's Guide	B10734-01	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 Recovery Manager Reference	B10770-01	This book describes RMAN syntax and recovery catalog views.
Data Warehousing		
Oracle Data Warehousing Guide	B10736-01	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 Warehouse Builder User's Guide	B12146-01	Describes Oracle Warehouse Builder functionality and how to define data objects, design ETL processes, deploy target systems, and manage metadata.
Oracle Warehouse Builder Release Notes	B12149-01	Provides important last minute information not included in the Oracle Database Documentation Library.
Oracle Warehouse Builder Installation and Configuration Guide	B12150-01	Describes the steps to install Oracle Warehouse Builder software, configure the Design and Runtime Repositories, and integrate Warehouse Builder with other products.
Oracle Warehouse Builder Transformation Guide	B12151-01	Describes the functions and procedures that characterize Oracle Warehouse Builder transformations. Warehouse Builder includes several predefined transformations, as well as a library of predefined functions and procedures to transform data. In addition, Warehouse Builder also uses SQL and PL/SQL as transformation languages.

Table 2–1 (Cont.) Oracle Database 10g Server and SQL*Plus Documentation

Documentation	Part Number	Description
Oracle Warehouse Builder Scripting Reference	B12152-01	Describes using OMB Plus to create, modify, delete, and retrieve object metadata in Warehouse Builder design and runtime repositories.
Oracle Warehouse Builder Java API Reference	B12155-01	Describes Oracle Warehouse Builder APIs.
Oracle Data Mining Administrator's Guide	B10697-01	Describes how to install the Oracle Data Mining (ODM) software and how to perform a other administrative functions common to all ODM administration.
Oracle Data Mining Concepts	B10698-01	Discusses the basic concepts underlying Oracle Data Mining.
Oracle Data Mining Application Developer's Guide	B10699-01	Describes using the Oracle Data Mining Java and PL/SQL APIs to perform data mining tasks for business applications, bioinformatics, and text mining.
Streams and Replication		
Oracle Streams Concepts and Administration	B10727-01	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 Streams Replication Administrator's Guide	B10728-01	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 Database Advanced Replication	B10732-01	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	B10733-01	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.

Table 2–1 (Cont.) Oracle Database 10g Server and SQL*Plus Documentation

Documentation	Part Number	Description
SQL*Plus		
*SQL*Plus User's Guide and Reference*	B12170-01	Provides introductory and how-to sections describing the SQL*Plus command-line interface, the Windows graphical user interface, and the web-based *i*SQL*Plus user interface and how to configure and use them. Reference and error sections provide SQL*Plus command descriptions and list SQL*Plus error messages with associated cause and action text.
*SQL*Plus Quick Reference*	B12171-01	Shows *i*SQL*Plus syntax, navigation, and buttons, and SQL*Plus command syntax.
*SQL*Plus Release Notes*	B12172-01	Provides important last minute information not included in the Oracle Database Documentation Library.
Data Guard		
Oracle Data Guard Broker	B10822-01	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 Data Guard Concepts and Administration	B10823-01	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.
Networking and Security		
Oracle Net Services Administrator's Guide	B10775-01	Explains how to plan, configure, and manage enterprise-wide connectivity with Oracle Net Services.
Oracle Net Services Reference Guide	B10776-01	Contains a complete listing and description of the control utility commands and configuration file parameters available for managing components of Oracle Net Services.
Oracle Security Overview	B10777-01	Introduces the basic concepts of system security. It outlines the data security risks which are prevalent today, and the industry-standard technologies available to address them. It then presents the suite of Oracle products you can use to implement these security technologies.

