

Bonus Glossary

ADO (active data objects): A set of programmatic interfaces that Microsoft SQL Server talks to.

Alert: An event that responds to criteria that you specify.

Alias: An alternative name for a database object to make referring to said object easier.

Ambiguous join: A join between two tables that contain the same column name in each table but does not specify the table to which each column belongs.

Articles: Tables or stored procedures published in the course of replication or data warehousing.

Backup: The archiving of data in case of data loss.

Base Level Objects: Elemental SQL Server objects, such as databases and tables.

Batch: One or more processes designed to execute as a single unit.

Bind: To assign a rule or default to one or more columns in a table.

Bound: A rule or default assigned to a column in a table.

Clause: An identifiable part of a SQL statement that begins with a SQL-reserved keyword.

Client: The computer via which users connect to a server.

Clustered index: An index that physically sorts records in order.

Column: Also known as a field; stores the individual attributes of a row of data.

Complete backup: A backup of the entire set of data.

Complex select: A SELECT statement that consists of joining multiple tables.

Constraint: Any limitation placed on a table, such as an index.

Cursor: A virtual window that allows you to look at a specific number of rows in a result set.

DAO (data access objects): A set of programmatic objects that manipulate SQL Server objects.

Data warehouse: A method of storing data, normally in an external database, that allows for reporting without slowing down the main database.

Database: A container for objects that store data.

Database context: Term used to mean that a database must be selected before running a specific SQL Server wizard.

Datatype: Specification for the type of data a column is to represent.

DBA (database administrator): An individual who designs and administers all components of a database.

DB-Library: A DLL that knows how to communicate with a specific type of database, such as Microsoft SQL Server.

DBO: Database owner, or one who creates a database.

Deallocate: To free computer memory when releasing an object.

Default: The value for SQL Server 2000 to use for a column in a table when a SQL statement lacks a specified value.

Delete: SQL statement that specifies data to delete from one or more tables.

Delimiter: A character that serves to separate columns of data in a text file.

Denormalize: To construct data in tables so that data repeats; opposite of normalize.

Derived field: A field that does not actually exist in the database, but rather, consists of one or more fields that do exist in the database.

Destination: Indicates where the data is going.

Differential backup: A backup (archiving) of only data that has changed since the last backup.

Disabled: A field into which users cannot enter data and which appears grayed out on-screen.

Distributed architecture: Many computers networked together, with each computer serving a specific purpose. All the computers together serve the needs of the enterprise.

Distributor: A computer that acts as an intermediary between publishers and subscribers. Can be the same as the publishing computer.

DLL (dynamically linked library): DLLs contain many procedures that achieve some functionality. DLLs can be called by extended stored procedures.

DRI (declared referential integrity): A means of ensuring that a column of data in a table has a null or matching value in a corresponding table; normally enforced with foreign keys.

Drill-down: A method of expanding folders in a hierarchical tree to find specific folders or leaves.

DTS (Data Transformation Services): A way that Microsoft SQL Server allows you to import or export data.

Duty schedule: Indicates hours during which you can notify an operator about a job alert.

Dynamic: A process or data that can change while it is being evaluated; the opposite of static.

E-mail operator: An individual who may initiate an e-mail message when scheduling a job.

Enabled: A field on-screen that allows the user to type data into it.

Enterprise Manager: A program (which comes with Microsoft SQL Server) that enables you to administer the server, databases, and objects within the databases.

Event: A process that occurs when specific conditions exist.

Extended stored procedure: A stored procedure that runs a function contained within a DLL.

Field: See Column.

Filegroup: A group of files treated as a single unit.

Fixed field: Indicates that each value in a column is stored at the same size or length.

Folder: Item in a tree that indicates that there are more folders or leaves below it in the hierarchical format.

Foreign key: A column in a table that relates to a primary key in another table; used to enforce referential integrity (see also DRI).

Hierarchical: A way to display data in an outline format in such a way that items are stored in layers of categories.

HTML (Hypertext Markup Language): A text-based language that browsers use to display pages over the Internet.

Identity: A type of column in a table that increments automatically as new rows are added to the table.

Index: A database object that speeds up queries.

Inner join: Type of link between tables that does not guarantee every row of data that you expect will be returned.

INSERT: SQL statement that specifies data to insert into one or more tables.

Job: Task that you can perform automatically.

Join: The link between two tables in a relational database.

Key: Fields in a table that normally indicate the primary key.

Leaf: Item in a tree that indicates the lowest level available.

Legacy data: Data stored in another, possibly obsolete, database system.

Logical name: Name given to an object that you will use when you refer to that object after it is created. A logical name is different from a filename.

Login ID: Also known as a logon ID; indicates which specific user is logging in to the database.

Look-up table: A table that contains data that serves only to validate and provide a list of valid values for a column in another table.

Many-to-many: A type of relationship in which many values in one table can have many corresponding values in another table.

MMC (Microsoft Management Console): The familiar environment by which the SQL Server Enterprise Manager and other tools appears to the user.

Naming convention: A method for naming database objects, making it easy for others to recognize the purpose and scope of those objects.

Net-library: A DLL that knows how to communicate with a specific network protocol, such as Named Pipes or TCP/IP.

Net send operator: An individual who may initiate a message over the network when scheduling a job.

