Oracle® Database

New Features Guide

11*g* Release 2 (11.2)

E17128-04

October 2010

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

ORACLE®

Oracle Database New Features Guide, 11g Release 2 (11.2)

E17128-04

Contents

2 Oracle Database 11*g* Release 2 (11.2.0.2) New Features

3   Oracle Database 11*g* Documentation

x

List of Tables

Preface

This Preface contains these topics:

- Audience
- Documentation Accessibility
- Related Documents
- Conventions

Audience

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

Documentation Accessibility

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

Accessibility of Code Examples in Documentation

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

Accessibility of Links to External Web Sites in Documentation

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

Deaf/Hard of Hearing Access to Oracle Support Services

To reach Oracle Support Services, use a telecommunications relay service (TRS) to call Oracle Support at 1.800.223.1711. An Oracle Support Services engineer will handle

technical issues and provide customer support according to the Oracle service request process. Information about TRS is available at http://www.fcc.gov/cgb/consumerfacts/trs.html, and a list of phone numbers is available at http://www.fcc.gov/cgb/dro/trsphonebk.html.

Related Documents

For more information, see the following documents in the Oracle Database 11*g* Release 2 (11.2) documentation set:

- *Oracle Database Error Messages*
- *Oracle Database Administrator's Guide*
- *Oracle Database Concepts*
- *Oracle Database Reference*

Conventions

The following text conventions are used in this document:

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

1

Oracle Database 11*g* Release 1 (11.2.0.1) New Features

This chapter contains descriptions of all of the features that are new to Oracle 11*g* Database, Release 2. This chapter contains the following sections:

- Application Development
- Availability
- Business Intelligence and Data Warehousing
- Clustering
- Database Overall
- Diagnosability
- Performance
- Security
- Server Manageability
- Unstructured Data Management

1.1 Application Development

The following sections describe the new application development features for Oracle Database 11*g* Release 2 (11.2).

1.1.1 Oracle Application Express

The following sections describe Oracle Application Express features.

1.1.1.1 Application Date Format

You can now define a date format to be used throughout an application. This date format is used to alter the `NLS_DATE_FORMAT` database session setting prior to showing or submitting any page within the application. This format is used by all reports showing dates and is also picked up by form items of type "Date Picker (use Application Date Format)".

The ability to specify a date format at the application level ensures consistency across the application. Therefore, whenever dates are displayed or input, they are in the same format.

1.1.1.2 Custom Themes

In addition to the default themes provided with Oracle Application Express, you can create your own customized themes. You can either start with one of the twenty standard themes available with Oracle Application Express and modify the underlying templates or define your own templates from scratch. Each theme consists of a set of templates defined with cascading style sheets (CSS) and HTML.

The ability to publish custom themes enables you to design a specific look and feel to meet your corporate requirements and then publish them as a theme for all other applications to use.

1.1.1.3 Declarative BLOB Support

Declarative BLOB support enables files to be declaratively uploaded in forms, and downloaded or displayed using reports. BLOB display and download can also be authored procedurally using PL/SQL.

The storing of binary large objects (BLOBs) within the database is growing in popularity due to the many advantages over storing content on disparate file systems. By incorporating declarative support for managing BLOBs into Application Express, the loading and manipulating of content is greatly simplified.

1.1.1.4 Documented JavaScript Libraries

This release includes an improved framework for advanced Oracle Application Express developers to build and leverage custom Web 2.0 capabilities, improving performance and enabling developers to create more dynamic application widgets. Oracle Application Express also includes the ability to suppress standard JavaScript and CSS files. All included JavaScript files are now compressed to improve page load time.

Many developers want to extend their applications to include additional Web 2.0 capabilities or to minimize the page weight for use on mobile devices such as iphones and smartphones. The documentation and declarative capabilities allow developers to design applications for these disparate requirements.

1.1.1.5 Enhanced Report Printing

Release 3.1 includes XML as a download format and supports multiple SQL statements.

Oracle Application Express interactive reporting provides the ability to manipulate the way in which the data is displayed on the screen. Users can also download this data in various formats including PDF, RTF, XLS and now XML.

1.1.1.6 Forms Conversion

Forms Conversion captures the design of existing Oracle Forms and automatically converts some components, primarily the user interface. Other components, such as complex triggers, need to be manually converted post-generation.

Moving to native HTML is not seamless and changes to the user interface are required to deliver optimal Web interactivity.

The Oracle Application Express Forms Conversion enables you to take advantage of Oracle Application Express dynamic HTML capabilities, including interactive reports. Given the similarities between Oracle Forms and Oracle Application Express development (both use SQL and PL/SQL), retraining requirements are also low.

1.1.1.7 Improved Security

Oracle Application Express offers a number of security enhancements. Key enhancements include the ability to declaratively encrypt session state and specify session time outs for maximum idle time and maximum session duration as well as create new password item types that enable users to enter passwords without ever saving them to session state.

Other features include reducing the privileges required by the Oracle Application Express database account, disabling database monitoring by default, and the ability to specify HTTPS for administration. In addition, administrators can now restrict password reuse. This release also includes a new Hidden and Protected item type. This item type greatly simplifies the developer's task of protecting item session state. This, together with other minor improvements, makes the default security functionality more robust within Oracle Application Express.

The additional declarative security capabilities make it easier for developers and administrators to harden the security of their applications and the development environment. These new capabilities complement existing Oracle Application Express security features some of which include flexible authentication, authorization schemes, and URL tampering protection.

1.1.1.8 Interactive Reporting Region

Interactive Reporting Regions enable end users to customize reports. Users can alter the layout of report data by choosing the columns they are interested in, applying filters, highlighting, and sorting. They can also define breaks, aggregations, different charts, and their own computations. Users can create multiple variations of the report and save them as named reports and download to various file formats including comma-delimited file (CSV) format, Microsoft Excel (XLS) format, Adobe Portable Document Format (PDF), and Microsoft Word Rich Text Format (RTF).

Oracle Application Express Interactive Reporting enables developers to quickly develop reports that can be manipulated by end users to meet a wide range of reporting requirements. Therefore, instead of developers having to define specific report layouts for different users or groups, they can define a common report that can be used to meet the majority of the different requirements.

1.1.1.9 Runtime-Only Installation

For testing and production instances, Oracle Application Express now supports the ability to install a runtime version of Oracle Application Express. This minimizes the installed footprint and privileges. Scripts are also provided to remove or add the developer interface from an existing instance.

The ability to implement a runtime-only environment improves application security as developers cannot inadvertently or maliciously update a production application.

1.1.2 Other General Development Features

The following sections describe new features in the areas of OCI, Pro*C, JDBC, and other development APIs.

1.1.2.1 Support WITH HOLD Option for CURSOR DECLARATION in Pro*C

The `WITH HOLD` option can now be specified during cursor declaration.

This new option provides easy migration of Pro*C applications.

See Also:

Pro*C/C++ Programmer's Guide for details

1.1.2.2 Pro*C Support for 8-Byte Native Numeric Host Variable for INSERT and FETCH

Oracle Call Interface (OCI) now provides Pro*C support for 8-byte native numeric host variable for `INSERT` and `FETCH` on 32-bit and 64-bit platforms.

Fusion applications need Pro*C to be able to support 8-byte native data type for bind/define while inserting or fetching data to and from a `NUMBER(18)` column.

See Also:

Pro*C/C++ Programmer's Guide for details

1.1.2.3 Pro*COBOL Support for 8-Byte Native Numeric Host Variable for INSERT and FETCH

Oracle Call Interface (OCI) now provides Pro*COBOL support for 8-byte native numeric host variable for `INSERT` and `FETCH` on 32-bit and 64-bit platforms.

Fusion applications need Pro*COBOL to be able to support 8-byte native data type for bind/define while inserting or fetching data to and from a `NUMBER(18)` column.

1.1.2.4 JDBC Support for Time Zone Patching

The JDBC driver is updated to conform with the new time zone upgrading scheme.

This feature provides a simplified time zone patching process. As a result, Java applications using the `TIMESTAMP WITH TIME ZONE` data type are immune to Daylight Saving Time (DST) changes.

1.1.2.5 JDBC Support for SecureFile Zero-Copy LOB I/O and LOB Prefetching

JDBC now supports SecureFile zero-copy LOB I/O and LOB prefetching.

This feature allows performant and secure Java access to structured (relational) and unstructured data.

See Also:

Oracle Database JDBC Developer's Guide for details

1.1.2.6 OCI Support for 8-Byte Integer Bind/Define

Oracle Call Interface (OCI) now provides support for 8-byte integer bind/define on 32-bit and 64-bit platforms.

Fusion applications need Pro*C or Pro*COBOL to be able to support 8-byte native data type for bind/define while inserting or fetching data to and from a `NUMBER(18)` column. Pro*C or Pro*COBOL need this support from OCI to be able to pass it on to application developers.

See Also:

Oracle Call Interface Programmer's Guide for details

1.2 Availability

The focus of this Availability section is aimed towards providing capabilities that keep the Oracle database available for continuous data access, despite unplanned failures and scheduled maintenance activities. These various capabilities form the basis of Oracle Maximum Availability Architecture (MAA), which is the Oracle blueprint for implementing a highly available infrastructure using integrated Oracle technologies.

1.2.1 Backup and Recovery

The following sections describe new features in this release that provide improvements in the area of backup and recovery.

1.2.1.1 Automatic Block Repair

Automatic block repair allows corrupt blocks on the primary database or physical standby database to be automatically repaired, as soon as they are detected, by transferring good blocks from the other destination. In addition, `RECOVER BLOCK` is enhanced to restore blocks from a physical standby database. The physical standby database must be in real-time query mode.

This feature reduces time when production data cannot be accessed, due to block corruption, by automatically repairing the corruptions as soon as they are detected in real-time using good blocks from a physical standby database. This reduces block recovery time by using up-to-date good blocks from a real-time, synchronized physical standby database as opposed to disk or tape backups or flashback logs.

> **See Also:**
>
> *Oracle Data Guard Concepts and Administration* for details

1.2.1.2 Backup to Amazon Simple Storage Service (S3) Using OSB Cloud Computing

Oracle now offers backup to Amazon S3, an internet-based storage service, with the Oracle Secure Backup (OSB) Cloud Module. This is part of the Oracle Cloud Computing offering.

This feature provides easy-to-manage, low cost database backup to Web services storage, reducing or eliminating the cost and time to manage an in-house backup infrastructure.

> **See Also:**
>
> *Oracle Database Backup and Recovery Reference* for details

1.2.1.3 DUPLICATE Without Connection to Target Database

`DUPLICATE` can be performed without connecting to a target database. This requires connecting to a catalog and auxiliary database.

The benefit is improved availability of a `DUPLICATE` operation by not requiring connection to a target database. This is particularly useful for `DUPLICATE` to a destination database where connection to the target database may not be available at all times.

> **See Also:**
>
> *Oracle Database Backup and Recovery User's Guide* for details

1.2.1.4 Enhanced Tablespace Point-In-Time Recovery (TSPITR)

Tablespace point-in-time recovery (TSPITR) is enhanced as follows:

- You now have the ability to recover a dropped tablespace.

- TSPITR can be repeated multiple times for the same tablespace. Previously, once a tablespace had been recovered to an earlier point-in-time, it could not be recovered to another earlier point-in-time.

- `DBMS_TTS.TRANSPORT_SET_CHECK` is automatically run to ensure that TSPITR is successful.

- `AUXNAME` is no longer used for recovery set data files.

This feature improves usability with TSPITR.

1.2.1.5 New DUPLICATE Options

The following are new options for the `DUPLICATE` command:

- `NOREDO`

  `NOREDO` indicates that archive logs are not applied. Because targetless DUPLICATE does not connect to the target database, it cannot check if the database is running in `NOARCHIVELOG` mode. It can also be used during regular duplication to force a database currently in `ARCHIVELOG` mode to be recovered without applying archive logs (for example, because it was in `NOARCHIVELOG` mode at the point-in-time it is being duplicated).

- `UNDO TABLESPACE <tsname> [ , <tsname> ... ]`

  When not connected to a recovery catalog and not connected to an open target database, RMAN cannot obtain the list of tablespaces with undo segments, therefore, you must specify them with this clause.

This feature improves the usability of the DUPLICATE command.

1.2.1.6 New SET NEWNAME Clauses and Format Options

The following are new clauses and format options for the `SET NEWNAME` command:

- A single `SET NEWNAME` command can be applied to all files in a tablespace, or for all files in the database. For example:

  ```
  SET NEWNAME FOR TABLESPACE <tsname> TO <format>;
  ```

  Or,

  ```
  SET NEWNAME FOR DATABASE TO <format>;
  ```

- New format identifiers for `SET NEWNAME...<format>` are as follows:

  – `%U`

    Unique identifier. `data_D-%d_I-%I_TS-%N_FNO-%f`

  – `%b`

    UNIX base name of the original data file name. For example, if the original data file name was `ORACLE_HOME/data/tbs_01.f`, then `%b` is `tbs_01.f`.

The benefit is improved flexibility of `RESTORE`, `DUPLICATE`, and TSPITR.

> **See Also:**
>
> *Oracle Database Backup and Recovery User's Guide* for details

1.2.1.7 Tablespace Checks in DUPLICATE

The `DUPLICATE...TABLESPACE` and `DUPLICATE... SKIP TABLESPACE` commands now perform the following initial checks:

- Excluded tablespaces are checked to see if they contain any objects owned by `SYS`.

- `DBMS_TTS.TRANSPORT_SET_CHECK` is run to ensure that the set of tablespaces being duplicated are self-contained before the actual duplicate process.

These checks are not possible for a targetless `DUPLICATE` as they are required to be run at the target database.

This feature improves usability of DUPLICATE. Any tablespace issues are immediately identified prior to commencement of the actual duplicate operation.

1.2.2 Online Application Maintenance and Upgrade

The following sections describe online application maintenance and upgrade features.

1.2.2.1 Edition-based Redefinition

Edition-based redefinition allows an application's database objects to be changed without interrupting the application's availability by making the changes in the privacy of a new edition. Every database has at least one edition. The DBA creates a new edition as a child of the existing one. The changes are made in the child edition while you continue to use the parent edition. When needed, changes to data are made safely by writing only to new columns or new tables not seen by the old edition. Editioning views expose a different projection of each changed table into each edition to allow each to see just its own columns. Crossedition triggers propagate data changes made by the old edition into the columns of the new edition. When the installation of the changes is complete, some users start to use the new edition while others drain off the old edition. Here, crossedition triggers propagate data changes made by the new edition into the columns of the old edition.

Large, mission critical applications are often unavailable for long periods of time while database objects are patched or upgraded. Edition-based redefinition allows this cost to be avoided.

> **See Also:**
>
> *Oracle Database Advanced Application Developer's Guide* for details

1.2.2.2 Enhance CREATE or REPLACE TYPE to Allow FORCE

The `FORCE` option can now be used in conjunction with the `CREATE` or `REPLACE TYPE` command.

This feature provides enhanced usability and allows a `CREATE` or `REPLACE TYPE` operation to be performed even when `TYPE` dependent objects are present. However, if at least one `TABLE` dependent is present, then `FORCE` does not allow `CREATE` or `REPLACE TYPE` to succeed.

> **See Also:**
>
> *Oracle Database Object-Relational Developer's Guide* for details

1.2.2.3 Fine-Grained Dependencies for Triggers

Oracle Database 11*g* Release 1 (11.1) brought both fine-grained dependency tracking and the new possibility that a trigger might be a dependency parent by virtue of the new `FOLLOWS` keyword.

In release 11.1, dependents on triggers did not have fine-grained dependency. In release 11.2, this fine-grained dependence exists. (Release 11.2 also provides the new PRECEDES keyword which also allows trigger-upon-trigger dependencies.)

See Also:

Oracle Database Advanced Application Developer's Guide for details

1.2.2.4 IGNORE_ROW_ON_DUPKEY_INDEX Hint for INSERT Statement

With INSERT INTO TARGET...SELECT...FROM SOURCE, a unique key for some to-be-inserted rows may collide with existing rows. The IGNORE_ROW_ON_DUPKEY_ INDEX allows the collisions to be silently ignored and the non-colliding rows to be inserted. A PL/SQL program could achieve the same effect by first selecting the source rows and by then inserting them one-by-one into the target in a block that has a null handler for the DUP_VAL_ON_INDEX exception. However, the PL/SQL approach would take effort to program and is much slower than the single SQL statement that this hint allows.

This hint improves performance and ease-of-programming when implementing an online application upgrade script using edition-based redefinition.

See Also:

Oracle Database SQL Language Reference for details

1.2.3 Oracle Data Guard

The following sections describe new features in this release that provide improvements in Oracle Data Guard.

1.2.3.1 Compressed Table Support in Logical Standby Databases and Oracle LogMiner

Compressed tables (that is, tables with compression that support both OLTP and direct load operations) are supported in logical standby databases and Oracle LogMiner.

With support for this additional storage attribute, logical standby databases can now provide data protection and reporting benefits for a wider range of tables.

See Also:

Oracle Data Guard Concepts and Administration for details

1.2.3.2 Configurable Real-Time Query Apply Lag Limit

A physical standby database can be open for read-only access while redo apply is active only if the Oracle Active Data Guard option is enabled. This capability is known as real-time query.

The new STANDBY_MAX_DATA_DELAY session parameter can be used to specify a session-specific apply lag tolerance, measured in seconds, for queries issued by non-administrative users to a physical standby database that is in real-time query mode.

This capability allows queries to be safely offloaded from the primary database to a physical standby database, because it is possible to detect if the standby database has become unacceptably stale.

> **See Also:**
>
> *Oracle Data Guard Concepts and Administration* for details

1.2.3.3 Integrated Support for Application Failover in a Data Guard Configuration

Applications connected to a primary database can transparently failover to the new primary database upon an Oracle Data Guard role transition. Integration with Fast Application Notification (FAN) provides fast failover for integrated clients.

Flexibility and manageability of disaster recovery configurations using Oracle Data Guard is improved.

> **See Also:**
>
> *Oracle Data Guard Broker* for details

1.2.3.4 Support Up to 30 Standby Databases

The number of standby databases that a primary database can support is increased from 9 to 30 in this release.

The capability to create 30 standby databases, combined with the functionality of the Oracle Active Data Guard option, allows the creation of reader farms that can be used to offload large scale read-only workloads from a production database.

> **See Also:**
>
> *Oracle Data Guard Concepts and Administration* for details

1.3 Business Intelligence and Data Warehousing

The following sections describe new Business Intelligence and Data Warehousing features for Oracle Database 11*g* Release 2 (11.2).

1.3.1 Improved Analytics

The following sections describe new and improved analytical capabilities in this release.

