

Database objects

The following table describes the database objects displayed in DB Optimizer and contains information regarding each one, including object name, DBMS platform, and any notes pertaining to the specified object.

In DB Optimizer, database objects are stored in Data Source Explorer as subnodes of individual, pertinent databases.

Database Object	DBMS Platforms	Notes
Aliases	DB2	An alias is an alternate name that references a table, view, and other database objects. An alias can also reference another alias as long as the aliases do not reference one another in a circular or repetitive manner. Aliases are used in view or trigger definitions in any SQL statements except for table check-constraint definitions. (The table or view name must be referenced in these cases.) Once defined, an alias is used in query and development statements to provide greater control when specifying the referenced object. Aliases can be defined for objects that do not exist, but the referenced object must exist when a statement containing the alias is compiled. Aliases can be specified for tables, views, existing aliases, or other objects. Create Alias is a command available on the shortcut menu.
Check Constraints	All	A check constraint is a search condition applied to a table. When a check constraint is in place, Insert and Update statements issued against the table will only complete if the statements pass the constraint rules. Check constraints are used to enforce data integrity when it cannot be defined by key uniqueness or referential integrity restraints. A check condition is a logical expression that defines valid data values for a column.
Clusters	Oracle	A cluster is a collection of interconnected, physical machines used as a single resource for failover, scalability, and availability purposes. Individual machines in the cluster maintain a physical host name, but a cluster host name must be specified to define the collective as a whole. To create a cluster, you need the CREATE CLUSTER or CREATE ANY CLUSTER system privilege.
Database Links	Oracle	A database link is a network path stored locally, that provides the database with the ability to communicate with a remote database. A database link is composed of the name of the remote database, a communication path to the database, and a user ID and password (if required). Database links cannot be edited or altered. To make changes, drop and re-create.
Foreign Keys	All	A foreign key references a primary or unique key of a table (the same table the foreign key is defined on, or another table and is created as a result of an established relationship). Its purpose is to indicate that referential integrity is maintained according to the constraints. The number of columns in a foreign key must be equal to the number of columns in the corresponding primary or unique key. Additionally, the column definitions of the foreign key must have the same data types and lengths. Foreign key names are automatically assigned if one is not specified.
Functions	DB2, Oracle	A function is a relationship between a set of input data values and a set of result values. For example, the TIMESTAMP function passes input data values of type DATE and TIME, and the result is TIMESTAMP. Functions can be built-in or user-defined. Built-in functions are provided with the database. They return a single value and are part of the default database schema. User-defined functions extend the capabilities of the database system by adding function definitions (provided by users or third-party vendors) that can be applied in the database engine itself. A function is identified by its schema, a function name, the number of parameters, and the data types of its parameters. Access to functions is controlled through the EXECUTE privilege. GRANT and REVOKE statements are used to specify who can or cannot execute a specific function or set of functions.
Groups	All	Groups are units that contain items. Typically, groups contain the result of a single business transaction where several items are involved. For example, a group is the set of articles bought by a customer during a visit to the supermarket.

Indexes	All	An index is an ordered set of pointers to rows in a base table. Each index is based on the values of data in one or more table columns. An index is an object that is separate from the data in the table. When an index is created, the database builds and maintains it automatically. Indexes are used to improve performance. In most cases, access to data is faster with an index. Although an index cannot be created for a view, an index created for the table on which a view is based can improve the performance of operations on that view. Indexes are also used to ensure uniqueness. A table with a unique index cannot have rows with identical keys. DB2: Allow Reverse Scans, Percent Free (Lets you type or select the percentage of each index page to leave as free space when building the index, from 0 to 99), Min Pct Used (Lets you type or select the minimum percentage of space used on an index leaf page. If, after a key is removed from an index leaf page, the percentage of space used on the page is at or below integer percent, an attempt is made to merge the remaining keys on this page with those of a neighboring page. If there is sufficient space on one of these pages, the merge is performed and one of the pages is deleted. The value of integer can be from 0 to 99. Oracle: The Logging, No Sort, Degrees, and Instances properties are documented in the editor.
Java Classes	Oracle	A model or template, written in Java language, used to create objects with a common definition and common properties, operations and behavior. Java classes can be developed in Eclipse (or another Java development environment such as Oracle JDeveloper) and moved into an Oracle database to be used as stored procedures. Java classes must be public and static if they are to be used in this manner. When writing a class to be executed within the database, you can take advantage of a special server-side JDBC driver. This driver uses the user's default connection and provides the fastest access to the database. Java classes become full-fledged database objects once migrated into the database via the loadjava command-line utility or the SQL CREATE JAVA statement. A Java class is published by creating and compiling a call specification for it. The call spec maps a Java method's parameters and return type to Oracle SQL types. Once a Java class is developed, loaded, and published – the final step is to execute it.
Java Resources	Oracle	A Java resource is a collection of files compressed in a .jar file.
Libraries	Oracle	A library is a configurable folder for storing and sharing content with an allocated quota. Multiple libraries may exist in the same database environment. A library is a special type of folder in Oracle Content Services. Unlike Containers and regular folders, each library has a Trash Folder and an allocated amount of disk space. A library is composed of a name (mandatory), description, quota, path, and library members. The library service allows you to create folders, list quotas, and manage categories, workflow, trash folders, and versioning. The Library service does not allow you to create or upload files.
Materialized Views	Oracle	A database object that contains the results of a query. They are local copies of data located remotely, or are used to create summary tables based on aggregations of table data. Materialized views are also known as snapshots. A materialized view can query tables, views, and other materialized views. Collectively, these are called master tables (a replication term) or detail tables (a data warehouse term). For replication purposes, materialized views allow you to maintain copies of remote data on your local node. These copies are read-only. If you want to update the local copies, you need to use the Advanced Replication feature. You can select data from a materialized view as you would from a table or view. For data warehousing purposes, the materialized views commonly created are aggregate views, single-table aggregate views, and join views.

