

ad hoc query A query that is executed infrequently. Ad hoc queries can sometimes be troublesome because untrained users can create queries that are inefficient, causing the server to suffer significant performance degradation.

aggregate A group such as grouped data. For example, when aggregating data, we are grouping data. A common aggregate function is Avg (average). It looks at a group of data (an aggregate) and provides an average.

alerts, Performance A Performance alert is an alert configured within the operating system Performance Logs and Alerts tool. Performance alerts can be configured to log an entry into the application event log, send a network message, start a performance log, or run a program. They are not directly connected to SQL Server Agent.

alerts, SQL Server Agent alerts A SQL Server Agent alert is an automated response to an event. SQL Server Agent monitors the application log and if an alert has been specified, the SQL Server Agent fires the alert. Alerts can notify operators or launch jobs.

alias Within SELECT statements there are two possible definitions of an alias. First, the column header on the output can be labeled with something different than the name of the column in the SELECT statement by using an alias. Second, tables can be identified with an alias. This is useful in a SELECT statement joining multiple tables.

ALTER A DDL statement used to modify objects within a database, or objects on the server.

baseline A baseline is a known starting point for something. In the context of the MCITP Database Developer certification, it's a known starting point for a server. For example, when creating a performance baseline, we would measure the four core resources of a system: CPU, memory, disk, and network. A performance baseline would take a snapshot of the resources (perhaps every 30 minutes) over a period of about a week. Six months later, another counter log could be created, and by comparing it to the baseline, an administrator can identify what has changed.

bcp The Bulk Copy Program (bcp) utility bulk copies data between SQL Server 2005 and a data file, using the format specified by the user (such as commas or other characters employed to separate the fields). It is intended to be an easy-to-use tool, but not when using the queryout option.

BI Business Intelligence (BI) is a group of different applications and technologies used to collect data, access data, and analyze data about the business. BI is used to help decision makers make better business decisions.

BIDS Business Intelligence Development Studio (BIDS) is installed with SQL Server 2005. It's used to create SQL Server Integration Services (SSIS) packages and SQL Server Reporting Services (SSRS) reports.

BLOB A binary large object. Large value data types [varchar(max), nvarchar(max), and varbinary(max)] are stored as BLOBs. Within SQL Server 2005, BLOBs can be as large as 2GB.

blocking lock A lock that blocks another process is called a blocking lock. Locks are placed on data to prevent it from being viewed or modified by one process as it is being viewed or modified by another process. Locks and blocking locks are not necessarily a problem. It's only when blocking locks are held for a long period of time and stopping a process from moving forward that they become problematic.

BOL Books Online (BOL) is a built-in help resource for SQL Server 2005. It can be read as a book using chapters from the contents tab, or searched for specific information. When connected to the Internet, BOL also returns information from MSDN Online, the Codezone Community.

bulk-logged recovery model In the bulk-logged recovery model, all transactions are logged into the transaction log EXCEPT bulk-logged operations. It is used when we want the capability to recover our data up to the point of failure, including all transactions in the transaction log, except for bulk-logged operations that either take up too many resources to log, or are easily reproduced.

CLR Common language runtime (CLR) is an execution environment for program code developed by Microsoft. Programs are written and developed in any language, and then compiled to a Microsoft Intermediate Language (MSIL). When run, they are submitted to the CLR to be fully compiled and run. CLR code can be integrated into SQL and it's highly recommended to do so for computation-intensive purposes.

clustered index A special type of index used to optimize queries. Tables can have only one clustered index. The data is ordered the same way the index is ordered. It's similar to a dictionary or a phone book. When you find the word or name, you're pointed right at the data.

collation Identifies how data is stored, sorted, compared, and presented. A collation is a set of rules and can be set at the database and column levels. It includes locale, sort order, and case sensitivity.

Common Table Expression (CTE) A temporary result set that is defined and used within the execution scope of another T-SQL statement. It is similar to a derived table in that it isn't persisted beyond the query. Unlike a derived table, a CTE can reference itself recursively within the same query.

composite index An index that is composed of more than one column. Both clustered and nonclustered indexes can be created as composite indexes.

computed column A column in a table that displays the result of an expression instead of stored data. For example, $\text{InventoryCost} = \text{QuantityOnHand} * \text{ProductCost}$. A calculated column could be calculated on-the-fly with the results not being stored, or the data can be persisted, where the computed data is held within the table.

concurrency Provides multiple users with the ability to access data simultaneously. Locks are used automatically by SQL Server and can be configured by setting transaction isolation levels.

constraint Constraints are special database objects used to enforce different types of integrity on data. While data types enforce basic data integrity (domain integrity), constraints allow much more sophisticated checking (entity integrity and referential integrity). Types of constraints include PRIMARY KEY, FOREIGN KEY, CHECKS, and DEFAULTS.

