

New Features of Oracle SQL

In this chapter:

- What are the new features of Oracle SQL in Oracle Database 10g?
- What were the new features of Oracle SQL in Oracle Database 9i?
- What PL/SQL improvements are there?
- How is XML development better supported?
- What's new in Oracle SQL utilities?

This chapter takes a bird's-eye view of Oracle SQL changes in both Oracle Database 10g and Oracle Database 9i. Without further ado, let's get started with Oracle Database 10g.

2.1 New Features in Oracle Database 10g

Oracle Database 10g contains the following SQL and PL/SQL features.

2.1.1 Oracle SQL Improvements in Oracle Database 10g

- Oracle documentation states that case sensitivity is no longer required for filtering and sorting in SQL statements. Proving this point is a tuning exercise and does not belong in this book.
- The CONNECT BY clause now allows ancestor-descendant pairs as opposed to only parent-child pairs. In other words, pairs can be matched and returned where those pairs are not directly related within a hierarchy but related from the top to the bottom of a hierarchy (see Chapter 13).

- Object improvements include VARRAY resizing and splitting of nested table type columns into different tablespaces (see Chapter 16).
- A new row timestamp pseudocolumn called `ORA_ROWSCN` contains a commit point timestamp or system change number (SCN). For updates only, the SCN for a row must be retrieved to ensure that no row change occurred between a row `SELECT` and subsequent `UPDATE`.
- The following new datatypes have been added (see Chapter 16):
 - `BINARY_FLOAT` and `BINARY_DOUBLE` allow 32-bit single precision and 64-bit double precision floating-point numbers.
 - `SDO_GEORASTER` and `SI_STILLIMAGE` store raster and digital images, respectively (including object characteristics), for object-relational multimedia storage.
- A multitude of DDL commands have been altered and enhanced. Most DDL command changes are relevant to database administration and not Oracle SQL.
- Several `SELECT` statement and DML command syntax changes have been introduced:
 - `MERGE` allows insertions, updates, or both. Previously, the `MERGE` command always performed both insertions and updates. Additionally, `MERGE` can also delete rows from the target table (see Chapter 15).
 - `SELECT` can be executed as a flashback or versions query, retrieving data at a point in time in the past, based on an SCN or timestamp (see Chapter 13).
 - Grouped outer joins allow groupings on data where subset parts may not exist.

Note: (10g) Grouped outer joins are omitted from this book because syntax documentation was not available at the time of writing.

- The `SPREADSHEET` clause extends the `SELECT` statement, allowing multiple dimensional array query result output. Calculations between resulting rows can be performed much like cross-tabbing or interdimensional data warehouse reporting (see Chapter 11).
-

Note: (10g) The SPREADSHEET clause has been renamed to the MODEL clause.

- Oracle Database 10g has recycle bin technology (see Chapter 18):
 - Recovering a table from the recycle bin requires use of the FLASHBACK TABLE command.
 - The PURGE command is used to permanently destroy objects dropped into the recycle bin. Space is not released for dropped objects until they are purged.
- New built-in functions are as follows:
 - COLLECT creates a nested table from the row set result of a single column in a table (see Chapter 16).
 - Nested tables have new collection functions (see Chapter 16):
 - CARDINALITY returns the number of elements in a collection for each row.
 - POWERMULTISET returns all set elements in a collection.
 - POWERMULTISET_BY_CARDINALITY combines the previous two functions by returning all set elements with a specified number of entries, for each collection in each row. One could find every row in a table where that collection has a specified number of entries.
 - SET converts a nested table (collection) for each row to a set. A set contains unique values only such that duplicates are removed within each collection in each row.
 - SPREADSHEET clause functions are used to facilitate cross calculations between different rows. Functions include the following (see Chapter 11):
 - CURRENTV returns a dimensional value or current value.
 - PRESENTNNV returns one expression if a value exists, otherwise another.
 - PRESENTV is as for PRESENTNNV except allowing null values.
 - PREVIOUS returns a value at the beginning of each iteration or loop.
 - New binary floating-point number functions include the following (see Chapter 9):
 - TO_BINARY_DOUBLE and TO_BINARY_FLOAT allow for conversions.