1.3.1.1 Analytic Functions 2.0

New and enhanced analytical functions are introduced in this release. A new ordered aggregate, `LISTAGG`, concatenates the values of the measure column. The new analytic window function `NTH_VALUE` (a generalization of existing `FIRST_VALUE` and `LAST_VALUE` functions) gives users the functionality of retrieving an arbitrary (or nth) record in a window.

The `LAG` and `LEAD` functions are enhanced with the `IGNORE NULLS` option.

The new and enhanced SQL analytical functions allow more complex analysis in the database, using (simpler) SQL specification and providing better performance.

> **See Also:**
>
> *Oracle Database Data Warehousing Guide* for details

1.3.1.2 Recursive WITH Clause

The SQL `WITH` clause has been extended to enable formulation of recursive queries.

Recursive `WITH` clause complies with the American National Standards Institute (ANSI) standard. This makes Oracle ANSI-compatible for recursive queries.

See Also:

Oracle Database SQL Language Reference for details

1.3.2 Improved Data Loading

The following sections describe new and improved data loading capabilities in this release.

1.3.2.1 EXECUTE Privilege for DIRECTORY Objects

`EXECUTE` privilege is allowed for `DIRECTORY` objects in this release. The `ORACLE_ LOADER` access driver creates a process that runs a user-specified program. That program must live in a directory path specified by a directory object defined in the database. Only a user that has been given EXECUTE access to the directory object is allowed to run programs in it.

This feature allows the DBA to control who is allowed to run preprocessors as part of loading data with external tables. It also allows the DBA to restrict which programs those users can run. No existing users with access to the directory object are allowed to run any programs from that directory unless the DBA gives them EXECUTE access to that directory.

See Also:

Oracle Database Utilities for details

1.3.2.2 Preprocessing Data for ORACLE_LOADER Access Driver in External Tables

The syntax for the `ORACLE_LOADER` access driver is extended in this release to allow specification of a program to preprocess the data files that are read for the external table. The access parameters can specify the name of a directory object and the name of an executable file in that directory object. When the access driver needs to read data from a file, it creates a process that runs the specified program, passing in the name of the data file. The output from the program is passed into the access driver which parses the data into records and columns.

The initial use of this feature is by a customer who needs to load data that is stored in compressed files. The user specifies the name of the program used to decompress the file as part of the access parameters. The access driver reads the output of the decompression program.

Large customers want to load data from compressed files which requires less disk space and uses the I/O bandwidth between the disk and memory more efficiently.

See Also:

Oracle Database Utilities for details

1.3.3 Improved Partitioning

The following sections describe new and improved partitioning capabilities in this release.

1.3.3.1 Allow Virtual Columns in the Primary Key or Foreign Key for Reference Partitioning

Virtual columns can be used as the primary or the foreign key column of a reference partition table.

Allowing the use of virtual columns for reference partitioned tables enables an easier implementation of various business scenarios using Oracle Partitioning.

1.3.3.2 System-Managed Indexes for List Partitioning

System-managed domain indexes are now supported for list partitioned tables.

This feature provides enhanced completeness of domain-specific indexing support for partitioning to meet user requirements including Oracle XML DB. Performance of local domain indexes on list partitioned tables is improved in this release.

See Also:

Oracle Database SQL Language Reference for details

1.3.4 Improved Performance and Scalability

The following sections describe new and improved performance and scalability capabilities in this release.

1.3.4.1 In-Memory Parallel Execution

Traditionally, parallel execution has enabled organizations to manage and access large amounts of data by taking full advantage of the I/O capacity of the system. In-memory parallel execution harnesses the aggregated memory in a system to enhance query performance by minimizing or even completely eliminating the physical I/O needed for a parallel operation. Oracle automatically decides if an object being accessed using parallel execution benefits from being cached in the SGA (buffer cache). The decision to cache an object is based on a well defined set of heuristics including size of the object and the frequency that it is accessed. In an Oracle RAC environment, Oracle maps fragments of the object into each of the buffer caches on the active instances. By creating this mapping, Oracle knows which buffer cache to access to find a specific part or partition of an object to answer a given SQL query.

In-memory parallel query harnesses the aggregated memory in a system for parallel operations, enabling it to scale out with the available memory for data caching as the number of nodes in a cluster increases. This new functionality optimizes large parallel operations by minimizing or even completely eliminating the physical I/O needed because the parallel operation can now be satisfied in memory.

See Also:

Oracle Database VLDB and Partitioning Guide for details

1.3.4.2 Minimal Effort Parallel Execution - Auto Degree of Parallelism (DOP) and Queuing

When activated, Oracle determines the optimal degree of parallelism (DOP) for any given SQL operation based on the size of the objects, the complexity of a statement, and the existing hardware resources.

The database compensates for wrong or missing user settings for parallel execution, ensuring a more optimal resource consumption and overall system behavior.

See Also:

Oracle Database VLDB and Partitioning Guide for details

1.3.4.3 The DBMS_PARALLEL_EXECUTE Package

The `DBMS_PARALLEL_EXECUTE` package provides subprograms to allow a specified `INSERT`, `UPDATE`, `DELETE`, `MERGE`, or anonymous block statement to be applied in parallel chunks. The statement must have two placeholders that define the start and end limit of a chunk. Typically, these are values for the rowid or a surrogate unique key in a large table. But, when an anonymous block is used, the block can interpret the values arbitrarily. The package has subprograms to define ranges that cover the specified table. These include rule-based division of a table's rowid or key range and support user-defined methods. The SQL statement together with the set of chunk ranges define a task. Another subprogram starts the task. Each task is processed using a scheduler job and automatically commits when it completes. Progress is logged. Untried, successful, and failed chunks are flagged as such on task completion or interruption. Another subprogram allows the task to resume to try untried and failed chunks.

Many scenarios require the bulk transformation of a large number of rows. Using an ordinary SQL statement suffers from the all-or-nothing effect. In the common case, where the transformation of one row is independent of that of other rows, it is correct to commit every row that is transformed successfully and to roll back every row where the transformation fails. Some customers have implemented schemes to achieve this from scratch, using the Oracle Scheduler and suitable methods to record progress. This package provides a supported solution and adds database-wide manageability through new catalog views for parallel task metadata. The package is especially useful in online application upgrade scenarios to apply a crossedition trigger to all the rows in the table on which it is defined.

See Also:

Oracle Database PL/SQL Packages and Types Reference for details

1.3.4.4 Significant Performance Improvement of On-Commit Fast Refresh

Fast refresh of a materialized view is now significantly faster due to reducing the time spent on log handling.

This provides significantly reduced maintenance time and more fast refreshes are possible.

See Also:

Oracle Database Data Warehousing Guide for details

1.3.5 Oracle Warehouse Builder

The following sections describe improvements to the extraction, transformation, and loading (ETL) capabilities available with Oracle Warehouse Builder (OWB).

1.3.5.1 Advanced Find Support in Mapping Editor

The mapping editor has been enhanced with advanced find capabilities to make it easier to locate and make updates to operators, groups, and attributes in a mapping diagram, in the Available Objects tab, and in the Selected Objects tab.

This feature enhances ETL mapping developer productivity, especially on large and complex mappings and, for example, when working with complex data sources with large numbers of tables, views, or columns.

> **See Also:**
>
> *Oracle Warehouse Builder Data Modeling, ETL, and Data Quality Guide* for details

1.3.5.2 Business Intelligence Tool Integration

Oracle Warehouse Builder (OWB) now offers metadata integration with Oracle Business Intelligence Standard Edition (Discoverer) as well as Oracle Business Intelligence Enterprise Edition.

For Oracle Business Intelligence Enterprise Edition (OBI EE), this feature allows derivation of ready-to-use physical, business model and presentation layer metadata from a data warehouse design, visualization and maintenance of the derived objects from within OWB, and deployment of the derived objects in the form of an RPD file that can be loaded into OBI EE.

Oracle Discoverer integration was added in a previous release, and includes derivation of metadata for Discoverer from the data warehouse design, and deploying those derived objects into Discoverer. In this release, similar capabilities are now available for OBI Enterprise Edition. All business intelligence application objects are modeled in OWB and can be included in lineage and impact analysis at the column level.

Customers using Oracle business intelligence tools with their Oracle data warehouses can get better answers from their warehouses faster, with no additional design or development effort.

1.3.5.3 Copy and Paste of Operators and Attributes in Mapping Editor

In the mapping editor, users can now copy and paste operators within a mapping or across mappings, including attribute settings.

This enhancement saves time and reduces errors in the development of complex ETL mappings that reuse common or similar elements.

> **See Also:**
>
> *Oracle Warehouse Builder Data Modeling, ETL, and Data Quality Guide* for details

1.3.5.4 Current Configuration Dropdown List in Design Center Toolbar

In the Design Center, there is now a dropdown list that displays the active configuration of the user.

This feature improves usability of the multi-configuration feature.

> **See Also:**
>
> *Oracle Warehouse Builder Installation and Administration Guide for Windows and Linux* for details

1.3.5.5 Enhanced Support for Flat File Imports

There are numerous support improvements for importing flat files, including a simplified Flat File Sampling wizard, support for multi-character and hexadecimal

format delimiters and enclosures, simplified support for fixed format fields, and support for bulk flat file loads into heterogeneous targets.

Flat files are frequently used for simple and high performance data movement in ETL applications. These changes improve ETL developer productivity and provide flexible handling of flat files in more scenarios.

1.3.5.6 Enhanced Table Function Support

OWB now has improved support for table functions, including importing metadata for existing table functions, an editor for creating table functions from within OWB, and better support for table functions in mappings.

Improved support simplifies using table functions for much more flexible and powerful transformations, such as user-defined aggregations and data mining sampling operators.

1.3.5.7 Experts Available in Editor Menu

It is now possible to add OWB experts to the mapping editor menu.

This feature makes it possible to enhance and extend the functionality of the mapping editor, improving developer productivity.

1.3.5.8 Expression Editing in Operator Edit Dialog

Expressions associated with operator attributes can now be entered directly into an Operator Edit Dialog or Expression Editor, rather than requiring that these expressions be entered into a property in the Property Inspector.

Developers can finish more of their work in one place when creating operators in ETL mappings, thus improving their productivity.

> **See Also:**
>
> *Oracle Warehouse Builder Data Modeling, ETL, and Data Quality Guide* for details

1.3.5.9 Grouping and Spotlighting of Objects in Mapping Editor

In the mapping editor, users can now temporarily or permanently group objects in the mapping editor so that they are collapsed to a single icon. This hides complexity in mappings. Users can also spotlight a single operator, which temporarily hides all objects in the mapping except for those objects that connect directly to the operator.

These features improve productivity for developers working with complex mappings with large numbers of operators.

> **See Also:**
>
> *Oracle Warehouse Builder Data Modeling, ETL, and Data Quality Guide* for details

1.3.5.10 Improved Management of Locations Registered in Multiple Control Centers

The user interface for managing the registration of locations in control centers has been reworked to improve usability, especially when working with locations registered in multiple control centers.

This change improves productivity of OWB administrators responsible for managing locations across control centers.

1.3.5.11 Improved User Interface for Managing Locations

The user interface for managing OWB locations has been reworked to improve usability and support access to non-Oracle data sources using newly supported connectivity methods.

These changes improve Oracle Warehouse Builder administrator and developer productivity in heterogeneous and Oracle-only environments.

1.3.5.12 Key Lookup Operator Enhancements

Extensive changes have been made to the key lookup operator:

- More efficient use of screen real estate.

- Support for non-equality lookups.

- Dynamic lookups, where the lookup table may be modified during the mapping execution.

These changes make the lookup operator more powerful in many situations, including improving Type 2 slowly changing dimension support.

1.3.5.13 Mapping Debugger Enhancements

There are numerous enhancements to the OWB mapping editor, including:

- Improved support for watch points and enabling and disabling individual break points.

- Support for user-defined type columns.

- Enhanced support for numerous existing operators, such as `VARRAY`, `EXPAND`, and `CONSTRUCT`.

- Support for key lookup and table function operators.

- Support for correlated joins.

- Improved cleanup of debugger-specific objects.

These enhancements improve productivity for ETL mapping developers, especially when working with complex mappings where the mapping debugger adds the most value.

1.3.5.14 New JDeveloper-Style User Interface

The Oracle Warehouse Builder Design Center user interface has been updated to use the Fusion Client Platform, the same core Integrated Development Environment (IDE) platform as Oracle JDeveloper and Oracle SQL Developer.

The advantages of this user interface include:

- More efficient and flexible use of screen real estate.

- Support for opening multiple editors of the same type, for example, editing several ETL mappings at once in different windows.

- More consistent behavior across different parts of the OWB user interface.

This change brings Oracle Warehouse Builder Design Center in line with other development tools from Oracle. Developers experience increased productivity in the

Oracle Warehouse Builder environment, which now benefits from the usability research behind the Fusion Client Platform and consistency with other Oracle products.

> **See Also:**
>
> *Oracle Warehouse Builder Sources and Targets Guide* for details

1.3.5.15 Operator References Included in Generated PL/SQL Code

PL/SQL code generated for OWB ETL mappings now includes detailed comments to help developers associate specific operators in a mapping with sections of the generated code.

Developers can more easily troubleshoot issues with OWB-generated code that can only be detected when the code is deployed. This additional information enhances developer productivity.

1.3.5.16 Quick Mapper

In this release, Oracle Warehouse Builder (OWB) introduces a new spreadsheet-like dialog for connecting operators in a mapping. This functionality replaces the existing auto mapping dialog.

This improvement saves developer time and reduces errors when working with operators with a large number of inputs or outputs.

> **See Also:**
>
> *Oracle Warehouse Builder Data Modeling, ETL, and Data Quality Guide* for details

1.3.5.17 Repository Browser Changes

The Repository Browser has been updated to support foldering, expose the new types of metadata associated with the release 11.2 feature set, and support OC4J 10.3.3.

These changes improve Oracle Warehouse Builder manageability.

1.3.5.18 Simplified Oracle Warehouse Builder Repository Upgrades

The repository upgrade automatically upgrades an Oracle Warehouse Builder (OWB) repository to the current release with less user intervention.

This feature simplifies the task of upgrading from one release to the next.

1.3.5.19 Support for Extracting Data From Tables Containing LONG Data Type

Oracle Warehouse Builder can now generate SQL*Plus code to extract data from database schemas supporting the deprecated `LONG` data type, such as occurs in PeopleSoft application data sources.

Support for `LONG` data types used in PeopleSoft data enables OWB users to integrate more effectively with PeopleSoft data or any other data source that uses the `LONG` data type.

1.3.5.20 Support for Subqueries in Join Operator

The join operator in Oracle Warehouse Builder (OWB) now supports several new behaviors related to the use of subqueries in joins:

- Specifying subqueries using `EXISTS`, `NOT EXISTS`, `IN`, and `NOT IN`.

- Specifying outer joins using the input role instead of the + (plus) sign.

- Generating ANSI SQL syntax for all join types instead of only outer joins.

More flexible handling for join operations improves developer productivity and makes possible more flexible data transformations.

1.4 Clustering

The following sections describe new clustering features for Oracle Database 11*g* Release 2 (11.2).

1.4.1 Oracle Real Application Clusters Ease-of-Use

This release of Oracle Real Application Clusters (Oracle RAC) provides many features to dramatically simplify installation and on-going management of a cluster and Oracle RAC database, making it easy for the novice to adopt clustering and Oracle RAC and reap the benefits of this technology.

The following sections describe ease-of-use features for Oracle RAC.

1.4.1.1 Configuration Assistants Support New Oracle RAC Features

Database Configuration Assistant (DBCA), Database Upgrade Assistant (DBUA), and Net Configuration Assistant (NETCA) have been updated to support all of the new features of this release and provide a best practice implementation.

Configuration Assistants automate the configuration of the environment ensuring the correct steps are taken. The assistants simplify the implementation of clusters and clustered databases.

> **See Also:**
>
> *Oracle Real Application Clusters Installation Guide for Linux and UNIX* for details

1.4.1.2 Enhanced Cluster Verification Utility

Additional functionality has been added to the Cluster Verification Utility (CVU) in regard to checking certain storage types and configurations. Furthermore, it gives more consideration to user-specific settings.

These enhancements provide easier implementation and configuration of cluster environments and improved problem diagnostics in a cluster environment.

> **See Also:**
>
> *Oracle Real Application Clusters Installation Guide for Linux and UNIX* for details

1.4.1.3 Integration of Cluster Verification Utility and Oracle Universal Installer

The Cluster Verification Utility (CVU) is now fully integrated with the installer so that checks are done automatically for all nodes included in the installation.

This integration improves Oracle RAC manageability and deployment by ensuring that any problems with cluster setup are detected and corrected prior to installing Oracle software.

See Also:

Oracle Grid Infrastructure Installation Guide for Linux for details

1.4.1.4 Cluster Time Service

The Cluster Time Service synchronizes the system time on all nodes in the cluster. A synchronized system time across the cluster is a prerequisite to install and successfully run an Oracle cluster.

This feature simplifies management, maintenance, and support of an Oracle cluster and an Oracle RAC environment by providing an out-of-the-box time server. It also improves the reliability of Oracle RAC environments.

See Also:

Oracle Clusterware Administration and Deployment Guide for details

1.4.1.5 Oracle Cluster Registry (OCR) Enhancements

There have been improvements in this release in the way the Oracle Cluster Registry (OCR) is accessed. These improvements include:

- Faster relocation of services on node failure.
- Support for up to 5 copies of the OCR for improved availability of the cluster.
- Storage of OCR in Automatic Storage Management (ASM).

The tools to manage the OCR have changed to support the new management options.

These enhancements improve performance in Oracle Clusterware and Oracle Real Application Clusters environments and provide easier management of the cluster through consistent storage management automation

See Also:

Oracle Clusterware Administration and Deployment Guide for details

1.4.1.6 Grid Plug and Play (GPnP)

Grid Plug and Play (GPnP) eliminates per-node configuration data and the need for explicit add and delete nodes steps. This allows a system administrator to take a template system image and run it on a new node with no further configuration. This removes many manual operations, reduces the opportunity for errors, and encourages configurations that can be changed easily. Removal of the per-node configuration makes the nodes easier to replace, because they do not need to contain individually-managed state.

Grid Plug and Play reduces the cost of installing, configuring, and managing database nodes by making their per-node state disposable. It allows nodes to be easily replaced with regenerated state.

See Also:

Oracle Grid Infrastructure Installation Guide for details

1.4.1.7 Oracle Restart

Oracle Restart improves the availability of your single-instance Oracle database. Oracle Restart automatically restarts the database instance, the Automatic Storage Management (ASM) instance, the listener, and other components after a hardware or software failure or whenever your database host computer restarts. Server Control

(SRVCTL) is the command line interface to manage Oracle processes that are managed by Oracle Restart on a standalone server.