Table 2–1 (Cont.) Oracle Database 10g Server and SQL*Plus Documentation

Documentation	Part Number	Description
Oracle Database Security Guide	B10773-01	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 Advanced Security Administrator's Guide	B10772-01	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 Label Security Administrator's Guide	B10774-01	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 Real Application Clusters		
Oracle Real Application Clusters Installation and Configuration Guide	B10766-01	Explains Oracle Real Application Clusters (RAC) setup and configuration procedures as well as how to use Oracle tools for software installation and database creation. This RAC install book is new for Oracle Database10*g* and includes RAC installation information for all platforms.
Oracle Real Application Clusters Administrator's Guide	B10765-01	Explains how to administer Oracle Real Application Clusters.
Oracle Real Application Clusters Deployment and Performance Guide	B10768-01	Provides a high-level explanation of deployment practices for various system types such as e-commerce, data warehousing, and online transaction processing. Gives an in-depth examination of Oracle Real Application Clusters performance measurement and tuning methods to maximize clustered environment performance.
Application Development		
Oracle Database Application Developer's Guide - Fundamentals	B10795-01	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.

Table 2–1 (Cont.) Oracle Database 10g Server and SQL*Plus Documentation

Documentation	Part Number	Description
Oracle Database Application Developer's Guide - Large Objects	B10796-01	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 Database Application Developer's Guide - Object-Relational Features	B10799-01	Describes user-defined object datatypes and how to use these datatypes to model complex real-world entities as objects in the database.
Oracle Database Application Developer's Guide - Workspace Manager	B10824-01	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.
Oracle Database Application Developer's Guide - Expression Filter	B10821-01	Provides usage and reference information about how to manage, index, and evaluate conditional expressions in relational tables.
Oracle Streams Advanced Queuing User's Guide and Reference	B10785-01	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.
Oracle Data Cartridge Developer's Guide	B10800-01	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 XML DB Developer's Guide	B10790-01	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.

Table 2–1 (Cont.) Oracle Database 10g Server and SQL*Plus Documentation

Documentation	Part Number	Description
Oracle HTML DB User's Guide	B10992-01	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.
Oracle XML Developer's Kit Programmer's Guide	B10794-01	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.
Language and Interface		
Oracle Call Interface Programmer's Guide	B10779-01	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 C++ Call Interface Programmer's Guide	B10778-01	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.
PL/SQL User's Guide and Reference	B10807-01	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.
*Pro*C/C++ Programmer's Guide*	A97269-03	Describes how to develop C++ programs that use the SQL and PL/SQL database languages to access and manipulate Oracle data.
*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.

Table 2–1 (Cont.) Oracle Database 10g Server and SQL*Plus Documentation

Documentation	Part Number	Description
Programmer's Guide to the Oracle Precompilers	A42525-1	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.
*Pro*FORTRAN Supplement to the Oracle Precompilers Guide*	A42523-1	Describes how to write FORTRAN programs that use SQL to access and manipulate Oracle data.
*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 an Oracle database. A number of complete examples using Module Language, Ada code, and stored database procedures are provided.
*Pro*PL/1 Supplement to the Oracle Precompilers Guide*	A87540-03	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.
Application Reference		
PL/SQL Packages and Types Reference	B10802-01	Describes the PL/SQL packages and defined types supplied with the Oracle Database. Packages are listed alphabetically with syntax, procedures and functions, and parameters described for each package.
Oracle XML API Reference	B10789-01	Describes Oracle XML Developer's Kit (XDK) and Oracle XML DB APIs for developers building XML applications on the Oracle Database and provides syntax and a brief description of functions, methods, and procedures associated with them.
***inter*Media**		
Oracle interMedia User's Guide	B10840-01	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 interMedia Reference	B10829-01	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.
Oracle interMedia Java Classes Reference	B10830-01	Describes Java Classes to enable users to write Java applications using Oracle *inter*Media objects.