- NANVL returns a replacement value if the initial value is not a number.
- REMAINDER is a remainder or modulus function specifically for binary floating-point numbers.
- Regular expression functions are REGEXP_INSTR, REGEXP_REPLACE, and REGEXP_SUBSTR. These functions essentially expand the search-and-replace capabilities of INSTR, REPLACE, and SUBSTR to full pattern-matching regular expression capabilities. For more information, refer to the section titled “Expressions” (on page 44) in this chapter and see Chapter 14.

Note: The essential difference between simple pattern matching and regular expression matching is that simple pattern matching searches for patterns. Regular expression matching searches for patterns allowing for replacement and return of specific values found within a pattern.

- New statistical aggregation functions are covered in detail in Chapter 11. Statistical analysis can be extremely useful in data warehouse and reporting environments. There is now extensive capability in Oracle SQL for OLAP-type inter-row and cross-row analysis, which was previously only available in expensive software packages and add-ons such as Formula1.
 - CORR_{S | K} calculate Pearson’s correlation coefficient, measuring the strength of a linear relationship between two variables. Plotting two variables on a graph results in a lot of dots plotted from two axes. Pearson’s correlation coefficient can tell you how good the straight line is.
 - MEDIAN returns a median, middle, or interpolated value. Quite literally, a median is the middle sequenced value in a set of values. If a distribution is discontinuous and skewed or just all over the place, the median will not be anywhere near a mean or average of a set of values. A median is not always terribly useful.
 - Other statistical functions begin with “STATS.” The syntax appears like this:

```
STATS_{BINOMIAL_TEST | CROSSTAB | F_TEST | KS_TEST |
MODE | MW_TEST | ONE_WAY_ANOVA | STATS_T_TEST_* |
STATS_WSR_TEST}
```

These functions provide various statistical goodies.

- The `ORA_HASH` function returns a hash value for an expression.
- New and enhanced operators are as follows:
 - Collections (nested tables and `VARRAY` objects) can now be compared using equality (`=`) and inequality (`<>` | `!=`) operators.
 - `CONNECT_BY_ROOT` helps extend hierarchical queries from parent-child connections only, on to root and below connections (see Chapter 13).
 - Multiset operators `MULTISET {EXCEPT | INTERSECT | UNION}` combine results of two collections. `EXCEPT` is similar to the outer part of an outer join, including all elements in one collection and not another. `INTERSECT` is the intersection of two collections (the unique list of common values). `UNION` combines all elements in both collections (see Chapters 7 and 16).
- New pseudocolumns are as follows (see Chapter 7):
 - Hierarchical pseudocolumns `CONNECT_BY_{ISLEAF | ISCYCLE}` give an indication of contained child elements in a hierarchy.
 - Version query pseudocolumns provide versioning information for flashback version queries.
- New conditional operators are as follows (see Chapter 7):
 - Floating-point conditions `IS [NOT] {NAN | INFINITE}` allow undefined and infinite checks against floating-point number expressions.
 - `IS [NOT] A SET` implies that a collection is a set because it contains unique values only.
 - `IS ANY` qualifies `SPREADSHEET` clause dimensional values.
 - `IS [NOT] EMPTY` checks for an empty collection, a nested table containing no elements whatsoever, essentially a collection not as yet instantiated.
 - `IS PRESENT` ensures that a cell exists before the execution of a `SPREADSHEET` clause.
 - `[NOT] MEMBER OF collection` attempts to validate membership within a collection.
 - `REGEXP_LIKE` utilizes regular expressions as opposed to simple pattern matching.
 - `SUBMULTISET` indicates if one or more collection items are a subset of another collection. See Chapters 7 and 16 for details on collections such as nested tables and `VARRAY` objects.