This feature provides improved reliability and automated management of a single-instance Oracle database and the management of any process or application running on the database server.

See Also:

Oracle Database Administrator's Guide for details

1.4.1.8 Policy-Based Cluster and Capacity Management

Oracle Clusterware allocates and reassigns capacity based on policies defined by you. This enables faster resource failover and dynamic capacity assignment using a policy-based management.

Policy-Based Cluster and Capacity Management allows the efficient allocation of all kinds of applications in the cluster. Various applications can be hosted on a shared infrastructure being isolated regarding their resource consumption by policies and, therefore, behave as if they were deployed in single system environments.

See Also:

Oracle Clusterware Administration and Deployment Guide for details

1.4.1.9 Improved Clusterware Resource Modeling

In this release, there are now more options for managing all types of applications and creating dependencies among them using Oracle Clusterware.

Improved Clusterware Resource Modeling enables a granular definition of dependencies among applications or processes to manage them as one entity.

See Also:

Oracle Clusterware Administration and Deployment Guide for details

1.4.1.10 Role-Separated Management

Role-separated management for Oracle Clusterware allows certain administrative tasks to be delegated to different people, representing different roles in the company. It is based on the idea of a clusterware administrator. The administrator may grant administrative tasks on a per resource basis. For example, if two databases are placed into the same cluster, the clusterware administrator can manage both databases in the cluster. But, the clusterware administrator may decide to grant different administrative privileges to each DBA responsible for one of those databases.

Role-separated management allows multiple applications and databases to share the same cluster and hardware resources, but ensures that different administration groups do not interfere with each other.

See Also:

Oracle Clusterware Administration and Deployment Guide for details

1.4.1.11 Agent Development Framework

Oracle Clusterware provides an agent framework for managing all kinds of applications with Oracle Clusterware. Using the agent framework provides optimized application startup, checking, and stopping based on user-defined scripts.

Making it easy to protect applications with Oracle Clusterware reduces costs allowing you to efficiently enable high availability for applications.

See Also:

Oracle Clusterware Administration and Deployment Guide for details

1.4.1.12 Zero Downtime Patching for Oracle Clusterware and Oracle RAC

The patching of Oracle Clusterware and Oracle Real Application Clusters can now be completed without taking the entire cluster down. Patchsets are now installed as out-of-place upgrades to the Oracle Grid infrastructure for a cluster software (Oracle Clusterware and Automatic Storage Management) and Oracle Database.

Now you can reduce your unplanned downtime of clustered databases and applications running in a cluster.

1.4.1.13 Enterprise Manager-Based Clusterware Resource Management

New in this release is an Enterprise Manager graphical user interface (GUI) to manage various Oracle Clusterware resources with full lifecycle support. In addition to allowing the creation and configuration of resources within Oracle Clusterware, it also helps to monitor and manage resources once deployed in the cluster.

Using Oracle Enterprise Manager as a GUI to monitor and manage various Oracle Clusterware resources eases the daily management in high availability environments.

See Also:

Oracle Clusterware Administration and Deployment Guide for details

1.4.1.14 Enterprise Manager Provisioning for Oracle Clusterware and Oracle Real Application Clusters

Enterprise Manager provisioning introduces procedures to easily scale up or scale down Oracle Clusterware and Oracle Real Application Clusters.

Ease-of-implementation and management for a clustered database environment can be achieved through utilizing the Enterprise Manager provisioning framework.

See Also:

Oracle Grid Infrastructure Installation Guide for Linux for details

1.4.1.15 Enterprise Manager Support for Grid Plug and Play

Oracle Enterprise Manager, the graphical user interface (GUI) for managing Oracle RAC, provides management and monitoring for the Grid Plug and Play environment.

Enterprise Manager is the standard GUI interface for Oracle Database. This integration provides an easy-to-use interface that customers are familiar with to manage Grid Plug and Play environments.

See Also:

Oracle Database 2 Day + Real Application Clusters Guide for details

1.4.1.16 Enterprise Manager Support for Oracle Restart

Enterprise Manager provides support for Oracle Restart and the configuration with single-instance databases. This is a change in configuration, monitoring, and administration to enable Oracle Restart.

Enterprise Manager provides a graphical user interface (GUI) interface to easily manage Oracle databases. This additional functionality enables you to restart your Oracle databases.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for details

1.4.1.17 Configuration Assistant Support for Removing Oracle RAC Installations

Database Configuration Assistant (DBCA), Database Upgrade Assistant (DBUA), and Net Configuration Assistant (NETCA) have been updated to support the complete deinstallation and deconfiguration of Oracle RAC databases and listeners.

This support improves the manageability of an Oracle RAC environment through automation of deinstallation and deconfiguration of Oracle RAC databases.

> **See Also:**
>
> *Oracle Real Application Clusters Installation Guide for Linux and UNIX* for details

1.4.1.18 Oracle Universal Installer Support for Removing Oracle RAC Installations

The installer can clean up a failed Oracle Clusterware installation or upgrade of an environment prior to reattempting the operation. This ensures that the reattempted operation is done over a clean environment, thereby eliminating the chances of errors related to environmental inconsistencies.

Easily cleaning up an environment provides improved Oracle RAC manageability and deployment.

> **See Also:**
>
> *Oracle Grid Infrastructure Installation Guide for Linux* for details

1.4.1.19 Improved Deinstallation Support With Oracle Universal Installer

The installation of Oracle Clusterware and Oracle RAC now have recovery points. If a failure occurs during installation, you can rollback to the closest recovery point and restart the installation once the problem has been corrected.

Installation rollback and recovery make the installation and configuration of Oracle Clusterware and Oracle RAC easier. It reduces project time lines by making it easy to recover from installation failures.

> **See Also:**
>
> *Oracle Real Application Clusters Installation Guide for Linux and UNIX* for details

1.4.1.20 Downgrading Database Configured With DBControl

Scripts are included to support DBControl downgrade as part of database downgrade.

If an upgrade is deemed unsuccessful, the system needs to be returned to the starting release. In order to maintain the reliability of management when modifying software releases, DBControl must be at the same release as the database that it is monitoring.

> **See Also:**
>
> *Oracle Database Upgrade Guide* for details

1.4.1.21 Oracle Restart Integration with Oracle Universal Installer

Oracle Restart requires a separate installation from Oracle Database. This installation is the Oracle Grid infrastructure for a cluster installation for standalone servers which includes Oracle Restart and Oracle Automatic Storage Management (ASM). This allows separation of roles such that the system administrator can manage the infrastructure and the database administrator can manage the database.

Oracle Universal Installer is the tool to install Oracle software. This improves the manageability of the Oracle environment on a standalone server allowing separation of roles and improved resiliency of the Oracle software.

1.4.1.22 Out-of-Place Oracle Clusterware Upgrade

A new version of Oracle Clusterware is now installed into a separate home from the current installation. This reduces the downtime required to upgrade a node in the cluster and facilitate the provisioning of clusters within an enterprise.

The benefit is a reduction in planned outage time required for cluster upgrades which assists in meeting availability service levels. This also makes it easier to provide a standard installation across the enterprise.

> **See Also:**
>
> *Oracle Grid Infrastructure Installation Guide for Linux* for details

1.4.1.23 OUI Support for Out-of-Place Oracle Clusterware Upgrade

You can now perform out-of-place upgrade of Oracle Clusterware software. The new version can be installed in a separate directory and pointed to during deployment.

Out-of-place upgrades provide easier Oracle RAC and grid deployment and manageability, as well as better testing for controlled application migration.

> **See Also:**
>
> *Oracle Grid Infrastructure Installation Guide for Linux* for details

1.4.1.24 Server Control (SRVCTL) Enhancements

The server control (SRVCTL) commands have been enhanced to manage the configuration in a standalone server with Oracle Restart as well as the new style of cluster management (Policy-Based Cluster and Capacity Management).

This feature provides easier management of Oracle Database through a consistent interface which can be used from the console or scripted.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for details

1.4.1.25 Server Control (SRVCTL) Enhancements to Support Grid Plug and Play

The command-line interface (CLI) for Oracle Clusterware and Oracle Real Application Clusters has been updated to support the new features of this release.

The CLI provides the ability to manage the cluster using a command line from a single point in the cluster and allows you to manage the cluster as a single entity. This reduces the management complexity for clusters and clustered databases. All changes to the cluster must be reflected in the management tool.

See Also:

Oracle Real Application Clusters Administration and Deployment Guide for details

1.4.1.26 SRVCTL Support for Single-Instance Database in a Cluster

Using SRVCTL, you can register a single-instance database to be managed by Oracle Clusterware. Once registered, Oracle Clusterware starts, stops, monitors, and restarts the database instance.

This feature provides an improved management interface which makes it easy to provide higher availability for single-instance databases that are running on a server that is part of a cluster.

See Also:

Oracle Database Administrator's Guide for details

1.4.1.27 Universal Connection Pool (UCP) Integration with Oracle Data Guard

In this release, Java applications that use the Oracle Universal Connection Pool (UCP) for Java now have fast connection failover when the primary site fails. When Data Guard fails over or switches over to the standby database site, the connection pool cleans up connections to the primary site, terminates active transactions, and creates connections to the standby database.

This feature provides increased availability for Java applications using UCP with Oracle RAC and Oracle Data Guard. Applications can easily mask failures to the end user.

1.4.1.28 UCP Integration With Oracle Real Application Clusters

Universal Connection Pool (UCP) is the new Java connection pool. It has many features that make it easy for Java applications to manage connections to an Oracle Real Application Clusters database such as Web Session Affinity, XA Affinity, Runtime Connection Load Balancing, and Fast Connection Failover.

This feature provides a robust connection pool for Java applications with improved throughput and fast failover in an Oracle Real Application Clusters environment.

1.4.1.29 Universal Connection Pool (UCP) for JDBC

Universal Connection Pool for JDBC supersedes Implicit Connection Cache and provides the following functions:

- Connection labeling, connection harvesting, logging, and statistics
- Performance and stabilization enhancements
- Improved diagnostics and statistics or metrics

UCP for JDBC provides advanced connection pooling functions, improved performance, and better diagnosability of connection issues.

See Also:

Oracle Universal Connection Pool for JDBC Developer's Guide for details

1.4.1.30 Java API for Oracle RAC FAN High Availability Events

A new Java API allows Oracle RAC customers who are not using an Oracle connection pool to receive Fast Application Notification (FAN) events (for example, DOWN and UP) and then process these events, clean up or add connections when an instance, service or node leaves or joins the cluster.

Applications using this API can be notified quickly when a failure occurs in the cluster.

> **See Also:**
>
> *Oracle Real Application Clusters Administration and Deployment Guide* for details

1.5 Database Overall

The following sections describe new database features for Oracle Database 11*g* Release 2 (11.2).

1.5.1 General

The following sections provide new feature information for Flashback Data Archive and instance caging.

1.5.1.1 Flashback Data Archive Support for DDLs

Oracle Database 11*g* Release 2 (11.2) users can now use most DDL commands on tables that are being tracked with Flashback Data Archive. This includes:

- Add, Drop, Rename, Modify Column
- Drop, Truncate Partition
- Rename, Truncate Table
- Add, Drop, Rename, Modify Constraint

For more complex DDL (for example, upgrades and split table), the Disassociate and Associate PL/SQL procedures can be used to temporarily disable Total Recall on specified tables. The Associate procedure enforces schema integrity after association; the base table and history table schemas must be the same.

This feature makes it much easier to use the Total Recall option with complex applications that require the ability to modify the schema.

1.5.1.2 Instance Caging

Instance Caging allows the DBA to limit the CPU usage of an Oracle instance by setting the CPU_COUNT initialization parameter and enabling CPU resource management.

With Instance Caging, users can partition CPU resources among multiple instances running on a server to ensure predictable performance.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for details

1.5.2 Improvements to Oracle Scheduler

The following sections detail improvements made to Oracle Scheduler.

1.5.2.1 E-mail Notification

Oracle Database 11*g* Release 2 (11.2) users can now get e-mail notifications on any job activity.

This feature improves efficiency by enabling users to be notified of any job activity that is of interest to them without having to constantly monitor the job.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for details

1.5.2.2 File Watcher

File watcher enables jobs to be triggered when a file arrives on a given machine.

This feature improves efficiency and ease-of-use. Jobs with file dependencies are automatically triggered when the specified file is received instead of constantly monitoring for the file.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for details

1.5.2.3 Multiple Destination Jobs

This feature enables users to specify multiple destinations for a job.

This is a key feature for Enterprise Manager scheduling. It improves efficiency and ease-of-use by enabling a job to be run on multiple nodes while managing it as one entity from a central location.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for details

1.5.2.4 Remote Database Jobs

This feature enables users to run PL/SQL blocks or stored procedures that reside in a remote database as a job.

This is a key feature for Enterprise Manager scheduling. It improves efficiency and ease-of-use by enabling job scheduling in a distributed environment to be managed centrally.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for details

1.5.3 Improvements to Utilities

The following features provide improvements to the various utilities in Oracle Database 11*g* Release 2 (11.2).

1.5.3.1 Data Pump Legacy Mode

Data Pump Legacy Mode provides backward compatibility for scripts and parameter files used for original Export and Import scripts.

This feature enables users to continue using original Export and Import scripts with Data Pump Export and Import. Development time is reduced as new scripts do not have to be created.

Oracle Database Utilities for details

1.5.4 IPv6 Support

The following sections describe improvements in IPv6 networking support.

1.5.4.1 Complete IPv6 Support for JDBC Thin Clients

JDBC supports Internet Protocol Version 6 (IPv6) style addresses in the JDBC URL and machine names that resolve to IPv6 addresses. For example:

```
2001:0db8:0000:0000:0000:0000:0000:0001
1080:0:0:0:8:800:200C:417A
```