Nonclustered index: An index that you place on a table to speed up queries but that does not physically sort records in order.

Normalize: To construct data in tables so that data does not repeat. Typically requires more tables, resulting in slower queries but less disk space. The opposite of denormalize.

Not null: Indicates that a column cannot be left blank, or null.

Null: Indicates that a column can be left blank.

ODBC (Open Database Connectivity): A standard by which a program can communicate with Microsoft SQL Server and other RDBMSs.

One-to-many: Type of relationship in which a value in one table can have many corresponding values in another table.

One-to-one: Type of relationship in which a value in one table can have only one corresponding value in another table.

Operator: An individual who is notified of the status of a job. The operator receives a specific type of notification within the time period as defined by the job.

Outer join: Type of link between tables that guarantees the return of every row of data but that also can require more resources to execute.

Overhead: The amount of resources (in terms of memory or disk space) required for a procedure or operation.

Owner: One who creates the job or is responsible for maintaining the job.

Page operator: An individual who may initiate a page when scheduling a job.

Parameters: Arguments used in stored procedures to accept and return variables.

Password: Security measure that prompts users individually for a means to identify the user (usually for a word that the user types). Only the user knows it and it typically is encoded so that a user cannot see the password typed on the screen.

Performance hit: An adverse condition resulting in degraded performance in Microsoft SQL Server.

Physical name: A physical location assigned to an object.

Primary key: A column whose values uniquely identify a row in a table.

Publication: A series of one or more articles slated for publishing.

Publisher: A computer that publishes data to another computer.

Pull subscription: A way of transferring published data from the server to the client at the client's request.

Push subscription: A way of transferring published data from the server to the client at the server's request.

Query: A SQL statement that serves to retrieve data from one or more tables; typically includes the SELECT statement.

Query plan: The method that Microsoft SQL Server uses to access data in the fastest way possible.

RDBMS (Relational Database Management System): A database program that stores data in tables that relate to one another.

RDO (remote data objects): Similar to DAO but more efficient.

Read-only: A mode in which a row or column in a table cannot be updated, only read.

Read-write: A mode in which a row or column in a table can be updated or read.

Record: See Row.

Recursive: An operation that can actually call itself. If recursive operations are not handled carefully, your computer can lock up.

Referential integrity: See DRI.

Relational database: A database system that contains tables that relate to each other in some way defined by its creator.

Replication: A means of duplicating data from one or more tables or databases into another; typically used in a data warehouse.

Restore: Retrieving data from a backup.

Result set: The rows of data that a query returns.

Row: Also known as a record; represents the individual items of data that make up a table.

Rule: A constraint that indicates what values are valid for a column in a table.

Schema: The structure of a database.

Server: The computer that runs Microsoft SQL Server.

Server context: Term used to mean that a server must be selected before running a specific SQL Server wizard.

Service: A process or program that runs in the background.

Severity: A measure of how critical an error is.

Simple select: A SELECT statement that consists of only one table.

Source: Indicates where data is coming from.

SQL (Structured Query Language): The standard language used by all relational databases, including Microsoft SQL Server.

Static: A process (or data) that does not change as it is being evaluated; opposite of dynamic.

Stored procedure: A compiled set of SQL logic, stored within Microsoft SQL Server, that can be called from within any SQL statement.

String literal: Series of characters that must be used exactly as entered; cannot be substituted the way a variable can.

Subscriber: A computer that receives published data from another computer.

Subset: A part of a result set or other mass of data.

Tab: A visual folder at the top of the screen that categorizes data entry fields into logical groups.

Table: A storage device that is contained within a database to allow for a two-dimensional array of data to be stored permanently.

Table scan: A method by which Microsoft SQL Server does not use an index, but reads the entire table into memory.

Ticks: Single quotation marks.

Tier: Layer of functionality that divides, or encapsulates, business rules into logical groups.

Transaction: Process by which every SQL statement contained within the transaction must succeed. If even one of the SQL statements fail, the transaction is rolled back to the state the database was in before the transaction began.

Transaction log: A database file where temporary transaction data is stored so that the data can be rolled back if a transaction fails.

Transact-SQL: Version of SQL that Microsoft SQL Server uses. It contains some specific keywords and statements that only Microsoft SQL Server understands.

Tree: Visual hierarchical display of categorized items.

Trigger: Database object that executes (or fires) SQL logic that you define when data is inserted, updated, and/or deleted (depending on how you design the code).

Unicode: Specification for a character in Microsoft SQL Server so that it can be correctly translated in all languages. To store Unicode data, you must use a Unicode datatype.

UPDATE: SQL statement that specifies data to update in one or more tables.

URL (Uniform Resource Locator): The user-friendly address associated with a Web page.

USE: SQL keyword for switching to a specific database.

User: A person who accesses Microsoft SQL Server.

User-defined types: A way to define your own datatypes in Microsoft SQL Server. Also known as UDTs.

Variable: A placeholder in a SQL procedure that is substituted for an actual value when the procedure executes.

View: A database object that allows you to define the columns and rows that a specific user can see; also can serve as a tool for enforcing security within your database.

Visual Basic: A Microsoft product that enables you to write Windows applications that can access data in Microsoft SQL Server.