- Expressions and the new EVALUATE operator permit what would previously have been multiple-line SQL statements to be placed into a single line of SQL or PL/SQL code. The term used by Oracle documentation is “describing user’s interest in data” (see Chapter 14).
 - The Oracle Expression Filter uses an Expression datatype and the EVALUATE operator. The EVALUATE operator allows concise conditional expression evaluation.
 - Portable operating system interface (POSIX) standard regular expression capabilities allow search-and-replace functionality with changes to the LIKE operator, REPLACE, and INSTR functions. This search-and-replace capability is equivalent to search-and-replace power in Unix scripting languages or something like the SED editor or when using Perl.

2.1.2 PL/SQL Improvements in Oracle Database 10g

An intense examination of the details of PL/SQL is not required in an SQL reference-type book, so some of the items listed here are not covered in this book. However, a basic introduction to PL/SQL programming is covered in Chapter 24.

- Everything possible in Oracle SQL with respect to SQL coding can now be coded and executed from within PL/SQL. PL/SQL is now fully syntactically equivalent with Oracle SQL. In other words, all Oracle SQL commands can be coded into PL/SQL scripts.
- The PL/SQL compiler is better optimized including bulk binding and native compilation. Native compilation stores PL/SQL units in BLOB objects as a compiled binary form. Previously, PL/SQL was interpreted PL/SQL code. Interpretation implies compilation or conversion to binary at run-time; in the case of PL/SQL, coded commands were read and parsed for every execution. Binary compilation simply executes binary code at run-time and is therefore potentially much quicker to execute.
- Using binary datatypes can help number-crunching performance.

Note: Number crunching or heavily computational code should not really be constructed using a language such as PL/SQL. Java or even C is better suited, and most commonly at the application level.

- Extensive collection set operation capability encapsulates collection testing and verification coding into single commands. See the previous section on Oracle SQL improvements in Oracle Database 10g.
- The PLSQL_WARNINGS database configuration parameter or the DBMS_WARNINGS package can be used to enable or disable PL/SQL compilation warnings.
- PL/SQL quoting of strings within strings no longer requires the use of multiple sets of single quotation marks. A string delimiter character can now be specified.
- The collection iteration FORALL statement is improved.
- SCN_TO_TIMESTAMP and TIMESTAMP_TO_SCN functions can help with setting up flashback queries.
- The packages UTL_COMPRESS and UTL_MAIL are new. The UTL_COMPRESS package allows data compression. The UTL_MAIL package simplifies e-mail from within PL/SQL, where underlying protocol detail is not required.

2.1.2.1 Java Improvements in Oracle Database 10g

The Oracle Database kernel JVM is improved in Oracle Database 10g for compliance with the latest version of Java, driver enhancements, connection caching, and passing of parameters by name for PL/SQL, among various other improvements. A discussion of the Oracle kernel JVM is a topic in itself that is beyond the scope of this book.

2.1.3 XML Improvements in Oracle Database 10g

New operators can be used to convert between XML and SQL, allowing creation of highly complex XML object document structures, and storage of those XML documents. See the later section on Oracle SQL improvements in Oracle Database 9i for a synopsis of XML functionality. Using XML in Oracle SQL is covered in detail in Chapter 17.

2.1.4 Some Utility Improvements in Oracle Database 10g

- SQL*Plus:
 - The SPOOL [CREATE | REPLACE | APPEND] options enhance the SPOOL command in SQL*Plus.

- SET SQLPROMPT can be set to values such as a schema and the name of a database server.
- The login script in the \$ORACLE_HOME/sqlplus/admin directory GLOGIN.SQL is now executed for every database connection, not only on opening the SQL*Plus utility.
- Contents of the recycle bin can be viewed.
- DBMS_OUTPUT functionality is more easily provided.
- iSQL*Plus now allows prompts for input values.

Now let's look at changes made to various database objects, those directly related to Oracle SQL, not database administration.

2.1.5 Database Object Improvements in Oracle 10g

- Tables can be purged such that they are dropped without being stored in the recycle bin. Thus the DROP TABLE command now has a PURGE clause, and a new command called the PURGE command has been introduced (see Chapter 18).
- The FLASHBACK TABLE command can be used to restore a previous version of a table back to an SCN or timestamp, perhaps even recover a mistakenly dropped table (see Chapter 18).
- These next two changes are interesting but more applicable to general database administration than to Oracle SQL specifically:
 - Multiple temporary tablespaces using tablespace groups can now be set for a user (schema) within the CREATE USER command syntax.
 - Nested table and VARRAY types can now be changed.