A JDBC URL would look like the following:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)
(HOST=[2001:0db8:0000:0000:0000:0000:0000:0001]) (PORT=5521))
(CONNECT_DATA=(SERVICE_NAME=boston.us.example.com)))
```

This feature provides Java interoperability with IPv6.

1.5.4.2 Complete IPv6 Support for JVM and Java Debuggers

IPv6 now provides support for the database resident Java virtual machine.

This allows Java applications running in the database to connect to and accept connections from both IPv4 and IPv6 hosts.

See Also:

Oracle Database Java Developer's Guide for details

1.6 Diagnosability

The following sections describe diagnosability features for Oracle Database 11*g* Release 2 (11.2).

1.6.1 Support Workbench

Enterprise Manager Support Workbench is a GUI workbench for customers and support to ease diagnosis and resolution of database errors.

1.6.1.1 Enterprise Manager Support Workbench for ASM

Enterprise Manager Support Workbench (Support Workbench) has been enhanced to help diagnose and package incidents to Oracle support for Automatic Storage Management (ASM) databases.

This feature extends the benefit of Enterprise Manager Support Workbench to ASM by helping customers package all necessary diagnostic data for incidents.

See Also:

Oracle Automatic Storage Management Administrator's Guide for details

1.7 Information Integration

This is area 2581

The following sections describe new information integration features for Oracle Database 11*g* Release 2 (11.2).

1.7.1 Oracle Database XStream

The following section describes the new Oracle Database XStream feature.

1.7.1.1 XStream

XStream provides application programming interfaces (APIs) that enable client applications to receive data changes from an Oracle database and to send data changes to an Oracle database. These data changes can be shared between Oracle databases and other systems. The other systems include non-Oracle databases, non-RDBMS Oracle products, file systems, third party software applications, and so on. The client application is designed by the user for specific purposes and use cases.

XStream consists of two components: XStream Out and XStream In. XStream Out provides APIs that enable you to share data changes made to an Oracle database with other systems. XStream In provides APIs that enable you to share data changes made to other systems with Oracle databases.

> **See Also:**
>
> *Oracle Database XStream Guide* for details

1.8 Performance

The following sections describe improvements in performance of the database and functionality of performance-related database features.

1.8.1 General Server Performance

The following sections describe general server performance enhancements.

1.8.1.1 Database Smart Flash Cache

New in Oracle Database 11*g* Release 2 (11.2), the Database Smart Flash Cache feature is a transparent extension of the database buffer cache using solid state device (SSD) technology. The SSD acts as a Level 2 cache to the (Level 1) SGA.

Database Smart Flash Cache can greatly improve the performance of Oracle databases by reducing the amount of disk I/O at a much lower cost than adding an equivalent amount of RAM.

1.8.1.2 Stored Outlines Migration to SQL Plan Management

Stored outlines can be migrated for future and enhanced usage with SQL Plan Management (SPM).

Stored outlines lack the flexibility and adaptability of SQL Plan Management. By providing a migration path, old applications using stored outlines can be transparently migrated and can instantaneously take advantage of the enhanced functionality of SPM.

See Also:

Oracle Database Performance Tuning Guide for details

1.8.1.3 Client Result Cache Using Table Annotations Support

Table annotations support provides the ability to annotate a table as being cache worthy, which enables applications to leverage client and server result caching through deployment time knobs as opposed to making application changes. In addition, this feature provides automatic client cache invalidation.

This feature allows non-intrusive application performance acceleration using client and server result caches.

See Also:

Oracle Call Interface Programmer's Guide for details

1.8.1.4 Support 4 KB Sector Disk Drives

Today, disk drives have 512 byte sectors. Disk drive manufacturers are moving to 4 KB sector drives because it allows them to offer higher capacity with lower overhead. If customers use 4 KB sector drives as 512 byte sector drives, then there is likely to be a performance penalty (because they have to run in 512 byte emulation mode). This feature allows Oracle to work with 4 KB (and 512 byte) sector drives without a performance penalty. There is also the capability in Automatic Storage Management (ASM) to allow migration of a disk group from 512 byte sector drives to 4 KB sector drives.

This feature allows customers to take full advantage of new generation, higher capacity disk drives.

See Also:

Oracle Database SQL Language Reference for details

1.9 Security

The new features discussed in the following section cover areas that include encryption and auditing. Significant new encryption key management functionality has been introduced in Oracle Database 11*g* Release 2 to enable complete integration with Hardware Security Modules and increased performance for Transparent Data Encryption. Audit Management has been simplified through the introduction of a new package for managing audit data on the Oracle database.

1.9.1 Audit Data Management

Simplify the management of audit data created by the Oracle Database to facilitate compliance with various privacy and compliance mandates such as SOX, HIPAA, and PCI.

1.9.1.1 Audit Trail Cleanup

Audit Trail Cleanup provides the ability to manage the Oracle database audit trail by:

- Automating the periodic deletion of audit records from the database tables and operating system files after they have been securely backed up or are no longer needed.

- Controlling the size and age of the audit trail written to operating system files before a new operating system audit trail file is created.

- Moving the database audit trail tables out of the `SYSTEM` tablespace to a different tablespace.

Audit Trail Cleanup reduces the time and cost required to manage the Oracle database audit content. It enables you to dedicate an optimized tablespace for audit records and move the audit tables out of the `SYSTEM` tablespace for improved performance. In addition, it provides automated deletion of audit records from the database tables and operating system files.

1.9.2 Encryption Key Management

Encryption key management provides the ability to change the master key associated with transparent data encryption (TDE) encrypted tablespaces. The tablespace master key is used to encrypt the encryption keys associated with individual tablespaces. This is commonly referred to as a 2-tier key architecture. Prior to Oracle Database 11*g* Release 2 (11.2), changing the master key was only possible when using TDE column encryption.

1.9.2.1 Tablespace Master Key Rekey

In Oracle Database 11*g* Release 2 (11.2), Oracle Advanced Security capability allows customers to change the master key used to protect the encryption keys used to encrypt Oracle tablespaces.

Industry initiatives, such as the Payment Card Industry Data Security Standard (PCI DSS), mandate periodic rotation of encryption keys associated with credit card data. This support is now available in this release.

> **See Also:**
>
> *Oracle Database Advanced Security Administrator's Guide* for details

1.10 Server Manageability

The following sections describe server manageability features for Oracle Database 11*g* Release 2 (11.2).

1.10.1 Automatic Storage Management for All Data

These features extend the capabilities of Automatic Storage Management (ASM) to support all types of data including database files, clusterware files, and file system data such as Oracle homes and binaries.

The following sections describe ASM features.

1.10.1.1 ASM Cluster File System (ACFS)

The ASM Cluster File System (ACFS) extends Automatic Storage Management (ASM) by providing a robust, modern, general purpose file system for files beyond the Oracle database files. ACFS provides support for files such as Oracle binaries, report files, trace files, alert logs, and other application data files. With the addition of the Oracle ASM Cluster File System, ASM becomes a complete storage management solution for both Oracle database and non-database files.

ACFS supports large files with 64-bit file and file system data structure sizes leading to exabyte-capable file and file system capacities. ACFS scales to hundreds of nodes and

uses extent-based storage allocation for improved performance. A log-based metadata transaction engine is used for file system integrity and fast recovery. The ACFS on-disk structure supports endian neutral metadata. ACFS file systems can be exported to remote clients through industry standard protocols such as NFS and CIFS.

Oracle ASM Cluster File System (ACFS) complements and leverages Automatic Storage Management (ASM) and provides a general purpose journaling file system for storing and managing non-Oracle database files. This eliminates the need for expensive third-party cluster file system solutions while streamlining, automating and simplifying all file type management in a single node as well as Oracle RAC and Oracle Grid infrastructure for a cluster computing environments.

ACFS supports dynamic file system expansion and contraction without any downtime. ACFS is highly available leveraging the ASM mirroring and striping features in addition to hardware Redundant Array of Inexpensive Disks (RAID) functionality.

> **See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide* for details

1.10.1.2 ASM Dynamic Volume Manager (DVM)

The ASM Dynamic Volume Manager (DVM) is a kernel-loadable device driver that provides a standard device driver interface to clients (for example, ACFS). File systems or other processes can do I/O to this device driver as they would to any other disk device driver on the system. DVM is the primary I/O interface for ACFS to perform I/O and build a file system leveraging ASM as a volume manager. DVM is loaded on ASM start up. The device driver is cluster-aware and communicates with ASM for extent map information, extent rebalancing, and I/O failures.

The ASM Dynamic Volume Manager (DVM) provides a standard I/O interface allowing general purpose file systems to leverage the full functionality of ASM as a volume manager. Oracle database files as well as non-Oracle database files, for example Oracle binaries, can now reside on ACFS eliminating the need for third-party file systems or volume managers to host general purpose files.

> **See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide* for details

1.10.1.3 ASM FS Snapshot

ASM FS Snapshot is a point-in-time copy of a file system and can provide up to 64 snapshot images. ASM FS Snapshot performs fast creation of persistent ASM FS images at a specific point-in-time with low overhead leveraging the Copy on Write technology.

Read-only ASM FS Snapshots can be generated on an interval basis. They may reside in existing ASM FS storage or in an additional storage device and persist following a system restart.

Even as the file system changes, the snapshot does not, giving you the ability to view the file system as it was at the time the snapshot was created. Initially, snapshots are read-only, which preserves their point-in-time capture. The following are the benefits of ASM FS Snapshots:

- ASM FS Snapshots can be used as a source for backup. The original file system can continue to change but the static nature of the snapshot makes them ideal as a source for backup without keeping the original file system offline.

- ASM FS Snapshots can be used as a means for you to recover accidentally deleted or modified files.

- ASM FS Snapshots can be used as a source for data mining or report applications which need to work on a static, point-in-time data set.

> **See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide* for details

1.10.1.4 Oracle Cluster Registry (OCR) and Voting Disk on ASM

Automatic Storage Management (ASM) disks are used to store the Oracle Cluster Registry (OCR) and the voting disks. ASM Partnership and Status Table (PST) is replicated on multiple disks and is extended to store the OCR. Consequently, the OCR tolerates loss of the same number of disks as the underlying disk group. OCR is relocated in response to disk failures.

ASM reserves a number of blocks at a fixed location of every ASM disk for storing the voting disk. Should the disk holding the voting disk fail, ASM selects another disk to store this data.

Storing the OCR and the voting disk on ASM eliminates the need to use expensive third-party cluster volume managers or deal with the complexity of managing disk partitions for OCR and voting disks in Oracle RAC configurations.

> **See Also:**
>
> *Oracle Grid Infrastructure Installation Guide for Linux* for details

1.10.1.5 ASM Intelligent Data Placement

Disk drives have higher transfer rates and bytes per track on the outer tracks. This makes it preferable to keep the hotter data closer to the edge of the disk; that is, the lower numbered blocks. This feature enables ASM to identify higher performance disk regions. Most frequently accessed ASM files can be marked to be moved into the hot region and take advantage of higher I/O performance (for example, hot tablespaces and indices) and able to better meet the application I/O demand. This feature is only applicable when whole physical disks are presented to ASM versus local unit numbers (LUN).

Most frequently accessed Oracle database files in ASM disk groups (ASM files) can be placed in hot disk regions to deliver higher bandwidth and reduce seek latency to meet the application I/O performance requirements.

> **See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide* for details

1.10.1.6 ASM Storage Management Configuration Assistant

ASM Storage Management Configuration Assistant was previously known as Enterprise Manager Integration with ASM Optimal Disk Placement.

This release now allows the configuration, monitoring, and management of Optimal Disk Placement with Enterprise Manager which is the graphical user interface (GUI) for management.

This GUI manages ASM which makes storage management easier in an Oracle environment.

See Also:

Oracle Automatic Storage Management Administrator's Guide for details

1.10.1.7 Automatic Storage Management (ASM) File Access Control

Automatic Storage Management (ASM) on UNIX platforms implements access control on its files to isolate different database instances from each other and prevent unauthorized access. ASM implements new SQL statements to grant, modify, and deny file permissions. The new security model and syntax is coherent with those already implemented for the objects represented in Oracle Database.

Multiple database instances can store ASM files in the same disk group and, therefore, are able to consolidate multiple databases with security. This prevents unauthorized database instances from accessing or overwriting each other's files.

See Also:

Oracle Automatic Storage Management Administrator's Guide for details

1.10.1.8 ASMCMD Command Extensions

The ASMCMD tool is extended to include management of ASM disks, disk groups, and ASM instance in addition to managing ASM files. This is a comprehensive command-line interface that parallels the SQL*Plus command functionality and provides an easy user interface for the system and storage administrators to manage ASM.

The ASMCMD extensions provide the system and storage administrators with a comprehensive and user friendly command-line interface to manage ASM from all perspectives.

See Also:

Oracle Automatic Storage Management Administrator's Guide for details

1.10.1.9 Enterprise Manager Support for ASM Cluster File System (ACFS)

Oracle Enterprise Manager provides a graphical user interface (GUI) to manage the ASM Dynamic Volume Manager and ASM cluster file system (ACFS) as part of the Automatic Storage Management (ASM) solution.

Enterprise Manager provides a graphical user interface which makes is easier to manage the environment whether it is a standalone server or a cluster deployment of ASM. The centralized console provides a consistent interface for managing volumes, database files, and file systems as well as the Oracle Database.

See Also:

Oracle Automatic Storage Management Administrator's Guide for details

1.10.1.10 Enterprise Manager Integration for ASM File Access Control

In this release, Oracle provides a graphical user interface (GUI) for managing File Access Control for Automatic Storage Management (ASM) files.

This GUI simplifies management of ASM for the DBA, system administrator, or storage administrator.

> **See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide* for details

1.10.2 Database Management

The following sections describe general database management features to ease database management.

1.10.2.1 EMCA Supports New Oracle RAC Configuration for Enterprise Manager

Enterprise Manager Configuration Assistant (EMCA) has been updated to support the new configuration required for Enterprise Manager to support new features of the release.

Configuration assistants automate the configuration of the environment ensuring the correct steps are taken. The assistants simplify the configuration of Enterprise Manager in clusters and clustered database settings.

> **See Also:**
>
> *Oracle Enterprise Manager Configuration Plug-In chapter* for details

1.10.2.2 Patch Application with DBControl

Enterprise Manager DBControl manages the application of patches to a single-instance database.

Using Enterprise Manager to apply patches simplifies software maintenance.

> **See Also:**
>
> *Oracle Database 2 Day + Real Application Clusters Guide* for details

1.10.2.3 Automatic Patching of Time Stamp With Time Zone Data

Time stamp with time zone data could become stale in the database tables when the time zone version file is updated. Today, users have to manually fix the affected data. This feature updates the system and user data transparently with minimal downtime and provides automatic and transparent patching of time stamp with time zone data whenever a time zone file is updated.

Also, when a server time zone version is patched, all of the clients that communicate with the server need to be patched as well. With this feature, OCI, JDBC, Pro*C, and SQL*Plus clients can now continue to work with the server without having to update their client-side files.

This new feature provides automatic and transparent patching of time stamp with time zone data whenever a time zone file is updated.

> **See Also:**
>
> *Oracle Database Globalization Support Guide* for details

1.10.2.4 Prevent Data Loss for Time Zone with Local Time Zone Data Type

This feature makes `TIMESTAMP WITH TIME ZONE` data type immune to Daylight Saving Time (DST) changes and reduces the overhead of patching time zone data file and upgrading data on disk.

The benefit of this feature is the elimination of the processing cost and the complexity of maintaining `TIMESTAMP WITH TIME ZONE` data type whenever there are new changes to DST transition rule and time zones.

> **See Also:**
>
> *Oracle Call Interface Programmer's Guide* for details

1.10.2.5 Segment Creation on Demand

The initial segment creation for nonpartitioned tables and indexes can be delayed until data is first inserted into an object.

Several prepackaged applications are delivered with large schemas containing many tables and indexes. Depending on the module usage, only a subset of these objects are really being used. With delayed segment creation, empty database objects do not consume any space, reducing the installation footprint and speeding up the installation.

> **See Also:**
>
> *Oracle Database SQL Language Reference* for details

1.10.2.6 Zero-Size Unusable Indexes and Index Partitions

Unusable indexes and index partitions no longer consume space in the database because they become segmentless.

Unusable indexes and index segments are not usable for any data access. Any space allocated by this unusable (dead) object is freed as soon as an object is marked unusable.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for details

1.10.2.7 Metadata SXML Comparison Tool

The Metadata API has been enhanced to provide a cross database comparison tool to compare object metadata of the same type from different databases. This comparison depends on an alternate XML representation, called SXML. Full XML is typically complex and opaque. In contrast, SXML is somewhat simplified and more closely maps to the SQL creation DDL. These SXML documents provide the building blocks for the new comparison tool in which two SXML documents of the same type can be compared and a new SXML document is produced which describes their differences.

This feature enables users to compare objects between databases to identify drift (that is, metadata changes over time) in objects of the same type.

> **See Also:**
>
> *Oracle Database PL/SQL Packages and Types Reference* for details

1.10.2.8 Compare Period Report

A Replay Compare Period Report performs a high-level comparison of workload replay to its capture or to another replay of the same capture. The Replay Compare

Period report contains a summary of the most important changes between the two runs in terms of performance, errors and data divergence. This makes it easier for Database Replay users to understand and test the impact of system changes.

Replay Compare Period Report simplifies understanding and assessment of the impact of system change in testing by providing summarized information on how the replay performed versus capture or other replays in terms of performance, errors and divergence.

> **See Also:**
>
> *Oracle Database Real Application Testing User's Guide* for details

1.10.2.9 Compare SQL Tuning Sets

Compare SQL Tuning Set feature of the SQL Performance Analyzer allows:

- Building a trial from SQL Tuning Set (STS).

- Comparison of two such trials built from two different STSs. A detailed comparison report including any new or missing SQL statements in one and not in another trial and any plan changes noticed in the compared trials is compiled.

Compare SQL Tuning Sets makes it possible for Database Replay users to perform SQL-centric analysis through the SQL Performance Analyzer report. The two STSs are captured as follows:

- One during workload capture on production.

- And, the other during replay on the test system.

Then, the two STSs are used to generate the SPA report.

Compare STS feature can also be used in non-Database Replay scenarios where customers already have existing test scripts and can capture the SQL into two STSs; one for before system change and one for after.

Compare SQL Tuning Sets feature simplifies assessment of system changes by providing detailed SQL-centric analysis when using Database Replay or other load testing mechanisms.

> **See Also:**
>
> *Oracle Database Real Application Testing User's Guide* for details

1.10.2.10 Enable Sampling for Active Data Guard

Active Session History (ASH) is now available on standby systems.

Having ASH data available on standby systems for Data Guard environments allows customers to troubleshoot performance problems specific to their standby environments.

> **See Also:**
>
> *Oracle Database High Availability Overview* for details

1.10.2.11 Exadata Simulation

For a given workload, you can now simulate the possible benefits in I/O interconnect throughput that can be obtained from migration to Exadata architecture. SQL Performance Analyzer, a feature of Oracle Real Application Testing, allows simulation to be performed on a non-Exadata installation without needing to provision the

Exadata system. The SQL Performance Analyzer Exadata simulation feature can be used to identify workloads that are good candidates for Exadata migration.