Table 2–1 (Cont.) Oracle Database 10g Server and SQL*Plus Documentation

Documentation	Part Number	Description
Oracle interMedia Java Classes API Reference	B12248-01	Oracle *inter*Media is a feature that enables the 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 Java Classes for Servlets and JSP API Reference	B12249-01	Describes how to use servlets and JSPs to facilitate the retrieval and uploading of multimedia data from and to an Oracle database.
Oracle interMedia Annotator User's Guide and Reference	B10831-01	Describes how to extract metadata from a multimedia file and storing both the metadata and the multimedia file in an Oracle database.
Oracle Text		
Oracle Text Application Developer's Guide	B10729-01	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	B10730-01	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.
Oracle Ultra Search User's Guide	B10731-01	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.
Oracle Workflow		
Oracle Workflow Administrator's Guide	B10283-02	Describes how to administer Oracle Workflow.
Oracle Workflow Developer's Guide	B10284-02	Describes how to develop applications using Oracle Workflow.
Oracle Workflow User's Guide	B10285-02	Describes how to use Oracle Workflow.
Oracle Workflow API Reference	B10286-02	Describes the APIs used by Oracle Workflow.
Oracle Workflow Installation Notes for Oracle Database	B12169-01	Describes information about installation of Oracle Workflow.

Oracle Data Mining Documentation

Table 2–2 Oracle Database 10g Data Mining Documentation

Documentation	Part Number	Description
Oracle Data Mining Concepts	B10698-01	Provides an overview of basic Oracle Data Mining concepts.
Oracle Data Mining Administrator's Guide	B10697-01	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 Application Developer's Guide	B10699-01	Describes how to use the Oracle Data Mining API (Java and PL/SQL) to perform data mining tasks, including building and testing models, computing lift, and scoring.

Oracle Spatial Documentation

Table 2–3 Oracle Spatial Documentation

Documentation	Part Number	Description
Oracle Spatial User's Guide and Reference	B10826-01	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 GeoRaster	B10827-01	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 Topology and Network Data Models	B10828-01	Provides usage and reference information about the topology data model and network data model capabilities of Oracle Spatial.

Oracle OLAP Documentation

Table 2–4 OLAP Documentation

Documentation	Part Number	Description
Oracle OLAP Application Developer's Guide	B10333-02	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	B10334-02	Explains the syntax of PL/SQL packages and types and the column structure of views related to Oracle OLAP.
Oracle OLAP DML Reference	B10339-02	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	B10335-02	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 an Oracle database. Describes the API and how to discover metadata, create queries, and retrieve data.
Oracle OLAP Java API Reference	B10994-01	Describes the classes and methods in the Oracle OLAP Java API for querying analytic workspaces and relational data warehouses.
Oracle OLAP Analytic Workspace Java API Reference	B12180-01	Describes the classes and methods in the Oracle OLAP Java API for building and maintaining analytic workspaces.

Oracle Database10*g* Documentation for Windows

Table 2–5 Oracle Windows Documentation

Documentation	Part Number	Description
Oracle Database Platform Guide for Windows	B10113-01	Provides platform-specific information for configuring and using your Oracle database software on Windows. Topics discussed include: architecture, services, tools, configuration, administration, tuning, backup, accounts, passwords, user authentication, registry parameters, and application development.
Oracle Database Installation Guide for Windows	B10130-01	Provides platform-specific preinstallation, installation, postinstallation, and component information for your single instance Oracle Database 10*g* database and client software on Windows.
Oracle Services for Microsoft Transaction Server Developer's Guide	B10114-01	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.
Oracle Provider for OLE DB Developer's Guide	B10115-01	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. This developer's guide describes provider-specific features and properties of OraOLEDB.
Oracle COM Automation Feature Developer's Guide	B10116-01	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 Data Provider for .NET Developer's Guide	B10117-01	Oracle Data Provider for .NET (ODP.NET) is an implementation of Microsoft's ADO.NET interface and provides access to Oracle databases. This developer's guide describes provider-specific features and properties of ODP.NET.
Oracle Objects for OLE Developer's Guide	B10118-01	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. This developer's guide describes the concepts important to OO4O and the object classes.
Oracle Objects for OLE C++ Class Library Developer's Guide	B10119-01	The Oracle Objects for OLE C++ Class Library is a collection of C++ classes that provide programmatic access to the OO4O automation server. This developer's guide describes the concepts and the object classes.