Now let's go backward in time and make a quick synopsis of Oracle SQL and PL/SQL features introduced in Oracle Database 9i, perhaps putting some of the changes for Oracle Database 10g into perspective.

2.2 New Features in Oracle Database 9i

Oracle Database 9i (Release 1 or Release 2) contained the following new features for SQL and PL/SQL.

2.2.1 Oracle SQL Improvements in Oracle Database 9i

- A new data warehousing command called MERGE.
- New features for data warehousing such as CUBE and ROLLUP GROUP BY extensions and the RANK function. These features are for use in summary and subtotal reports generated with SQL commands.
- Columns and constraints can be renamed.
- A query of data at a point in time in the past (called a Flashback query) is enhanced for use within an SQL command rather than requiring an environmental session change.
- A new datatype called TIMESTAMP supports timezone-sensitive dates and times.
- Nearly full support of SQL*Plus features in iSQL*Plus so you can create Web-based reports.
- Support for the ANSI standard JOIN command among other new ANSI standards.
- XML and partitioning enhancements when creating tables and views.
- XML functional capability enhancements allowing various new types of operations on XML datatypes:
 - XML conditions:
 - EQUALS_PATH tries to find something in an XML path.
 - UNDER_PATH finds something within an XML path.
 - DEPTH and PATH are ancillary functions of EQUALS_PATH and UNDER_PATH conditions. DEPTH is the number of levels within a path and PATH is a relative path specifier.

Note: An ancillary function is a subordinate part of whatever uses it. For example, the DEPTH function can only be used with EQUALS_PATH and is thus ancillary to the EQUALS_PATH function.

- XML functions:
 - DEPTH and PATH (see previous explanation).
 - Working with XML document objects:
 - EXTRACTVALUE returns a scalar value from an XML document object node.

- UPDATEXML returns an XML document object including a change.
- Generating XML from Oracle SQL code:
 - XMLAGG aggregates or merges multiple XML pieces from an SQL row set.
 - XMLCONCAT concatenates XML values from an SQL row set into XML elements.
 - XMLCOLATTVAL collates XML values into XML structural elements from row sets.
 - XMLELEMENT creates an XML element from an SQL row set.
 - XMLFOREST creates an XML hierarchical structure from an SQL row set with an object for each row, and a name-value pair (element) for each column value within that row.
- Returning data from XML documents:
 - XMLSEQUENCE returns an array of elements at a specific level or path (row identifier) within an XML document.
 - XMLTRANSFORM simply applies an XSL style sheet to an XML document object.

Note: XML capabilities with Oracle SQL are covered in Chapter 17.

- Expression capabilities have been added to include CASE statements, cursor expressions using CURSOR (subquery) syntax, and scalar subqueries returning an expression.
- Numerous new functions and enhancements to existing functions.
- Privilege and DDL command enhancements, most of which fall outside the scope of this book.

2.2.2 New PL/SQL Features in Oracle Database 9i

As already stated, an intense examination of the details of PL/SQL is not required in an SQL reference book. Items in this list may or may not be covered in this book in later chapters:

- A number of object-handling enhancements, a little beyond the scope of this book.
 - Most SQL syntax is supported in PL/SQL.
-

- CASE statement expression, as already mentioned.
- Performance is improved by native PL/SQL code compilation as opposed to direct interpretation.
- Temporary tables and cursors are no longer required to pass structured expressions between functions. Now a query can be executed against a returned set of rows.
- Bulk SQL operations can be executed using the EXECUTE IMMEDIATE command.

Note: Some Oracle Database 10g items listed in this chapter are not covered in other chapters of this book. Some things are too obscure. Any omissions are deliberate. Details can be found in Oracle documentation.

The next chapter introduces the basics of Oracle Database physical architecture. A general understanding of underlying architecture is essential to a thorough understanding of Oracle SQL.