This feature simplifies simulation and testing of workloads for Exadata migration system change without requiring provisioning of Exadata hardware.

See Also:

Oracle Database Real Application Testing User's Guide for details

1.10.2.12 Global Oracle RAC ASH Report + ADDM Backwards Compatibility

The Active Session History (ASH) report now includes cluster-wide information, greatly enhancing it's utility in identifying and troubleshooting performance issues that span nodes for a cluster database.

Automatic Database Diagnostic Monitor (ADDM) has been enhanced to be backward compatible allowing it to analyze archived data, or data preserved through database upgrades, allowing a customer to do performance comparisons over a longer time frame.

See Also:

Oracle Database Performance Tuning Guide for details

1.10.2.13 Oracle MTS and Streams Support

Database Replay supports capture and replay of workloads on shared server and Oracle Streams architecture.

Customers using shared server and Oracle Streams architecture can benefit from Database Replay testing and have the ability to adopt technology faster.

See Also:

Oracle Database Real Application Testing User's Guide for details

1.10.2.14 Parallel Query + Alt Plan Recording, Export STS Enhancements

This release includes the following enhancements to SQL Tuning Advisor:

- SQL Tuning Advisor may recommend accepting a profile that uses the Automatic Degree of Parallelism (Auto DOP) feature. A parallel query profile is only recommended when the original plan is serial and when parallel execution can significantly reduce the elapsed time for a long-running query.

- While tuning a SQL statement, SQL Tuning Advisor searches real-time and historical performance data for alternative execution plans for the statement. If plans other than the original plan exist, then SQL Tuning Advisor reports an alternative plan finding.

- You can transport a SQL tuning set to any database created in Oracle Database 10*g* (Release 2) or later. This technique is useful when using SQL Performance Analyzer to tune regressions on a test database.

These features are introduced to enhance the capabilities of the SQL Tuning Advisor. The newest version can recommend alternative plans that were seen at some time in the past, in case they perform better, as well as recommending queries to run in parallel if that is beneficial for the total runtime of each query.

The STS export to old releases is designed to help customers with upgrades so that they are using the most recent version of the database software.

> **See Also:**
>
> *Oracle Database Performance Tuning Guide* for details

1.10.2.15 Synchronization Controls

Synchronization controls have been enhanced to allow more concurrency, filtering, and scale-up during replay.

This gives you the ability to replay workload more realistically to identify the impact of system change and provide flexibility during replay.

> **See Also:**
>
> *Oracle Database Real Application Testing User's Guide* for details

1.11 Unstructured Data Management

The new features in the following sections describe the significant performance, developer productivity and advanced capabilities in Oracle Multimedia, Oracle Spatial, Oracle Database SecureFiles, and XML Database.

1.11.1 Enhanced Oracle Multimedia and DICOM Support

The following sections describe the new features and capabilities for Oracle Multimedia and Digital Imaging and Communications in Medicine (DICOM).

1.11.1.1 Attribute Extraction of Requested Attributes Only

Oracle Multimedia now allows extraction of a subset of DICOM metadata attributes as requested by a user or an application, without first extracting all DICOM attributes from the DICOM content.

DICOM content can contain hundreds of metadata attributes. Often only a few attributes are required for indexing, searching and partitioning. Extraction of a subset eliminates the previous requirement of extracting all attributes thus improving performance.

> **See Also:**
>
> *Oracle Multimedia DICOM Developer's Guide* for details

1.11.1.2 Client-Side DICOM Attribute Extraction

Oracle Multimedia now allows DICOM metadata extraction to be performed outside the database by a client tool or in the middle-tier.

This enables extraction of DICOM metadata before the data is loaded into the database, facilitating metadata-based partitioning of DICOM data in the database.

> **See Also:**
>
> *Oracle Multimedia DICOM Developer's Guide* for details

1.11.1.3 DICOM Enhancements

The following enhancements have been added to Oracle Multimedia:

- Metadata extraction

  Any portion of a DICOM attribute can now be extracted (for example, extraction of the last name portion of the patient name attribute).

■ Constraint definitions

Validation of recursive structures such as DICOM Structured Reports can now be specified, and new FOR EACH syntax has been added to allow iteration through all the components in a single predicate.

■ DICOM content processing:

– DEFLATE transfer syntax support has been added.

– RLE compression is now supported.

– Encoding of multi-bit monochrome images is now supported.

– YBR photometric interpretation is now supported.

– DICOM to AVI and DICOM to MPEG conversions are now supported.

– MPEG encapsulated into DICOM format is now supported.

These new features allow more complete and powerful operations on DICOM data that include:

■ Flexible extraction of metadata from DICOM data.

■ Support for DICOM Structured Reports.

■ Presentation of DICOM videos in any browser, without requiring specific DICOM support in the browser.

See Also:

Oracle Multimedia DICOM Developer's Guide for details

1.11.1.4 Watermarking and Image Processing Enhancements

Oracle Multimedia now includes a new applyWatermark method to add an image or text watermark to any supported image. It also supports new image processing operators to remove metadata when creating thumbnail images and to sharpen image quality.

Watermarking is commonly used to prevent misuse of copyrighted or trademarked images. With this feature, watermarking may be enforced by the database. Removing metadata when creating thumbnail images allows for the production of the smallest possible thumbnail images.

See Also:

Oracle Multimedia Reference for details

1.11.2 Enhanced Oracle Spatial Support

The following sections describe new features in Oracle Spatial support for 3D, geocoder, routing engine, GeoRaster, and network data model.

1.11.2.1 3D Visualization Framework

This release includes a set of metadata tables to describe themes, scenes, textures, viewpoints, light sources, non-geographic data, and other elements used to visualize three-dimensional (3D) content. It also delivers a number of performance improvements to the 3D analysis operations.

This metadata support for 3D content enables a consistent way to combine all 3D, raster, vector, and non-geometric data into a unified visualization framework.

Information may be logically grouped into themes to simplify the development, analysis, use, and maintenance of 3D applications.

See Also:

Oracle Spatial Developer's Guide for details

1.11.2.2 Network Data Model Enhancements

Oracle Spatial network data model (NDM) delivers numerous enhancements. These include a 30-50% more memory-efficient representation of user attribute data associated with the network, many additional highly requested path and subpath analysis functions such as traveling sales person (TSP), hierarchical shortest path (HSP), and K-shortest path (KSP).

This release also allows the logical partitioning of networks based on metrics appropriate to the application. For analysis of data associated with the network, NDM has added a network buffer feature to derive the zone of influence with coverage and cost information and a "minimum cost polygon" to allow for the association of geographic points of interest or coverages with a network defined region.

These enhancements allow NDM to support more completely a wider range of requirements found in utility networks, logistics and other applications dependent upon network-based analysis.

See Also:

Oracle Spatial Topology and Network Data Models Developer's Guide for details

1.11.2.3 New GeoRaster JAVA API

With Oracle Spatial, there is a new Java application programming interface (API) to all functions currently available in the existing PL/SQL interface. These include support for all the search, analysis and raster management features in Oracle Spatial GeoRaster.

In addition, this API includes calls to support the development of extraction, transformation, and loading (ETL) tools, Web applications and raster processing applications.

This new feature simplifies the development of Java applications that use, access, and manipulate raster and gridded data sets stored in Oracle Spatial.

See Also:

Oracle Spatial GeoRaster Developer's Guide for details

1.11.2.4 Raster Reprojections and Ground Control Point-Based (GCP) Georeferencing

Oracle Spatial currently includes support for over 4,000 coordinate systems when using vector data. Oracle GeoRaster now supports the reprojection of imagery to any of these 4,000 plus Oracle Spatial coordinate systems.

In this release the GeoRaster feature also supports native storage and georeferencing of Ground Control Point (GCP) data.

These capabilities remove the requirement for third-party tools currently required when using vector data and raster imagery in different coordinate systems. Oracle Spatial can now be used to perform these reprojections.

Ground Control Point-based georeferencing is used in data collection and processing applications. This native GCP storage and georeferencing can be used to georeference raw (non-rectified) and processed (rectified) raster data.

> **See Also:**
>
> *Oracle Spatial GeoRaster Developer's Guide* for details

1.11.2.5 Routing and Geocoding Enhancements

With this release, the Oracle Spatial routing engine is based on the network data model. This increases the ability of Oracle to support the restrictions and conditions required for advanced routing applications.

The Oracle Spatial 11.2 geocoder, in addition to the support for standard address geocoding based on interpolation, now supports point-based geocoding where data sets include the exact location of addresses, intersections, and points of interest.

Oracle Spatial now supports truck routing data sets to produce driving directions that include restrictions based on roads, weight, height, time of day, and other conditions applied to commercial and logistics applications.

Point-based geocoding is becoming increasingly popular because it allows for more accurate results and can be used in situations where interpolation is not possible.

> **See Also:**
>
> *Oracle Spatial Developer's Guide* for details

1.11.3 Oracle SecureFiles

The following sections describe improvements in Oracle SecureFiles.

1.11.3.1 Database File System (DBFS)

The Oracle Database File System (DBFS) enables the database to be used as a POSIX-compatible file system on Linux. This feature includes a PL/SQL package on the database server that enables the DBFS server functionality and a Linux client for DBFS (`dbfs_client`). The `dbfs_client` client is a utility that enables mounting of a DBFS file system as a mount point on Linux. It provides the mapping from file system operations to database operations. The dbfs_client client runs completely in user space and interacts with the kernel through the FUSE library infrastructure.

DBFS Hierarchical Store provides an easy and application-transparent way to archive SecureFiles data that is stored in DBFS file systems to secondary storage tiers such as tape and storage clouds, using DBFS Links. It also allows archived data to be dearchived and brought back into the database on demand.

DBFS makes it easy for files to be accessed by database applications, and for file-based tools to access files stored in the database. With DBFS, all important file data can be seamlessly stored in an Oracle database, providing the benefits of security, backup, performance, and scalability that are standard with the Oracle Database.

SecureFiles is a high performance solution for storing files or unstructured data in Oracle Database. Customers often need to store these files for long periods of time for business or compliance reasons. Consequently, customers are looking to transfer files to cheaper forms of storage in an application-transparent manner to reduce manageability and administration overhead. DBFS Hierarchical Store provides a seamless, automatic, and transparent way to archive cold file data to inexpensive storage.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for details

1.11.3.2 Support for Oracle SecureFiles

Oracle Database 11*g* Release 2 (11.2) introduces a new compression level for SecureFiles LOBs called `COMPRESS LOW`. This compression level introduces a lightweight compression option that removes the majority of the CPU cost that is typical with file compression. Compressed SecureFiles at the `LOW` level now provides a very efficient choice for SecureFiles storage. SecureFiles LOBs compressed at `LOW` generally consume less CPU than BasicFile (pre-11g Release 1) LOBs, consume less storage than BasicFile LOBs and typically makes the application run faster because of reduced disk I/O.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for details

1.11.4 Oracle XML DB Scale and Performance Improvements

The following sections address key customer requirements in the area of scalability and performance by delivering partitioning of XML tables, scaling on registering large XML Schemata, significant improvements in XML indexing and query performance for common real-world workloads, and significant improvements in performance of repository operations.

1.11.4.1 Binary XML Enhancements

This release includes support for partitioning of binary XMLType table and relation tables containing binary XML columns, where the partition key is derived from the XML content. Also included are guidelines on how to optimize performance of binary XML operations.

This new feature allows Oracle partitioning to be used with binary XML content, thereby allowing large volumes of XML data to be managed effectively.

> **See Also:**
>
> *Oracle XML DB Developer's Guide* for details

1.11.4.2 Oracle XML DB Repository Performance Improvements and Guidelines

Oracle XML DB Repository performance improvements include guidance on:

- How to optimize hierarchical queries using `EQUALS_PATH` and `UNDER_PATH` condition.
- How to optimize performance of hierarchical index when querying the repository.

The benefit is improved performance for repository operations.

> **See Also:**
>
> *Oracle XML DB Developer's Guide* for details

1.11.4.3 XMLIndex Enhancements

This release includes improvements to the ability for Oracle to index unstructured, semi-structured, and highly-structured XML documents stored using binary XML.

Also included is support for partitioned indexes and parallel operations. These enhancements incorporate all features of existing XMLIndex and XMLTable Index into a single unified index.

The benefits are high performance query, fragments, and scalar extraction operations on schema and schema-less binary XML storage.

> **See Also:**
>
> *Oracle XML DB Developer's Guide* for details

1.11.4.4 XMLType Partitioning

This feature allows partitioning of the nested tables that are used to manage collections of child elements when storing XML documents using object-based persistence in conjunction with nested tables.

XMLType partitioning enables all the advanced features of the Oracle partitioning option to be used to manage XMLType data.

> **See Also:**
>
> *Oracle Database VLDB and Partitioning Guide* for details

2

Oracle Database 11*g* Release 2 (11.2.0.2) New Features

This chapter contains descriptions of all of the features that are new to Oracle Database 11*g* Release 2 (11.2.0.2). This chapter contains the following sections:

- General
- ACFS Improvements
- Quality of Service (QoS) Management
- Database Replay
- Management

2.1 General

The following sections describe the new features for Oracle Database 11*g* Release 2 (11.2.0.2).

2.1.1 General

The following sections provide information on new features for 11.2.0.2.

2.1.1.1 Control File Updates Can Be Disabled During NOLOGGING Operations

Parameter `DB_UNRECOVERABLE_SCN_TRACKING = [ TRUE | FALSE ]` can be used to turn off control file writes to update fields that track the highest unrecoverable SCN and Time during a `NOLOGGING` direct path operation.

Performance of the `NOLOGGING` load operation could be limited by the control file write I/O.

2.1.1.2 New Package for Configuring Automatic SQL Tuning

A new PL/SQL package, `DBMS_AUTO_SQLTUNE`, has been introduced to provide more restrictive access to the Automatic SQL Tuning feature.

With this package, access to Automatic SQL Tuning can be restricted to DBAs so that only they can change its configuration settings that effect run-time behavior of the query optimizer, such as enabling or disabling automatic SQL profile creation.

> **See Also:**
>
> *Oracle Database Performance Tuning Guide* for details

2.1.1.3 Enhanced Security for DBMS_SCHEDULER E-Mail Notification

Encryption and authentication have been added to the Oracle Scheduler's e-mail notification feature.

E-mail notification on job failures was added in 11.2.0.1, but it did not support those e-mail servers that require either encryption or authentication. This feature adds this support in 11.2.0.2.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for details

2.1.1.4 Enhanced TRUNCATE Functionality

While truncating a table or partition, you can now specify whether or not to keep any segments. Truncating a table or partition with the new extended syntax removes all segments and does not use any space until new data is inserted.

All allocated space in a database can now be reclaimed by truncating tables or partitions with the new extended syntax, optimizing the space foot print of any application.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for details

2.1.1.5 Improvements to Reduce the Occurrences of "Existing State of Packages Discarded" Errors

When the body of a stateful package is recompiled by session A, then a different session B which has already instantiated that package will experience the `ORA-04068 existing state of packages… has been discarded` family of errors on its next reference to that package following the recompilation. The result is severe disruption for the user of session B.

Two improvements in this release now reduce the likelihood of these `ORA-04068` errors:

- A package, whose global variables is marked `constant`; and where the initial value can be computed at compile time, is now considered to be stateless.

- The new procedure `DBMS_SESSION.RESET_SELF()`, whose invocation must be in the package of interest, allows you to preempt the problem when it is safe to reinitialize the package state.

The advantage is that packages (and especially those that are not editioned) can be "hot patched" (that is, recompiled) when other sessions are using them with a greatly reduced likelihood of disrupting those sessions.

> **See Also:**
>
> *Oracle Database PL/SQL Language Reference* for details

2.1.1.6 Support for In-Place Upgrade of Clients

Both in-place and out-of-place upgrades are supported for client installations.

You now have the option of doing in-place client upgrades reducing the need for extra storage and simplifying the installation process.

2.1.1.7 Maintenance Package for Segment Creation on Demand

Customers can manage the space allocation of any application through extended functionality of the DBMS_SPACE package. This package can be used to remove the segments for all empty tables in a database, a user schema, or for specific tables. This package also provides the opposite functionality to materialize all segments for empty tables or partitions with deferred segment creation

The explicit management of deferred segment creation enables you to take advantage of this functionality at any given point in time, even after table or partition creation. This is especially useful for systems that were upgraded in-place and makes a re-creation of all empty objects unnecessary.

2.1.1.8 Maximum CPU Utilization Limit

Resource Manager provides a new directive called MAX_UTILIZATION_LIMIT that allows you to place a hard limit on the amount of CPU utilized by a consumer group.

The MAX_UTILIZATION_LIMIT directive is useful for limiting the CPU utilization of low priority workloads. This directive is also useful for providing more consistent performance for the workload in a consumer group, and it helps to build systems where end users experience consistent response times for each database operation.

2.1.1.9 Name Matching

This feature provides an efficient method for matching proper names (and words) that take a query as input and returns a ranked list of matches. The new operator NDATA is introduced for this functionality.

In today's multicultural society, a person accustomed to the spelling rules of one demographic may have difficulty applying those same rules to a name originating from a different culture.

Name matching provides a solution to match proper names that might differ in spelling due to orthographic variation.

2.1.1.10 Named Entity Extraction

Entity extraction is the recognition of entity names (people and organizations), places, temporal expressions, and types of numerical expressions such as currencies and measures.

The goal of entity extraction is to identify instances of a particular pre-specified class of entities in textual documents.

The benefit is to produce a "structured" view of a document that can later be used for text or data mining and more comprehensive intelligence analysis.

See Also:

Oracle Text Reference for details

2.1.1.11 Default Size of First Extent of Any New Segment for a Partitioned Table Has Changed

The default size of the first extent of any new segment for a partitioned table is now 8 MB instead of 64 K.

The goal is to improve I/O performance. However, under certain circumstances, loading a table will take significantly more disk space.

See Also:

Oracle Database Administrator's Guide for details

2.1.1.12 Parallel Statement Queuing

Parallel Statement Queuing ensures all statements run on a system get the appropriate parallel resources to perform well by allowing you to ensure that a system is neither overwhelmed nor starved for parallel server processes. Queuing can be implemented per resource group and allows for both prioritization of statements and the above mentioned management of a parallel workload. Parallel Statement Queuing works in conjunction with Automatic Degree of Parallelism.

Data Warehouses are evolving into systems that support both operational environments and the more classic strategic data warehouse workloads. These mixed workloads require active workload management. One of these resources that should be managed as part of the workload management process is the use of Parallel Server Processes. Parallel server resources are allocated by Automatic Degree of Parallelism (DOP). Statement queuing is then used to ensure that each statement can run with the optimal DOP within the system limits. Allowing each statement to run with the optimal DOP allows a system to:

- Perform well overall and avoid large wait times on system resources.

- Utilize all resources in an optimal manner without trashing the system in peak times or due to runaway queries.

- Balance overall performance to be much more predictable.

- Allocate appropriate resources based on policies, not based on user abuse.

See Also:

Oracle Database VLDB and Partitioning Guide for details

2.1.1.13 PMML Import