Materialized View Logs	Oracle	Because Materialized Views are used to return faster queries (a query against a materialized view is faster than a query against a base table because querying the materialized view does not query the source table), the Materialized View often returns the data at the time the view was created, not the current table data. There are two ways to refresh data in Materialized Views, manually or automatically. In a manual refresh, the Materialized View is completely wiped clean and then repopulated with data from the source tables (this is known as a complete refresh). If source tables have changed very little, however, it is possible to refresh the Materialized View only for changed records – this is known as a fast refresh. In the case of Materialized Views that are updated via fast refresh, it is necessary to create Materialized View Logs on the base tables that compose the Materialized View to reflect the changes. If the number of entries in this table is too high, it is an indication that you might need to refresh the Materialized Views more frequently to ensure that each update does not take longer than it needs. Select owner, then select from tables with Materialized Views, etc.
Oracle Job Queue	Oracle	The Oracle Job Queue allows for the scheduling and execution of PL/SQL stored procedures at predefined times and/or repeated job execution at regular intervals, as background processes. For example, you could create a job in the Oracle Job Queue that processed end-of-day accounting – a job that must run every weekday, but can be run unattended, or you could create a series of jobs that must be run sequentially – such as jobs that might be so large, that in order to reduce CPU usage, only one is run at a time. Runs PL/SQL code at specified time or on specified schedule, can enable/disable.
Outlines	Oracle	Oracle preserves the execution plans of "frozen" access paths to data so that it remains constant despite data changes, schema changes, and upgrades of the database or application software through objects named stored outlines. Outlines are useful for providing stable application performance and benefit high-end OLTP sites by having SQL execute without having to invoke the cost-based optimizer at each SQL execution. This allows complex SQL to be executed without the additional overhead added by the optimizer when it performs the calculations necessary to determine the optimal access path to the data.
Packages	All	A package is a procedural schema object classified as a PL/ SQL program unit that allows the access and manipulation of database information. A package is a group of related procedures and functions, together with the cursors and variables they use, stored together in the database for continued use as a unit. Similar to standalone procedures and functions, packaged procedures and functions can be called explicitly by applications or users. DB applications explicitly call packaged procedures as necessary with privileges granted, a user can explicitly execute any of the procedures contained in it. Packages provide a method of encapsulating related procedures, functions, and associated cursors and variables together as a unit in the database. For example, a single package might contain two statements that contain several procedures and functions used to process banking transactions. Packages allow the database administrator or application developer to organize similar routines as well as offering increased functionality and database performance. Packages provide advantages in the following areas: encapsulation of related procedures and variables, declaration of public and private procedures, variables, constraints and cursors, separation of the package specification and package body, and better performance. Encapsulation of procedural constructs in a package also makes privilege management easier. Granting the privilege to use a package makes all constructs of the package assessable to the grantee. The methods of package definition allow you to specify which variables, cursors, and procedures are: public, directly accessible to the users of a package, private, or hidden from the user of the package.
Package Bodies	Oracle	A package body is a package definition file that states how a package specification will function. In contrast to the entities declared in the visible part of a package, the entities declared in the package body are only visible within the package body itself. As a consequence, a package with a package body can be used for the construction of a group of related subprograms in which the logical operations available to clients are clearly isolated from the internal entities.
Primary Keys	All	A key is a set of columns used to identify or access a row or rows. The key is identified in the description of a table, index, or referential constraint. The same column can be part of more than one key. A unique key is a key that is constrained so that no two of its values are equal. The columns of a unique key cannot contain NULL values. The primary key is one of the unique keys defined on a table, but is selected to be the key of the first importance. There can only be one primary key on a table. Oracle: If an index constraint has been defined for a table, the constraint status for the table's primary key cannot be set to Disabled.