CREATE A DDL statement used to create new objects within a database, or objects on the server.

cross-tab report A cross-tab report is one where the results have been rotated so the rows are displayed as columns, or the columns are displayed as rows.

DAC A dedicated administrator connection (DAC) is used to connect to a SQL Server instance that is otherwise unresponsive. Needs SQL Server Browser service to run remotely.

DatabaseMail A service that can be used to send e-mail messages from within SQL Server 2005. It uses SMTP. Intended to replace SQL Mail, which has been deprecated.

database mirror A new feature of SQL Server 2005 where a database can be mirrored on another instance or another server. When configured with a witness server and the mirror session placed in synchronous mode, automatic failover can be achieved.

database securable Any object within a SQL Server database that can have access regulated. Securables include tables, VIEWS, stored procedures, functions, and even security principals.

database snapshot A new feature of SQL Server 2005 where a snapshot of a database can be created at any given time to preserve the state of the database. The snapshot can be queried if desired and/or the entire database can be restored from the snapshot.

data mining The process of extracting valid, authentic, and actionable information from large databases. The primary tool for data mining in SQL Server 2005 is SQL Server Analysis Services (SSAS), though SQL Server Integration Services (SSIS) can often be used also.

DBA A database administrator (DBA) is typically responsible for administering a database and the database server, as opposed to a database developer, who does much of the development work for a database. Administering responsibilities often include recoverability (backups and restores), maintaining integrity, managing security, ensuring availability, and maximizing performance.

DBCC Database Console Command (DBCC) was previously known as Database Consistency Checker. Four categories of DBCC commands include many methods that check on the various elements of a database.

DBMS A database management system (DBMS) is composed of software necessary to organize, store, and retrieve data in a database. Access, SQL Server, Oracle, and DB2 are all database management systems.

DDL Data Definition Language (DDL). This includes the statements of CREATE, ALTER, and DROP and applies to database objects (not data) such as tables, triggers, functions, stored procedures and more.

DDL trigger A trigger that fires in response to a DDL statement. DDL triggers can only be configured as AFTER triggers, not INSTEAD OF triggers. DDL triggers are configured at the database level or server level.

deadlock A locking condition where two processes both have locks on a resource and are waiting for the other process to release a resource before they can continue. For example, Sally could have a lock on the Sales table and be waiting to obtain a lock on the Product table. At the same time, Joe has a lock on the Product table and is waiting to obtain a lock on the Sales table. As a result, neither process can move forward. SQL Server automatically detects deadlocks and picks one of the processes as a deadlock victim and rolls it back.

DELETE A DML statement used to delete data within a table or VIEW. Note that data is deleted, but objects are dropped.

denormalized Process of optimizing a database by adding redundant data. A normalized database typically conforms to first normal form, second normal form, and third normal form, though more normal forms exist. By denormalizing a database, performance gains can sometimes be achieved.

deterministic Typically refers to functions such as deterministic functions. A deterministic function always returns the same result when called with the same input values. Indexed VIEWS require any functions used within the indexed VIEW to be deterministic. As a comparison, see nondeterministic.

differential backup A backup type that backs up all the changes since the last full backup. Since the differential backup only backs up the changes, it can be done much quicker than a full backup. A possible backup strategy might include performing a full backup once a week and doing differential backups daily.

Distributor Used in replication. In the Publisher metaphor, a Distributor is the process that transfers data from the Publisher to the Subscriber.

DMF Dynamic management function (DMF). Dynamic management VIEWS and functions are new to SQL Server 2005. They are used to return server state information to monitor the health of a server instance, diagnose problems, and tune performance. There are about 12 different categories of both DMVs and DMFs.

DML Data Manipulation Language (DML) includes the statements SELECT, INSERT, DELETE, and UPDATE and applies to data (not database objects) within tables and VIEWS.

DML trigger A trigger that fires in response to a DML statement on a table or view. DML triggers can be AFTER triggers on a table, or INSTEAD OF triggers on a table or a VIEW.