This release adds support for importing external data mining models (linear and binary logistic regression) using the Data Mining Group Predictive Model Markup Language (PMML) standard. The imported models become native Oracle Data Mining (ODM) models capable of Exadata offload.

If you use an external data mining product to generate models, you could encounter difficulty when deploying those models into their production databases. The current process of deploying such models is expensive, error prone, and non-performant. This feature streamlines the movement of external models into production Oracle systems and leverages optimized performance of the ODM option.

> **See Also:**
>
> *Oracle Database PL/SQL Packages and Types Reference* for details

2.1.1.14 Result Set Interface

The client interface `CTX_QUERY.RESULT_SET` executes a query and generates a result set. The components of the result set are:

- Documents.

- Support order by `SDATA`.

- A total estimated count of number of matching documents.

- A count, broken down by metadata value, of matching documents in each category.

A page of search results consist of many disparate elements (for example, metadata of the first few documents, snippet, total hit counts, and so on). Instead of accessing the database to construct bits of the search results, it would be useful to have a clean result set mechanism. The result set interface is able to produce the various kinds of data needed for a page of search results all at once, improving performance by sharing overhead. The result set interface can also return data views which are difficult to express in SQL, such as top n by category queries.

> **See Also:**
>
> *Oracle Text Application Developer's Guide* for details

2.1.1.15 Segment Creation On Demand for Partitioned Tables

The initial segment creation for partitioned tables and indexes can be deferred until data is first inserted into an object. Individual partitions will not be physically created before data is inserted for the first time.

Several prepackaged applications are delivered with large schemas containing many partitioned tables and indexes. With deferred segment creation for partitioned tables, empty database objects do not consume any space, reducing the installation footprint and speeding up the installation.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for details

2.1.1.16 Simplification of XML and XQuery Interfaces

This feature extends the XQuery 1.0 standard's operator `fn:doc` and `fn:collection` to allow direct access to collections of XML documents stored in the database.

Direct access to XML content in tables and views is provided by extending `fn:doc` and `fn:collection` to support DBUri-style paths through the pseudo protocol `xdb://`.

Simplification of Oracle XML and XQuery interfaces provides standard mechanisms, allows building of portable XML applications that are easier to maintain, and deprecates redundant or unused functionality.

> **See Also:**
>
> *Oracle XML DB Developer's Guide* for details

2.1.1.17 SMTP Authentication

Starting with this release, you can configure the UTL_SMTP PL/SQL package for use on both Transport Layer Security (TLS) and Secure Sockets Layer (SSL) servers."

This allows the package to be used to send to SMTP servers that require authentication to combat spam.

> **See Also:**
>
> *Oracle Database PL/SQL Packages and Types Reference* for details

2.1.1.18 SMTP Encryption

UTL_SMTP is extended in this release to provide Secure Sockets Layer (SSL) and Transport Layer Security (TLS) support.

This allows the package to be used to send to SMTP servers using SSL and TLS to ensure channel integrity.

> **See Also:**
>
> *Oracle Database PL/SQL Packages and Types Reference* for details

2.1.1.19 SPA Support for Active Data Guard Environment

If you are using Oracle Active Data Guard physical standby database, you already have full dataset or clone or both of the production environment that can be leveraged for testing with SQL Performance Analyzer (SPA). Using remote test execution SPA trial method, you can connect to a physical standby database in read-only mode and use it for testing. The physical standby database continues to be in read-only and standby mode (changes are being applied) during SPA testing. The SPA analysis and reports are available from the remote database that is orchestrating the SPA trials. The orchestrating database (SPA system) can be the primary database or any remote database running Oracle Database 11*g* and higher releases.

This feature allows customers to leverage existing Active Data Guard physical standby databases for SQL Performance Analyzer Testing.

> **See Also:**
>
> *Oracle Database Real Application Testing User's Guide* for details

2.1.1.20 The EDITION Attribute of a Database Service

The EDITION attribute of a database service specifies the initial session edition for a session that is started using that service. If the program that creates a new session does not specify the initial session, then the edition name specified by the service is used. If the service does not specify the edition name, then the initial session edition is the database default edition.

When an edition-based redefinition exercise is implemented to support hot rollover, some clients to the database will want to use the pre-upgrade edition and others will want to use the post-upgrade edition. In this scenario, the database default edition is insufficient because, by definition, it denotes a single edition. The EDITION attribute of a database service provides a way to allow the client to specify the edition it wants using environment data rather than by changing the client code.

> **See Also:**
>
> *Oracle Database PL/SQL Packages and Types Reference* for details

2.1.1.21 Using Binary XML with SecureFiles as the XMLType Default Storage

In this release, the default storage model has changed for XMLType from `STORE AS CLOB` to `STORE AS SECURE FILE BINARY XML`. This affects the storage used when an explicit `STORE AS` clause is not supplied when creating an XMLType table or column. Not specifying a `STORE AS CLAUSE` indicates that it is left to the database to determine what the optimal storage model should be.

Prior to database release 11.2.0.2, the default storage model was `STORE AS BASICFILE CLOB`. In 11.2.0.2, the default is changed to `STORE AS SECUREFILE BINARY XML`.

This change requires the installation of the XDB feature in order to work correctly. Customers that choose not to install the XDB feature must explicitly add `STORE AS CLOB` to any DLL statements that create XMLType table or columns to avoid DDL errors. Note that the use of XMLType without having the XDB installed is not a supported configuration as of 11.1.0.1.

No data migration takes place when databases are upgraded to 11.2.0.2.

Binary XML with SecureFiles provides efficient storage, retrieval, and DML capabilities for semi-structured and unstructured XML data. Changing the default storage for XMLType to binary XML with SecureFiles helps customers to adopt best practices.

> **See Also:**
>
> *Oracle XML DB Developer's Guide* for details

2.1.1.22 JDBC 4.0 SQLXML

This feature implements the JDBC 4.0 specification of the `SQLXML` interface for managing the XML data type in the database.

This feature allows Java applications using JDBC-Thin or JDBC-OCI to manage the XML data type in the database, using the standard `SQLXML` type (`java.sql.SQLXML`).

> **See Also:**
>
> *Oracle Database JDBC Developer's Guide* for details

2.1.1.23 ID Key LCRs in XStream

ID key LCRs enable an XStream client application to process changes to rows that include unsupported data types. ID key LCRs do not contain all of the columns for a row change. Instead, they contain the rowid of the changed row, a group of key columns to identify the row in the table, and the data for the scalar columns of the table that are supported by XStream Out. ID key LCRs do not contain columns for unsupported data types.

This feature enables XStream users to capture database changes that cannot be supported using Oracle Streams.

> **See Also:**
>
> *Oracle Database XStream Guide* for details

2.1.2 ACFS Improvements

The following sections provide information on ACFS improvements for 11.2.0.2.

2.1.2.1 ACFS, ADVM and Snapshots on Solaris and AIX

Oracle ACFS, Oracle ASM Dynamic Volume Manager (Oracle ADVM) and Snapshots were delivered in Oracle Database 11*g* Release 2 (11.2.0.1) on Windows NT and Linux platforms.

Oracle Database 11*g* Release 2 (11.2.0.2) now provides a general purpose cluster file system which leverages the capabilities of Oracle ASM on Solaris and AIX platforms.

> **See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide* for details

2.1.2.2 Oracle ACFS Replication

The Oracle Automatic Storage Management Cluster File System (Oracle ACFS) Replication feature supports asynchronous replication of an ACFS file system from a primary to standby site.

The Oracle ACFS Replication feature allows you to replicate ACFS file systems across the network to another (possibly distant) site. This provides a disaster recovery capability for the file system. This feature can be used in conjunction with Oracle Data Guard to replicate all Oracle data.

> **See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide* for details

2.1.2.3 Oracle ACFS Security and Encryption Features

Oracle ASM Cluster File System (Oracle ACFS) security feature provides realm-based security for Oracle ACFS.

Oracle ACFS encryption feature enables data stored on disk (data-at-rest) to be encrypted.

Oracle ACFS security feature provides the ability to create realms to specify security policies for users or groups for accessing file system objects. The Oracle ACFS security feature provides a finer-grained access control on top of the access control provided by the operating system.

Oracle ACFS encryption feature provides the ability to keep data in an Oracle ACFS file system in encrypted format to prevent unauthorized use of data in the case of data loss or theft.

> **See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide* for details

2.1.2.4 Oracle ACFS Tagging

The Oracle ACFS Tagging feature provides a method for relating a group of files based on a common naming attribute assigned to these files called a tag name.

You can use this feature alone or in conjunction with other features. For example, in conjunction with Oracle ACFS Replication, you can select specific files that you would like to replicate to a different remote cluster site by assigning a unique tag name to them. You would then instruct Oracle ACFS Replication to replicate files based upon this tag name. By using tagging in this respect, the need to replicate entire Oracle ACFS file systems is reduced.

> **See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide* for details

2.1.3 Quality of Service (QoS) Management

A new Quality of Service (QoS) Management Server enables run time management of service levels for hosted database applications on a shared infrastructure by cluster administrators. The goal is to present an easy-to-use, policy-driven management system that ensures meeting service levels if sufficient resources are available and when they are not, allocates resources to the most business critical workloads not meeting their service levels at the expense of the less critical ones.

The following sections describe Quality of Service Management Server features.

2.1.3.1 Database QoS Management Server

The Database Quality of Service (QoS) Management Server allows system administrators to manage application service levels hosted in Oracle Database clusters by correlating accurate run-time performance and resource metrics and analyzing with an expert system to produce recommended resource adjustments to meet policy-based performance objectives.

The Database QoS Management Server enables the pooling of resources to help ensure that, when sufficient resources are available, performance and availability objectives are met, even under demand surges. Managing resource allocations to match performance objectives using a set of predefined policies, the Database QoS Management Server greatly reduces system administrator and DBA time and expertise. By continuously monitoring the system performance based on real demand, it quickly identifies bottlenecks and potential problems that can be corrected before an actual outage occurs. This system cuts time to resolve service level violations as it provides detailed metrics and bottleneck identification along with recommendations for resolution. The end result is the stakeholders trust to share resources thus reducing capital and operational expenses.

> **See Also:**
>
> *Oracle Database Quality of Service Management User's Guide* for details

2.1.3.2 Database Quality of Service (QoS) Management Support

To support the Database Quality of Service (QoS) Management Server, the Oracle Database Resource Manager and metrics have been enhanced to support fine-grained performance metrics and now have the ability to manage workloads by user-defined performance classes.

By supporting the Database QoS Management Server, applications sharing a single database or multiple databases within a cluster can be managed discretely to monitor and maintain their service levels. This consolidation reduces hardware, software and management costs while maintaining business objectives.

> **See Also:**
>
> *Oracle Database Quality of Service Management User's Guide* for details

2.1.3.3 Enterprise Manager QoS Management Integration

The administration of the Database Quality of Service (QoS) Management Server is integrated into the new Cluster Administration section of Enterprise Manager. This is designed as a task-based interface to create policy sets using a wizard, manage application service levels using a dashboard, and monitor performance through historical graphs, logs and alerts.

This feature provides full task-based integration into Enterprise Manager, simplifying the administration tasks necessary to manage database application service levels using the Database QoS Management Server. It both reduces task and troubleshooting time as well as the level of training required thus reducing costs while maintaining application availability.

> **See Also:**
>
> *Oracle Database Quality of Service Management User's Guide* for details

2.1.3.4 Server Memory Stress Protection for Oracle Clusters

When QoS Management is enabled and managing an Oracle Clusterware server pool, it receives a metrics stream from the Cluster Health Monitor that provides real-time memory data including the amount available, in use, and swapped to disk for each server. Should a node be determined to be under memory stress, the CRS-managed database services are stopped on that node preventing new connections from being created thereby protecting existing sessions. Once the memory stress is relieved (for example, by either existing sessions closing or user intervention), the services are restarted automatically and the listener begins sending opening connections on that server.

Enterprise database servers can run out of available memory due to too many sessions or runaway workloads. This can result in failed transactions or, in extreme cases, a reboot of the server and loss of a valuable resource. Oracle Database QoS Management detects memory pressure in real-time and prevents the addition of new sessions from exhausting available memory thus protecting existing workloads and the availability of the server. This adds a new resource protection capability in managing Service Levels for Oracle RAC database-hosted applications.

2.1.4 Database Replay

The following sections provide information on new Database Replay features for 11.2.0.2.

2.1.4.1 Database Replay SQL Performance Analyzer (SPA) Integration

This feature allows you to perform SQL Tuning Set (STS) capture and workload capture or replay at the same time in a single process. STS is automatically exported when the AWR data for the capture or replay is exported into the specified directory object. By integrating SPA and Database Replay, you can analyze SQL-centric issues in the workload more easily than if they were to do this manually in separate steps. An SPA report can be generated at the end of workload replay to facilitate SQL-centric analysis. Oracle RAC is not yet supported.

Integration of SPA and Database Replay features provides the ability to perform SQL Tuning Set and workload capture or replay in one process and at the same time. As a result, an SPA report is available to help with SQL-centric analysis when workload replay is done.

See Also:

Oracle Database Real Application Testing User's Guide for details

2.1.4.2 Database Replay Timeout Function

During workload replay, it is sometimes possible that due to an execution plan, system change or otherwise, a replay call may hang or take a long time. You can specify a replay timeout parameter. If the call exceeds the timeout, that particular call is aborted. This is useful with workloads when one or a few calls result in the workload replay to run too long or hang. Aborting these will still provide a useful workload replay.

Database Replay timeout functionality provides the ability to control how long a long running or runaway replay call will take. Without this functionality, a replay call may take a long time or hang depending on the situation.

See Also:

Oracle Database Real Application Testing User's Guide for details

2.1.4.3 Database Replay Workload Analyzer

Database Replay Workload Analyzer is a tool that analyzes a captured workload and provides an assessment of how reliably it can be replayed. It highlights any potential problems that might be encountered during replay by outlining the parts that cannot be replayed accurately due to insufficient data, errors during capture, and usage of features that are unsupported by Database Replay.

This feature tells you, at the time of capture, whether the specific workload captured is something that can be relied upon for future testing.

See Also:

Oracle Database Real Application Testing User's Guide for details

2.1.5 Management

The following sections provide information on new management features for 11.2.0.2.

2.1.5.1 DBCA Support for Creating an Oracle RAC One Node Database

Support has been added in this release to Oracle Database Configuration Assistant (DBCA) to create an Oracle Real Application Clusters One Node (Oracle RAC One Node) database as part of the database creation process.

Oracle RAC One Node is a new option to the Oracle Enterprise Edition introduced with the Oracle Database 11.2.0.1. Oracle DBCA now recognizes Oracle RAC One Node databases and provides the required configuration options to ease the management of Oracle RAC One Node.

See Also:

Oracle Real Application Clusters Administration and Deployment Guide for details

2.1.5.2 Option of Downloading Latest Updates During Installation

This feature allows the installer to download mandatory patches for itself as well as for the base product at installation time so that they do not need to be applied later. It

also helps resolve installation issues at the middle of a release without either recutting the media or deferring the bug fix to a later release.

Currently, when there is a bug in the base installation, you have to wait until the next release before it can be fixed. This feature helps resolve installation issues at the middle of a release without either recutting the media or deferring the bug fix to a later release. The feature also applies mandatory patches for the base product, thereby creating more certified installations out-of-box.

> **See Also:**
>
> *Oracle Database Installation Guide for Linux* for details

2.1.5.3 Oracle ASM Configuration Assistant Support for Out-of-Place Upgrades

Oracle Grid Infrastructure for a Cluster 11*g* Release 2 supports out-of-place upgrades. The Oracle ASM Configuration Assistant (ASMCA) now fully supports out-of-place upgrades to this new release.

The graphical user interface (GUI) provides a simple interactive method for upgrading environments to this new release. To allow scripting, the assistant also provides an on-interactive method (silent) mode, which addresses various deployment scenarios used by customers.

> **See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide* for details

2.1.5.4 Oracle Database Upgrade Assistant Support for Out-of-Place Upgrades

Oracle Grid Infrastructure for a Cluster 11*g* Release 2 supports out-of-place upgrades. The Database Upgrade Assistant (DBUA) now fully supports out-of-place upgrades to this new release.

The graphical user interface (GUI) provides a simple interactive method for upgrading environments to this new release. To allow scripting, the assistant also provides an on-interactive method (silent) mode, which addresses various deployment scenarios used by customers.

> **See Also:**
>
> *Oracle Database Upgrade Guide* for details

2.1.5.5 Oracle Enterprise Manager DB Control Support for Oracle RAC One Node

Oracle Enterprise Manager DB Control provides support for Oracle RAC One Node databases.

Oracle RAC One Node is a new option to the Oracle Enterprise Edition introduced with the Oracle Database 11.2.0.1. Oracle Enterprise Manager DB Control now recognizes Oracle RAC One Node databases and provides the required configuration options in an easy-to-use graphical user interface (GUI), which simplifies the management of Oracle RAC One Node beyond the scope of the command-line tools that are already available.

2.1.5.6 Online Relocation of an Oracle RAC One Node Database

Oracle RAC One Node allows the online relocation of an Oracle RAC One Node database from one server to another. The migration period can be customized up to 12 hours.

Oracle RAC One Node allows the online relocation of an Oracle RAC One Node database from one server to another, which provides increased availability for applications based on an Oracle Database. You can now move a database for workload balancing as well as for performing planned maintenance on the server, on the operating system, or when applying patches to the Oracle software in a rolling fashion.

> **See Also:**
>
> *Oracle Real Application Clusters Administration and Deployment Guide* for details

2.1.5.7 SRVCTL-Based Management of Oracle RAC One Node Databases

Oracle RAC One Node is a new option to the Oracle Database Enterprise Edition. Oracle RAC One Node represents an Oracle RAC database that runs only one active database instance which can be managed using SRVCTL as any other Oracle RAC database.

Using SRVCTL simplifies and optimizes the management of Oracle RAC One Node databases.

> **See Also:**
>
> *Oracle Real Application Clusters Administration and Deployment Guide* for details

2.1.5.8 CRSCTL Command Enhancements

The CRSCTL command set has been enhancement to enable the management of various new Oracle Grid Infrastructure for a Cluster resources.

Using these new commands simplifies the management of Oracle Grid Infrastructure for a Cluster.

> **See Also:**
>
> *Oracle Clusterware Administration and Deployment Guide* for details

2.1.5.9 SRVCTL Command Enhancements

The SRVCTL command set has been enhancement to enable the management of various new Oracle Grid Infrastructure for a Cluster and Oracle RAC resources.

Using these new commands simplifies the management of Oracle RAC and Oracle Grid Infrastructure for a Cluster.

> **See Also:**
>
> *Oracle Real Application Clusters Administration and Deployment Guide* for details

2.1.5.10 Enhanced XStream Manageability

To increase the manageability of XStream, new process parameters are added to provide capabilities such as process memory control, changes to sequences, and the ability to exclude changes performed by specific users or transactions. Repositioning within the stream by either SCN or TIME is available. In addition, new views specific to XStream are provided such as V$XSTREAM_OUTBOUND_SERVER and V$XSTREAM_TRANSACTION, and existing views have been extended to provide additional information such as the client status or memory utilization of a process.