Procedures	All	A procedure is an application program that can be started through the SQL CALL statement. The procedure is specified by a procedure name, which may be followed by arguments enclosed within parenthesis. The argument or arguments of a procedure are individual scalar values, which can be of different types and can have different meanings. The arguments can be used to pass values into the procedure, receive return values from the procedure, or both. A procedure, also called a stored procedure, is a database object created via the CREATE PROCEDURE statement that can encapsulate logic and SQL statements. Procedures are used as subroutine extensions to applications, and other database objects that can contain logic. When a procedure is invoked in SQL and logic within a procedure is executed on the server, data is only transferred between the client and the database server in the procedure call and in the procedure return. If you have a series of SQL statements to execute within a client application, and the application does not need to do any processing in between the statements, then this series of statements would benefit from being included in a procedure.
Profiles	Oracle	Profiles are a means to limit resources a user can use by specifying limits on kernel and password elements. Additionally, Profiles can be used to track password histories and the settings of specific profiles may be queried. The following kernel limits may be set: maximum concurrent sessions for a user, CPU time limit per session, maximum connect time, maximum idle time, maximum blocks read per session, maximum blocks read per call, and maximum amount of SGA.
Roles	Oracle	A role is a set or group of privileges that can be granted to users to another role. A privilege is a right to execute a particular type of SQL statement or to access another user's object. For example: the right to connect to a database, the right to create a table, the right to select rows from another user's table, the right to execute another user's stored procedure. System privileges are rights to enable the performance of a particular action, or to perform a particular action on a particular type of object. Roles are named groups of related privileges that you grant users or other roles. Roles are designed to ease the administration of end user system and object privileges. However, roles are not meant to be used for application developers, because the privileges to access objects within stored programmatic constructs needs to be granted directly.
Sequences	DB2, Oracle	A sequence generates unique numbers. Sequences are special database objects that provide numbers in sequence for input into a table. They are useful for providing generated primary key values and for the input of number type columns such as purchase order, employee number, sample number, and sales order number. Sequences are created by use of the CREATE SEQUENCE command.
Structured Types	DB2	Structured Types are useful for modeling objects that have a well-defined structure that consists of attributes. Attributes are properties that describe an instance of the type. A geometric shape, for example, might have as attributes its list of Cartesian coordinates. A person might have attributes of name, address, and so on. A department might have a name or some other attribute.
Synonyms	Oracle	A synonym is an alternate name for objects such as tables, views, sequences, stored procedures, and other database objects. A synonym is an alias for one of the following objects: table, object table, view, object view, sequence, stored procedure, stored function, package, materialized view, java class, user-defined object type or another synonym.
Tables	All	Tables are logical structures maintained by the database manager. Tables are composed of columns and rows. The rows are not necessarily ordered within a table. A base table is used to hold persistent user data. A result table is a set of rows that the database manager selects or generates from one or more base tables to satisfy a query. A summary table is a table defined by a query that is also used to determine the data in the table.
Tablespaces	DB2, Oracle	A tablespace is a storage structure containing tables, indexes, large objects, and long data. Tablespaces reside in database partition groups. They allow you to assign the location of database and table data directly onto containers. (A container can be a directory name, a device name, or a file name.) This can provide improved performance and more flexible configuration.
Triggers	All	A trigger defines a set of actions that are performed when a specified SQL operation (such as delete, insert, or update) occurs on a specified table. When the specified SQL operation occurs, the trigger is activated and starts the defined actions. Triggers can be used, along with referential constraints and check constraints, to enforce data integrity rules. Triggers can also be used to cause updates to other tables, automatically generate or transform values for inserted or updated rows, or invoke functions to perform tasks such as issuing alerts.

Undo Segments	Oracle	In an Oracle database, Undo tablespace data is an image or snapshot of the original contents of a row (or rows) in a table. The data is stored in Undo segments in the Undo table space. When a user begins to make a change to the data in a row in an Oracle table, the original data is first written to Undo segments in the Undo tablespace. The entire process (including the creation of the Undo data is recorded in Redo logs before the change is completed and written in the Database Buffer Cache, and then the data files via the database writer (DBW) process.)
Unique Keys	All	A unique key is a key that is constrained so that no two of its values are equal. The columns of a unique key cannot contain null values. The constraint is enforced by the database manager during the execution of any operation that changes data values, such as INSERT or UPDATE. The mechanism used to enforce the constraint is called a unique index. Thus, every unique key is a key of a unique index. Such an index is said to have the UNIQUE attribute. A primary key is a special case of a unique key. A table cannot have more than one primary key. A foreign key is a key that is specified in the definition of a referential constraint. A partitioning key is a key that is part of the definition of a table in a partitioned database. The partitioning key is used to determine the partition on which the row of data is stored. If a partitioning key is defined, unique keys and primary keys must include the same columns as the partitioning key, but can have additional columns. A table cannot have more than one partitioning key. Oracle: You cannot drop a unique key constraint that is part of a referential integrity constraint without also dropping the foreign key. To drop the referenced key and the foreign key together, check the Delete Cascade option for the foreign key. Clustered: A cluster composes of a group of tables that share the same data blocks, and are grouped together because they share common columns and are often used together. Filegroup: Lets you select the filegroup within the database where the constraint is stored. Fill Factor: Lets you specify a percentage of how large each constraint can become.
Views	All	A view provides an alternate way of looking at the data in one or more tables. A view is a named specification of a result table and can be thought of as having columns and rows just like a base table. For retrieval purposes, all views can be used just like base tables. You can use views to select certain elements of a table and can present an existing table in a customized table format without having to create a new table.