Table 2–5 (Cont.) Oracle Windows Documentation

Documentation	Part Number	Description
*Pro*C/C++ Getting Started for Windows*	A96111-03	This guide is an addendum to the Pro*C / C++ Programmer's Guide containing information specific to the Windows platform.
*Pro*COBOL Getting Started for Windows*	A96113-03	This guide is an addendum to the Pro*COBOL Programmer's Guide containing information specific to the Windows platform.
Oracle Client Installation Guide for Windows	B10131-01	Describes how to install Windows clients.
Oracle Database Release Notes for Windows	B10132-01	Provides important last minute information not included in the Oracle Database Documentation Library.
Oracle HTML DB Quick Installation Guide for Windows	B12215-01	Describes how to install HTML DB.

Java Documentation

Table 2–6 Java Documentation

Documentation	Part Number	Description
Oracle Database JDBC Developer's Guide and Reference	B10979-01	A discussion of how to use Oracle JDBC drivers, including Oracle extensions and APIs.
Oracle Database JDBC Release Notes	B12143-01	Provides important last minute information not included in the Oracle Database Documentation Library.
Oracle Database JPublisher User's Guide	B10983-01	The Oracle Database JPublisher User's Guide is for JDBC and J2EE programmers who want Java classes in their applications to correspond to database entities such as object types, VARRAY types, nested table types, OPAQUE types, PL / SQL packages, server-side Java classes, or SQL queries and DML statements. It is also for programmers who want to publish such server-side entities to Web services. The manual assumes that you are an experienced Java programmer with knowledge of Oracle databases, SQL, PL / SQL, and JDBC.
Oracle Database Java Developer's Guide	B12021-01	This book provides a general overview on how to develop, load, and execute your Java applications in the database.
Oracle XML Java API Reference	B12024-01	Describes the Java APIs for XML.

Table 2–6 (Cont.) Java Documentation

Documentation	Part Number	Description
Oracle Data Cartridge Java API Reference	B12025-01	Describes the Java APIs for the Data Cartridge.
Oracle Ultra Search Java API Reference	B12028-01	Describes Java APIs for Ultra Search.
Oracle Ultra Search Java API Reference	B12030-01	Describes additional Java APIs for Ultra Search.
Oracle9iAS Certificate Authority Java API Reference	B12179-01	Describes the Java APIs for Oracle9*i*AS Certificate Authority.
Oracle Streams Advanced Queuing Java API Reference	B12023-01	Describes the Java APIs for Streams Advanced Queuing.

Oracle Enterprise Manager Documentation

Table 2–7 Oracle Enterprise Manager Documentation

Documentation	Part Number	Description
Oracle Enterprise Manager Grid Control Installation and Basic Configuration	B12012-01	Describes how to install and get up and running quickly with the new Oracle Enterprise Manager 10*g* Grid Control framework, which provides you with central management capabilities for managing all the services within your enterprise, including hosts, databases, listeners, application servers, HTTP Servers, and Web applications.
Oracle Enterprise Manager Concepts	B12016-01	Provides a high-level overview of the features and capabilities of the Oracle Enterprise Manager 10*g* product family.
Oracle Enterprise Manager Advanced Configuration	B12013-01	Describes how to perform advanced configuration tasks so you can take full advantage of Oracle Enterprise Manager 10*g* features, deploy Oracle Enterprise Manager efficiently throughout your organization, and become familiar with the directory structure, log files, and configuration options of the Enterprise Manager product set.
Oracle Universal Installer Concepts Guide	B12140-01	Describes how to use the Oracle Universal Installer to install Oracle and third-party software.

Index

Numerics

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

Z