These enhancements give the XStream user more control over and visibility into XStream processing.

See Also:

Oracle Database XStream Guide for details

2.1.5.11 Columnar Compression Support in Supplemental Logging and XStream

Columnar compression is now supported with Oracle Streams and XStream.

This feature enables logical replication of tables compressed using Exadata Hybrid Columnar Compression.

See Also:

Oracle Streams Concepts and Administration for details

2.1.5.12 Standalone Configuration Wizard for Post-Installation Cluster Configuration

The installation of Oracle Grid Infrastructure for a Cluster with Oracle Database 11*g* Release 2 includes a software-only option. This wizard assists the administrator with completing the cluster configuration independently of the software installation.

The configuration wizard provides an easy-to-use interface to configure the cluster independently of the software installation. Post-installation configuration of the software at the customer site is a standing requirement.

Customers that need to be able to mass deploy Oracle Grid Infrastructure for a Cluster or that need to support remote installations benefit from this feature.

See Also:

Oracle Clusterware Administration and Deployment Guide for details

2.1.5.13 Redundant Interconnect Usage

Oracle RAC requires a dedicated network connection between the servers of the Oracle RAC cluster. The dedicated network connection, called interconnect, is crucial for the communication in the cluster. Using redundant network connections for load balancing and for failure protection is recommended. While in previous releases, technologies like bonding or trunking had to be used to make use of redundant networks for the interconnect, Oracle Grid Infrastructure for a Cluster and Oracle RAC now provide a native way of using redundant network interfaces in order to ensure optimal communication in the cluster.

Using redundant interconnects optimizes the stability, reliability, and scalability of an Oracle RAC cluster.

See Also:

Oracle Grid Infrastructure Installation Guide for Linux for details

3

Oracle Database 11*g* Documentation

This chapter lists the books available online that document Oracle® Database 11*g*.

3.1 Documentation Titles and Online Documentation Structure

All books are available in electronic format.

The following categories of documentation are available:

- Oracle® Database 11g Server Documentation

  Subcategories include:

  - Administration, Concepts, and Reference
  - Application Development
  - Application Express
  - Backup and Recovery
  - Data Warehousing
  - Language and Interface
  - Networking and Security
  - Oracle Real Application Clusters
  - Streams and Replication
  - XStream

- Oracle® Database 11g Spatial Documentation
- Oracle® Database 11g Windows Documentation
- Oracle® Database 11g Gateways Documentation
- Oracle® Database 11g Communications Data Model Documentation
- Oracle® Database 11g Java Development Documentation
- Oracle® Database 11g Release Notes Documentation
- Oracle® Database 11g Data Mining Documentation
- Oracle® Database 11g TimesTen In-Memory Database Documentation
- Oracle® Database 11g Documentation Library
- Oracle® Database 11g Warehouse Builder Documentation
- Oracle® Database 11g OLAP Documentation

- Oracle® Database 11g Multimedia Documentation
- Oracle® Database 11g Installation Guide Documentation

3.1.1 Oracle® Database 11g Server Documentation

Table 3–1   Oracle® Database Server Documentation

Documentation	Part Number	Description
Administration, Concepts, and Reference		
Oracle Database Globalization Support Guide	E10729	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 Database 2 Day + Performance Tuning Guide	E10822	Explains how to perform day-to-day database performance tuning tasks using features provided by Oracle Diagnostics Pack, Oracle Tuning Pack, and Oracle Enterprise Manager.
Oracle Database Sample Schemas	E10831	Describes the sample schemas included in the seed database that ships with Oracle Database 11*g*. These schemas are used in examples in Oracle documentation, in curriculum materials, and in demos on Oracle Technology Network.
Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems	E10839	Explains how to create, configure, and manage Oracle Database on Linux and UNIX-based operating systems.
Oracle Database 2 Day DBA	E10897	Provides a single source reference for administering Oracle Database in a small to midsize environment, using the graphical user interface of Oracle Enterprise Manager Database Control. If you are familiar with computers, but unfamiliar with administering Oracle Database, and your computer has four CPUs or fewer, and you do not expect to support more than 500 users, this book is appropriate for you.
Oracle Database Heterogeneous Connectivity User's Guide	E11050	Describes Oracle's approach to information integration in a heterogeneous environment. Describes Oracle Database Gateways and is a user's guide for these Oracle products.
Oracle Universal Installer and OPatch User's Guide for Windows and UNIX	E12255	Describes how to manage Oracle and third-party software using the Oracle Universal Installer and OPatch. This guide provides instructions for installing, setting up, and starting the Oracle Universal Installer software. It also provides instructions on cloning and patching Oracle software.

Table 3–1   (Cont.) Oracle® Database Server Documentation

Documentation	Part Number	Description
Oracle Automatic Storage Management Administrator's Guide	E16102	Describes how to configure and administer Oracle Automatic Storage Management for system, database, and storage administrators who administer and manage storage.
Oracle Database Concepts	E16508	Describes Oracle Database architecture and essential topics for database administrators and developers.
Oracle Database Utilities	E16536	Describes how to use Oracle Database utilities to load data into a database, transfer data between databases, and maintain data. The topics discussed include Data Pump Export, Data Pump Import, SQL*Loader, external tables and associated access drivers, the Automatic Diagnostic Repository Command Interpreter (ADRCI), DBVERIFY, DBNEWID, LogMiner, the Metadata API, original Export, and original Import.
Oracle Database Real Application Testing User's Guide	E16540	Provides information about how to assure the integrity of database changes using Oracle Real Application Testing.
Oracle Database VLDB and Partitioning Guide	E16541	Provides conceptual, reference, and implementation material for using very large databases (VLDB). It emphasizes partitioning as a key component of a VLDB strategy.
Oracle Database Quality of Service Management User's Guide	E16542	Describes the concepts of Oracle Database Quality of Service Management (Oracle Database QoS Management), and instructs the reader on how to configure, administer, and troubleshoot Oracle Database QoS Management. Oracle Database QoS Management manages the resources that are shared across applications and adjusts the system configuration to keep the applications running at the performance levels needed by your business.
Oracle Database Vault Administrator's Guide	E16544	Explains how to configure an Oracle database to use Oracle Database Vault, which enables users to restrict and monitor administrative user access.
Oracle Text Reference	E16593	Provides reference information for building applications with Oracle Text.
*SQL*Plus User's Guide and Reference*	E16604	Contains a complete description of the SQL*Plus command-line interface to the Oracle Database, and a complete reference to SQL*Plus commands and syntax.
*SQL*Plus Quick Reference*	E16605	Contains a complete list of SQL*Plus command syntax.
Oracle Database Performance Tuning Guide	E16638	Describes how to tune the performance of Oracle Database and SQL statements using Oracle performance tools in the command-line interface.

Table 3–1   (Cont.) Oracle® Database Server Documentation

Documentation	Part Number	Description
Oracle Database PL/SQL Packages and Types Reference	E16760	Describes the packages provided with the Oracle database server. The manual is intended for programmers, systems analysts, project managers, and others interested in developing database applications.
Oracle Database Reference	E17110	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.
Oracle Database SQL Language Reference	E17118	Contains a complete description of Oracle SQL, the language used to manage information in Oracle Database.
Oracle Database SQL Language Quick Reference	E17119	Contains a high-level text description of the SQL syntax used to manage information in Oracle Database.
Oracle Database Administrator's Guide	E17120	Describes how to create, configure, and administer an Oracle database. Includes sections on Oracle Scheduler, Oracle Database Resource Manager, Oracle Restart, and distributed computing.
Oracle Database PL/SQL Language Reference	E17126	Presents PL/SQL, the Oracle procedural extension of SQL, an advanced fourth-generation programming language. Explains the concepts behind PL/SQL and illustrates every facet of the language.
Oracle Database New Features Guide	E17128	Provides an overview of features, options, and enhancements that are new in this release of the database.
Oracle Database High Availability Overview	E17157	Introduces Oracle best practices for deploying a highly available environment, helps you determine your availability requirements, and describes the database architectures that can help your business achieve high availability.
Oracle Database Upgrade Guide	E17222	Guides you through the process of planning and executing Oracle Database upgrades. In addition, this manual provides information about compatibility, upgrading applications, and important changes in the new Oracle Database release, such as initialization parameter changes and data dictionary changes.
Oracle Database Error Messages	E17766	Describes error messages that may appear while using products that are part of Oracle Database.

Application Development

Table 3–1 (Cont.) Oracle® Database Server Documentation

Documentation	Part Number	Description
Oracle XML Developer's Kit Programmer's Guide	E10708	Provides information to application developers who need to use components of the Oracle XML Developer's Kit (XDK) to generate and store XML data in a database or in a document outside the database. Examples and sample applications are introduced. Developers need familiarity with XML and a third-generation programming language such as Java, C, or C++.
Oracle Database Data Cartridge Developer's Guide	E10765	Describes how to implement custom indexing and query optimization services and how to package and use these as a server extension called a data cartridge.
Oracle Database 2 Day Developer's Guide	E10766	Explains general concepts behind development with Oracle Database, introduces basic features of SQL and PL/SQL, and provides references to in-depth information elsewhere in the Oracle Database library.
Oracle Database 2 Day + PHP Developer's Guide	E10811	Provides a tutorial on how to develop PHP scripts that use the database languages SQL and PL/SQL to access and manipulate Oracle data.
Oracle Database Object-Relational Developer's Guide	E11822	Describes user-defined object data types and how to use these data types to model complex real-world entities as objects in the database.
Oracle Database Workspace Manager Developer's Guide	E11826	Provides usage and reference information about Oracle Workspace Manager, which enables applications to create workspaces and group different versions of table row values in different workspaces.
Oracle Database SQL Developer User's Guide	E12152	Provides conceptual and usage information about Oracle SQL Developer, a graphical tool that enables you to browse, create, edit, and delete (drop) database objects; run SQL statements and scripts; edit and debug PL/SQL code; manipulate and export data; migrate third-party databases to Oracle; view metadata and data in third-party databases; and view and create reports.
Oracle Database Rules Manager and Expression Filter Developer's Guide	E14919	Provides usage and reference information about how to manage, index, and evaluate conditional expressions in relational tables.
Oracle Text Application Developer's Guide	E16594	Provides information for building applications with Oracle Text.
Oracle XML DB Developer's Guide	E16659	Describes Oracle XML DB, with guidelines and examples for storing, generating, accessing, searching, validating, transforming, evolving, and indexing XML data in Oracle Database.
Oracle Database Advanced Application Developer's Guide	E17125	Explains topics that experienced application developers reference repeatedly.

Table 3–1 (Cont.) Oracle® Database Server Documentation

Documentation	Part Number	Description
Oracle Database SecureFiles and Large Objects Developer's Guide	E18294	Describes use of large object data types BLOB, CLOB, NCLOB, and BFILE in application development. Presents APIs for working with LOBs in supported programmatic environments, and SecureFiles, that are LOBs that support compression, encryption, and deduplication. Describes database file systems and the Oracle Database File System Hierarchical Store package.
Application Express		
Oracle Application Express Advanced Tutorials	E11945	Features a series of tutorials that demonstrate how to use the Oracle Application Express development environment to develop Web applications.
Oracle Database 2 Day + Application Express Developer's Guide	E11946	Offers an introduction to application development using the Oracle Application Express development environment.
Oracle Application Express Application Builder User's Guide	E11947	Functions as a comprehensive, task-based reference on using the Application Builder development environment to build and deploy database-centric Web applications.
Oracle Application Express Installation Guide	E12196	Explains how to install and configure Oracle Application Express.
Oracle Application Express Application Migration Guide	E12509	Describes how to convert applications from either Microsoft Access or Oracle Forms to Oracle Application Express.
Oracle Application Express API Reference	E12510	Describes the Application Programming Interfaces, referred to as APIs, available when programming in the Oracle Application Express environment.
Oracle Application Express SQL Workshop and Utilities Guide	E12511	Describes how to use Oracle Application Express SQL Workshop and utilities to load and unload data from an Oracle database, generate DDL, view object reports, and restore dropped database objects.
Oracle Application Express Administration Guide	E12512	Describes how to perform administration tasks for an Oracle Application Express workspace, application, or an entire development instance.
Backup and Recovery		
Oracle Database Backup and Recovery User's Guide	E10642	Provides information for the backup and recovery of Oracle databases, including RMAN backup and recovery, RMAN data transfer, Oracle Flashback Technology, and user-managed backup and recovery.
Oracle Database Backup and Recovery Reference	E10643	Provides information for the backup and recovery of Oracle databases, including RMAN backup and recovery, RMAN data transfer, Oracle Flashback Technology, and user-managed backup and recovery.
Data Warehousing		

Table 3–1   (Cont.) Oracle® Database Server Documentation

Documentation	Part Number	Description
Oracle Database 2 Day + Data Warehousing Guide	E10578	Introduces the data warehousing solutions available in Oracle Database and discusses how to perform common day-to-day tasks necessary to implement and administer a data warehouse. This guide also provides information about design tasks needed to keep the data warehouse operational, including how to perform basic performance monitoring tasks.
Oracle Database Data Warehousing Guide	E16579	Provides conceptual, reference, and implementation material for using Oracle Database 11*g* in data warehousing. It covers the full range of data warehousing activities, from physical database design to advanced calculation techniques.
Oracle Data Guard Concepts and Administration	E17022	Provides a comprehensive overview of Oracle Data Guard concepts and describes how to configure and implement standby databases that can take over production operations if your production database becomes unusable. This guide includes several database scenarios such as creating, recovering, failing over, switching over, configuring, and backing up standby and primary databases.
Oracle Data Guard Broker	E17023	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.
Language and Interface		
Oracle Call Interface Programmer's Guide	E10646	Describes C programming techniques for including SQL and PL/SQL in applications, which can be used to unlock the full capabilities of the Oracle Database.
Oracle C++ Call Interface Programmer's Guide	E10764	Presents the Oracle C++ Call Interface, an application program interface 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, programmers may want to add code in the OTT-generated file. The OTT can distinguish between OTT-generated code and user-generated code by looking for some predefined markers (tags). Support for these tags has been added in this release.
Oracle Database XML C API Reference	E10770	Describes the C APIs for XML.
Oracle Database XML C++ API Reference	E10771	Describes the C++ APIs for XML.
*Pro*C/C++ Programmer's Guide*	E10825	Describes how to develop C++ programs that use the SQL and PL/SQL database languages to access and manipulate Oracle data.

Table 3–1   (Cont.) Oracle® Database Server Documentation

Documentation	Part Number	Description
*Pro*COBOL Programmer's Guide*	E10826	Describes how to develop COBOL programs that use the database languages SQL and PL/SQL to access and manipulate Oracle data.
*Oracle SQL*Module for Ada Programmer's Guide*	E10827	Describes how to use SQL*Module to develop Module programs and Ada code to access and manipulate Oracle data.
*Pro*FORTRAN Supplement to the Oracle Precompilers Guide*	E10828	Supplements the Oracle Database Programmer's Guide to the Oracle Precompilers to provide additional information on how to develop FORTRAN programs that use the database languages SQL and PL/SQL to access and manipulate Oracle data.
Oracle Database Programmer's Guide to the Oracle Precompilers	E10830	Describes how to develop applications using the Oracle Precompilers and the SQL and PL/SQL database languages to access and manipulate Oracle data.
Networking and Security		
Oracle Database 2 Day + Security Guide	E10575	Provides an introduction to securing an Oracle database. In addition to describing how to manage security in a default Oracle database, this guide includes beginning level tutorials for Oracle Label Security and Oracle Database Vault.
Oracle Database Enterprise User Security Administrator's Guide	E10744	Explains how to configure and use enterprise users, which are users defined in a directory. The guide also covers administrative tasks related to enterprise user security.
Oracle Label Security Administrator's Guide	E10745	Explains how to configure and use Oracle Label Security, which provides row-level access control based on the virtual private database technology of Oracle Database Enterprise Edition.
Oracle Database Advanced Security Administrator's Guide	E10746	Explains how to configure and use Oracle Database Advanced Security. It covers stored data encryption, network data encryption, and strong authentication.
Oracle Database Net Services Reference	E10835	Contains a complete listing and description of the control utility commands and configuration file parameters available for managing components of Oracle Net Services.
Oracle Database Net Services Administrator's Guide	E10836	Explains how to plan, configure, and manage enterprise-wide connectivity with Oracle Net Services.
Oracle Database Security Guide	E16543	Explains how to configure an Oracle database to use the default security features.
Oracle Real Application Clusters		

Table 3–1   (Cont.) Oracle® Database Server Documentation

Documentation	Part Number	Description
Oracle Clusterware Administration and Deployment Guide	E16794	Describes the Oracle Clusterware architecture and provides an overview of this product. This book also describes administrative and deployment topics for Oracle Clusterware.
Oracle Real Application Clusters Administration and Deployment Guide	E16795	Describes the Oracle Real Application Clusters (Oracle RAC) architecture and provides an overview of this product. This book also describes administrative and deployment topics for Oracle RAC.
Oracle Database 2 Day + Real Application Clusters Guide	E17264	Describes how to install, configure, and administer Oracle Clusterware and Oracle Real Application Clusters (Oracle RAC) on a two-node system using the Oracle Linux system.
Streams and Replication		
Oracle Streams Replication Administrator's Guide	E10705	Contains instructions for configuring and managing an Oracle Streams replication environment. It also includes best practices for Oracle Streams replication environments and instructions for migrating from Advanced Replication to Oracle Streams replication.
Oracle Database Advanced Replication	E10706	Contains conceptual information about Oracle Database Advanced Replication. It also includes information about planning an Advanced Replication environment and information about troubleshooting Advanced Replication problems.
Oracle Database Advanced Replication Management API Reference	E10707	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. It also includes reference information about data dictionary views that are important for Advanced Replication and security considerations for an Advanced Replication environment.
Oracle Streams Advanced Queuing Java API Reference	E10768	Describes the Java APIs for application development and integration using the Oracle messaging system, Advanced Queuing.
Oracle Streams Advanced Queuing User's Guide	E11013	Describes features of application development and integration using Oracle Streams Advanced Queuing (AQ). This information applies to versions of the Oracle Database server that run on all platforms, unless otherwise specified.
Oracle Streams Extended Examples	E12862	Contains extended examples that configure various types of Oracle Streams environments.

Table 3–1 (Cont.) Oracle® Database Server Documentation

Documentation	Part Number	Description
Oracle Streams Concepts and Administration	E17069	Contains conceptual information about Oracle Streams and information about administering, monitoring, and troubleshooting an Oracle Streams environment. It also contains instructions for using Oracle Streams to record table changes. It also contains conceptual information about information provisioning and instructions for using Oracle Streams for information provisioning. It also includes instructions for using Oracle Streams for database upgrades and database maintenance operations.
Oracle Database 2 Day + Data Replication and Integration Guide	E17516	Provides an introduction to completing common data replication and integration tasks. It includes instructions for accessing and modifying information in multiple databases, setting up and managing an Oracle Streams replication environment, setting up and managing a materialized view replication environment, setting up and managing a messaging environment, and comparing and converging data at two databases.
XStream		
Oracle Database XStream Guide	E16545	Contains conceptual information and reference information about XStream. Also contains instructions for using XStream. The reference information includes descriptions of the OCI API for XStream, PL/SQL packages related to XStream, and data dictionary views related to XStream.

3.1.2 Oracle® Database 11g Spatial Documentation

Table 3–2   Oracle® Database Spatial Documentation

Documentation	Part Number	Description
Oracle Spatial GeoRaster Developer's Guide	E11827	Provides usage and reference information for the GeoRaster feature of Oracle Spatial, which lets you store, index, query, analyze, and deliver raster data (raster image and gridded data and its associated metadata).
Oracle Database Semantic Technologies Developer's Guide	E11828	Provides usage and reference information about Oracle Database support for semantic technologies, including storage, inference, and query capabilities for data and ontologies based on Resource Description Framework (RDF), RDF Schema (RDFS), and Web Ontology Language (OWL).
Oracle Spatial Java API Reference	E11829	Provides Javadoc-generated information about interfaces and classes that support features available with Oracle Spatial, including GeoRaster and the topology and network data models, and with MapViewer, which is packaged with Oracle Application Server and Fusion Middleware.
Oracle Spatial Developer's Guide	E11830	Provides usage and reference information for indexing and storing spatial data and for developing spatial applications using Oracle Spatial and Oracle Locator.
Oracle Spatial Topology and Network Data Models Developer's Guide	E11831	Provides usage and reference information about the topology data model and network data model capabilities of Oracle Spatial.

3.1.3 Oracle® Database 11g Windows Documentation

Table 3–3   Oracle® Database Windows Documentation

Documentation	Part Number	Description
Oracle COM Automation Feature Developer's Guide for Microsoft Windows	E10591	Provides installation, post-installation configuration, and usage information for Oracle COM Automation Feature.
Oracle Database 2 Day + .NET Developer's Guide for Microsoft Windows	E10767	Serves as a quick start guide, which describes Oracle technologies for the Microsoft .NET Framework, including the key features of Oracle Data Provider for .NET and Oracle Developer Tools for Visual Studio. It leads you through installation and configuration, shows how to build basic applications using Oracle .NET products, and how to create and use both PL/SQL and .NET stored procedures.
Oracle Provider for OLE DB Developer's Guide for Microsoft Windows	E12243	Describes Oracle Provider for OLE DB, a high performance, feature rich mechanism for Microsoft ADO and OLE DB data access to Oracle databases, its usage, and its configuration.

Table 3–3   (Cont.) Oracle® Database Windows Documentation

Documentation	Part Number	Description
Oracle Objects for OLE Developer's Guide for Microsoft Windows	E12245	Describes Oracle Objects for OLE (OO4O), a COM-based data access driver that combines seamless and optimized access to Oracle databases with easy to use interfaces, and shows how to use and configure OO4O from virtually any Microsoft COM Automation programming or scripting language, such as Visual Basic.
Oracle Services for Microsoft Transaction Server Developer's Guide for Microsoft Windows	E12247	Describes Oracle Services for Microsoft Transaction Server, which allows Oracle databases to be used as resource managers in distributed transactions. The guide describes its usage and configuration in a Microsoft Distributed Transaction Coordinator environment.
Oracle Database Extensions for .NET Developer's Guide for Microsoft Windows	E12248	Provides introductory, installation, post-installation configuration, and usage information for Oracle Database Extensions for .NET, which makes it possible to build and run .NET stored procedures or functions with Oracle Database for Microsoft Windows.
Oracle Data Provider for .NET Developer's Guide	E12249	Describes provider-specific features and properties of ODP.NET. Oracle Data Provider for .NET is an implementation of Microsoft's ADO.NET interface and provides access to Oracle databases.
Oracle Providers for ASP.NET Developer's Guide for Microsoft Windows	E12250	Provides introductory, installation, postinstallation configuration, and usage information for using these providers with ASP.NET services and controls.

3.1.4 Oracle® Database 11g Gateways Documentation

Table 3–4   Oracle® Database Gateways Documentation

Documentation	Part Number	Description
Oracle Database Gateway Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64	E12013	Provides Oracle Database Gateway installation and configuration information.
Oracle Database Gateway for DRDA User's Guide	E12014	Provides users with transparent access to DB2.
Oracle Database Gateway Installation and Configuration Guide for Microsoft Windows	E12061	Assists Oracle database administrators who are responsible for installing and configuring Oracle Database Gateways.
Oracle Database Gateway for Informix User's Guide	E12066	Describes the Oracle Database Gateway for Informix, which enables Oracle client applications to access Informix data through Structured Query Language (SQL).

Table 3–4   (Cont.) Oracle® Database Gateways Documentation

Documentation	Part Number	Description
Oracle Database Gateway for Sybase User's Guide	E12067	Describes the Oracle Database Gateway for Sybase, which enables Oracle client applications to access Sybase data through Structured Query Language (SQL).
Oracle Database Gateway for Teradata User's Guide	E12068	Describes the Oracle Database Gateway for Teradata, which enables Oracle client applications to access Teradata data through Structured Query Language (SQL).
Oracle Database Gateway for SQL Server User's Guide	E12069	Describes the Oracle Database Gateway for SQL Server, which enables Oracle client applications to access SQL Server data through Structured Query Language (SQL).
Oracle Database Gateway for ODBC User's Guide	E12070	Describes the Oracle Database Gateway for ODBC, which enables Oracle client applications to access non-Oracle systems data through Structured Query Language (SQL).
Oracle Database Gateway for APPC User's Guide	E12071	Assists in determining hardware and software requirements, installing, configuring, or administering an Oracle Database Gateway for APPC, developing applications that access remote host databases through the Oracle Database Gateway for APPC using the SNA Communication Protocol or the TCP/IP communication protocol, determining security requirements, and determining and resolving problems.
Oracle Database Gateway for IMS User's Guide	E12072	Describes the Oracle Database Gateway for ODBC, which enables Oracle client applications to access non-Oracle systems data through Structured Query Language (SQL).
Oracle Database Gateway for VSAM User's Guide	E12073	Describes the Oracle Database Gateway for Adabas, which enables Oracle client applications to access Adabas data through Structured Query Language (SQL).
Oracle Database Gateway for Adabas User's Guide	E12074	Describes the Oracle Database Gateway for Adabas, which enables Oracle client applications to access Adabas data through Structured Query Language (SQL).
Oracle Connect for IMS, VSAM, and Adabas Gateways Installation and Configuration Guide for IBM z/OS	E12075	Describes the installation and configuration of Oracle Connect for IMS, VSAM, and Adabas Gateways, a prebuilt driver to mainframe data.
Oracle Database Gateway for IMS, VSAM, and Adabas Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64	E12076	Provides Oracle Database Gateway installation and configuration information.

Table 3–4 (Cont.) Oracle® Database Gateways Documentation

Documentation	Part Number	Description
Oracle Database Gateway for IMS, VSAM, and Adabas Installation and Configuration Guide for Microsoft Windows	E12077	Assists Oracle database administrators who are responsible for installing, configuring the components of the Oracle Database Gateway for IMS, VSAM, and Adabas and for diagnosing gateway errors.
Oracle Database Gateway for APPC Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64	E12078	Provides Oracle Database Gateway installation and configuration information.
Oracle Database Gateway for APPC Installation and Configuration Guide for Microsoft Windows	E12079	Assists in determining hardware and software requirements, installing, configuring, or administering an Oracle Database Gateway for APPC, developing applications that access remote host databases through the gateway using either the SNA communication protocol or the TCP/IP for IMS Connect communication protocol, determining security requirements, and determining and resolving problems.
Oracle Database Gateway for WebSphere MQ Installation and User's Guide	E12417	Provides access to WebSphere MQ services. This gateway requires a system that is capable of running 64Ebit applications.

3.1.5 Oracle® Database 11g Communications Data Model Documentation

Table 3–5 Oracle® Database Communications Data Model Documentation

Documentation	Part Number	Description
Oracle Communications Data Model Operations Guide	E15883	Describes how you create and maintain a data warehouse based on the Oracle Communications Data Model.
Oracle Communications Data Model Reference	E15886	Describes the data model structures for Oracle Communications Data Model.
Oracle Communications Data Model Installation Guide	E15887	Describes how to install and configure the Oracle Communications Data Model.
Oracle Communications Data Model Release Notes	E15888	Describes known issues with Oracle Communications Data Model

3.1.6 Oracle® Database 11g Java Development Documentation

Table 3–6   Oracle® Database Java Development Documentation

Documentation	Part Number	Description
Oracle Database JPublisher User's Guide	E10587	Describes the JPublisher utility useful for Java programmers who want classes in their applications to correspond to SQL or PL/SQL entities or server-side Java classes.
Oracle Database Java Developer's Guide	E10588	Describes how to develop, load, and run Java applications in Oracle Database.
Oracle Database SQLJ Developer's Guide	E10590	Provides guidance to developers on SQLJ programming. This document can be read by anyone with an interest in SQLJ programming, but assumes at least some prior knowledge of Java, SQLJ, PL/SQL, and Oracle database.
Oracle Globalization Development Kit Java API Reference	E10730	Describes APIs for the Oracle Globalization Development Kit.
Oracle Database XML Java API Reference	E10769	Describes the Java APIs for XML.
Oracle Database 2 Day + Java Developer's Guide	E12137	Guides application developers on how to use Java to access and modify data in Oracle Database.
Oracle Universal Connection Pool for JDBC Developer's Guide	E12265	Provides instructions for using Oracle Universal Connection Pooling API. The API is JDBC driver agnostic.
Oracle Universal Connection Pool for JDBC Java API Reference	E12826	Provides Java API reference for Universal Connection Pool.
Oracle Database RAC FAN Events Java API Reference	E13993	Provides easy access to Oracle Fast Application Notification (FAN) events.
Oracle Database JDBC Java API Reference	E13995	Provides Java API reference for Java Database connectivity.
Oracle Database JDBC Developer's Guide	E16548	Provides guidance to developers on JDBC-based applications and applets.
Oracle Database XStream Java API Reference	E16674	Describes the Java APIs for XStream.

3.1.7 Oracle® Database 11g Release Notes Documentation

Table 3–7   Oracle® Database Release Notes Documentation

Documentation	Part Number	Description
Oracle Warehouse Builder Release Notes	E10585	Contains late-breaking information about corrections, troubleshooting, and known issues for DB 11.2.0.2.
Oracle Application Express Release Notes	E11944	Contains important information not included in the Oracle Application Express documentation.
*SQL*Plus Release Notes*	E16608	Contains important information not included in the SQL*Plus User's Guide and Reference.
Oracle Database Release Notes for Linux	E16778	Contains important information not included in the Oracle Database documentation for Linux.

3.1.8 Oracle® Database 11g Data Mining Documentation

Table 3–8   Oracle® Database Data Mining Documentation

Documentation	Part Number	Description
Oracle Data Mining Application Developer's Guide	E12218	Discusses the PL/SQL and Java interfaces to Oracle Data Mining and the SQL functions for deploying Data Mining models.
Oracle Data Mining Java API Reference	E12219	Describes the classes and methods in the Oracle Data Mining Java API, the JDM JSR-73-compliant API for data mining. The Java API supports a full range of data mining activities, including model building and scoring, data preparation, and import/export of models.
Oracle Data Mining Administrator's Guide	E16807	Documents the installation, upgrade, and configuration of the Data Mining option.
Oracle Data Mining Concepts	E16808	Discusses the basic concepts underlying Oracle Data Mining. Includes information about mining functions, algorithms, data preparation, and predictive analytics.

3.1.9 Oracle® Database 11g TimesTen In-Memory Database Documentation

Table 3–9   Oracle® Database TimesTen In-Memory Database Documentation

Documentation	Part Number	Description
Oracle TimesTen In-Memory Database Installation Guide	E13063	Provides information about installing and upgrading the Oracle TimesTen In-Memory Database
Oracle TimesTen In-Memory Database Operations Guide	E13065	Provides information about how Oracle TimesTen In-Memory Database works. It includes step-by-step instructions and examples that show how to perform commonly needed tasks.
Oracle TimesTen In-Memory Database C Developer's Guide	E13066	Provides information about developing C applications to access the Oracle TimesTen In-Memory Database. It also includes information about TimesTen support for OCI and the Pro*C/C++ Precompiler.
Oracle TimesTen In-Memory Database Java Developer's Guide	E13068	Provides information about developing Java applications to access the Oracle TimesTen In-Memory Database
Oracle TimesTen In-Memory Database Reference	E13069	Provides a reference for attributes, built-in procedures and utilities used with the Oracle TimesTen In-Memory Database.
Oracle TimesTen In-Memory Database SQL Reference	E13070	Provides a reference for SQL statements, expressions, and functions used with the Oracle TimesTen In-Memory Database, including TimesTen SQL extensions. It also describes data types.
Oracle TimesTen In-Memory Database Error Messages and SNMP Traps	E13071	Provides a reference for Oracle TimesTen In-Memory Database error messages and SNMP traps
Oracle TimesTen In-Memory Database TimesTen to TimesTen Replication Guide	E13072	Provides information about designing and configuring TimesTen to TimesTen replication, including information about TimesTen integration with Oracle Clusterware
Oracle In-Memory Database Cache User's Guide	E13073	Provides information about designing and configuring the Oracle In-Memory Database Cache. It includes information about setting up a cache grid.
Oracle TimesTen In-Memory Database TTClasses Guide	E13074	Describes how to use the TimesTen C++ Interface Classes (TTClasses) library
Oracle TimesTen In-Memory Database Troubleshooting Procedures Guide	E13075	Provides information about utilities and procedures for troubleshooting problems that may be encountered when using the Oracle TimesTen In-Memory Database
Oracle TimesTen In-Memory Database PL/SQL Developer's Guide	E13076	Provides information about using PL/SQL with the Oracle TimesTen In-Memory Database
Oracle TimesTen In-Memory Database Release Notes	E13080	Provides information about changes in the latest release of Oracle TimesTen In-Memory Database

Table 3–9   (Cont.) Oracle® Database TimesTen In-Memory Database Documentation

Documentation	Part Number	Description
Oracle Enterprise Manager System Monitoring Plug-in for Oracle TimesTen In-Memory Database Installation Guide	E13081	Describes how to install the Oracle TimesTen In-Memory Database Plug-in for Oracle Enterprise Manager
Oracle Enterprise Manager System Monitoring Plug-in for Oracle TimesTen In-Memory Database Release Notes	E13082	Provides late-breaking information about the Oracle TimesTen In-Memory Database Plug-in for Oracle Enterprise Manager
Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference	E13083	Describes TimesTen extensions to JDBC classes and interfaces
Oracle TimesTen In-Memory Database JMS/XLA Java API Reference	E13084	Describes the TimesTen JMS/XLA package
Oracle TimesTen In-Memory Database PL/SQL Packages Reference	E14000	Provides reference information for packages provided with TimesTen PL/SQL
Oracle In-Memory Database Cache Introduction	E14261	Describes the features of Oracle In-Memory Database Cache and provides information to help developers plan an IMDB Cache application
Oracle TimesTen In-Memory Database System Tables and Limits Reference	E17114	Describes system tables and replication tables. It also has a list of system limits and defaults.

3.1.10 Oracle® Database 11g Documentation Library

Table 3–10   Oracle® Database Documentation Library

Documentation	Part Number	Description
Oracle Database Licensing Information	E10594	Describes Oracle Database licensing. If you have a question about your licensing needs, please contact your Oracle sales representative or refer to the resources listed in "Related Documents" for more information.
Oracle Database Readme	E17129	Describes important last-minute features and changes not included in Oracle Database Documentation Library.

3.1.11 Oracle® Database 11g Warehouse Builder Documentation

Table 3–11   Oracle® Database Warehouse Builder Documentation

Documentation	Part Number	Description
Oracle Warehouse Builder Help	B28692	Provides online versions of the complete Oracle Warehouse Builder documentation library, and context-sensitive help for all user-interface objects within OWB. This help system is available online within the OWB Help Center.
Oracle Warehouse Builder Concepts	E10581	Provides an introduction to the benefits of using Oracle Warehouse Builder and its architecture and components. It also includes "quick start" steps, a user interface tour, and provides information about OWB preferences and OWB editors and wizards. Contains fixes for 11.2.0.2.
Oracle Warehouse Builder Sources and Targets Guide	E10582	Describes data sources and targets for Oracle Warehouse Builder data warehouses and data integration systems. This guide also lists all the supported sources and targets and provides procedures for importing from sources and deploying to targets. Includes bug fixes for 11.2.0.2.
Oracle Warehouse Builder API and Scripting Reference	E10584	Describes the application programming interfaces and scripting language available with Oracle Warehouse Builder. It also includes information about creating Experts and Guided Assistants.
Oracle Warehouse Builder Data Modeling, ETL, and Data Quality Guide	E10935	Provides comprehensive procedures for loading data from sources, designing target schemas, transforming data, generating code, and performing all the tasks for optimizing and managing data quality in a data management system. Includes bug fixes for 11.2.0.2.
*Oracle Warehouse Builder OMB*Plus Command Reference*	E14406	Provides a complete language reference for OMB and OMU commands, which comprise OMB*Plus. OMB*Plus is the programming language included with Oracle Warehouse Builder.
Oracle Warehouse Builder Installation and Administration Guide for Windows and Linux	E17130	Describes installing the software and configuring Oracle Warehouse Builder (OWB). Part II of this guide provides procedures for administrating OWB.

3.1.12 Oracle® Database 11g OLAP Documentation

Table 3–12   Oracle® Database OLAP Documentation

Documentation	Part Number	Description
Oracle OLAP User's Guide	E17123	Explains how SQL applications can extend their analytic processing capabilities and manage summary data by using the OLAP option of Oracle Database.
Oracle OLAP Customizing Analytic Workspace Manager	E17237	Describes the Analytic Workspace Manager XML and Java plug-in interfaces and provides simple examples.
Oracle OLAP DML Reference	E17122	Contains a complete description of the OLAP Data Manipulation Language (OLAP DML) used to define and manipulate analytic workspace objects.
Oracle OLAP Java API Reference	E10794	Provides reference information for the Oracle OLAP Java API.
Oracle OLAP Java API Developer's Guide	E10795	Provides guidance for developing client applications using the Oracle OLAP Java API.

3.1.13 Oracle® Database 11g Multimedia Documentation

Table 3–13   Oracle® Database Multimedia Documentation

Documentation	Part Number	Description
Oracle Multimedia Reference	E10776	Presents reference information about using Oracle Database to store, manage, and retrieve images, audio, video, DICOM format medical images and other objects, or other heterogeneous media data in an integrated fashion with other enterprise information. (Oracle Multimedia was formerly known as Oracle *inter*Media.)
Oracle Multimedia User's Guide	E10777	Presents information about using Oracle Database to store, manage, and retrieve images, audio, video, DICOM format medical images and other objects, or other heterogeneous media data in an integrated fashion with other enterprise information. (Oracle Multimedia was formerly known as Oracle *inter*Media.)
Oracle Multimedia DICOM Developer's Guide	E10778	Presents reference and usage information about using Oracle Database to store, manage, and retrieve DICOM content such as single-frame and multiframe images, waveforms, slices of 3-D volumes, video segments, and structured reports in an integrated fashion with other enterprise information. (Oracle Multimedia was formerly known as Oracle *inter*Media.)

Table 3–13   (Cont.) Oracle® Database Multimedia Documentation

Documentation	Part Number	Description
Oracle Multimedia Java API Reference	E10779	Presents reference information about writing Java applications using Oracle Multimedia objects to manage multimedia content within Oracle Database. (Oracle Multimedia was formerly known as Oracle *inter*Media.)
Oracle Multimedia Servlets and JSP Java API Reference	E10780	Presents reference information about writing Java applications using Java servlets and JavaServer Pages (JSP) to retrieve and upload multimedia data from and to Oracle Database. (Oracle Multimedia was formerly known as Oracle *inter*Media.)
Oracle Multimedia DICOM Java API Reference	E10781	Presents reference information about writing Java applications using the Oracle Multimedia object designed to store Digital Imaging and Communications in Medicine (DICOM) data within Oracle Database. (Oracle Multimedia was formerly known as Oracle *inter*Media.)
Oracle Multimedia Mid-Tier Java API Reference	E17343	Presents reference information about writing Java applications using Oracle Multimedia to extract Digital Imaging and Communications in Medicine (DICOM) metadata outside of the database, either on the client or in the middle tier. (Oracle Multimedia was formerly known as Oracle *inter*Media.)

3.1.14 Oracle® Database 11g Installation Guide Documentation

Table 3–14   Oracle® Database Installation Guide Documentation

Documentation	Part Number	Description
Oracle Database Installation Guide for Linux	E16763	Provides configuration information for network and system administrators, and database installation information for database administrators (DBAs) who install and configure Oracle Database and Grid Infrastructure for a Single Server on Linux.
Oracle Database Client Installation Guide for Linux	E16765	Describes how to install and configure Oracle Database Client on Linux.
Oracle Grid Infrastructure Installation Guide for Linux	E17212	Describes how to install Oracle Grid Infrastructure for a Cluster software, which includes Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM), on systems running Linux.
Oracle Real Application Clusters Installation Guide for Linux and UNIX	E17214	Describes how to install and configure Oracle Real Application Clusters (Oracle RAC) on systems running Linux and UNIX.
Oracle Database Examples Installation Guide	E17861	Describes how to install and configure the products available on the Oracle Database Examples installation media.