DMV Dynamic management VIEWS (DMVs) and functions are new to SQL Server 2005. They are used to return server state information to monitor the health of a server instance, diagnose problems, and tune performance. About 12 different categories of both DMVs and DMFs exist.

DROP A DDL statement used to delete objects within a database, or objects on the server. Note that data is deleted, but objects are dropped.

DTA The Database Engine Tuning Advisor (DTA) is used to analyze the performance effects of T-SQL statements and provide recommendations to add, remove, or modify indexes, indexed VIEWS, and partitioning (physical design structures).

DTS Data Transformation Services (DTS) was the ETL tool available in SQL Server 7.0 and SQL Server 2000. SQL Server Integration Services (SSIS) is the successor to DTS. Legacy DTS packages can be migrated into SSIS, or run using the Execute DTS 2000 Package task.

dynamic SQL A SQL statement built and executed at runtime by concatenating different parts of the statement based on provided variables. Generally, dynamic SQL is not recommended, especially in web sites, due to SQL injection attacks. The solution is to use parameterized stored procedures that build the SQL statement differently.

endpoints A new feature in SQL Server 2005 used to manage connections. An endpoint is a SQL Server object used by SQL Server to communicate over the network. In database mirroring, a SQL Server instance uses a special purpose endpoint to receive database mirroring connections with other server instances. The two primary endpoints are HTTP endpoints and database mirroring endpoints.

ETL This is a Business Intelligence (BI) term representing Extract, Transform, and Load. We extract data out of one database or table, transform it to conform to the standards in the destination, and load it into the target database or table. SSIS is the primary tool used for ETL in SQL Server 2005.

EXCEPT The EXCEPT operator can be used within a SELECT statement. It will return any distinct values from a query to the left of the EXCEPT operator that are not also returned from the query to right of the EXCEPT operator.

federated database A database that is spread across multiple servers, often in multiple geographical locations, is called a federated database. The servers that hold the different parts of a federated database are referred to as a federation, or federated database servers. A federation of database servers is used to spread the processing load across a group of servers. The data is horizontally partitioned allowing each of the servers to be independently managed, but distributed queries can be used to process requests on the entire database.

filegroups A method of optimizing the performance of a database by controlling the placement of database files and database objects. By default, all data and database objects are placed into a single file, in a single filegroup. Unlike files, which can be viewed on the disk, a filegroup is conceptual.

fillfactor An index option that identifies how full an index will be when it is created. For tables that have a lot of INSERTS, setting an indexes fill factor to something other than 0 (indicating 100 percent full) will prevent excessive page splits and the resulting fragmentation of indexes.

FOREIGN KEY (FK) An FK is used to create a relationship between two tables, and typically points to a PK (PRIMARY KEY) in another table. The relationship enforces integrity between the two tables, allowing only entries in the FK table that exist in the PK table, and preventing deletions from the PK table if a related entry exists in the FK table.

fragmentation In databases, indexes can be fragmented similar to how a hard drive can be fragmented. A fragmented index results in slower performance of the database. Fragmentation can be reduced by setting a fill factor on an index so it has empty space. Fragmented indexes can be defragmented by using REORGANIZE (keeps the index online) or by using REBUILD (which defaults to offline but can be run online).

full backup A full backup backs up the complete database. This includes all data, all objects, and all files. A full backup also backs up the transaction log, but does not truncate it. Both differential and transaction log backups need to have a full backup done first.

full recovery model In the full recovery model, all transactions are logged into the transaction log. It is used when we want to be able to recover our data up to the point of failure, including all transactions in the transaction log.

full-text catalog Used to hold full-text indexes. A full-text catalog holds zero, one, or more full-text indexes.

full-text index A separate file that stores information about significant words in a column. Noise words (such as *the*, *and*, *a*, and so on) are not included in the index. Full-text indexes are used for complete full-text searches.

full-text search Allows faster searches of text columns (char, varchar, and nvarchar) and columns that include formatted binary data such as Microsoft Word documents held

in a varbinary(max) column. A full-text search is only possible on a column that has a full-text index created on it.

function Functions are routines that can accept parameters, perform an action, and return the result of that action. SQL Server includes many built-in functions. User-defined functions can be created to meet specific needs.

index A database object used to provide faster access to data in a table. SQL Server 2005 has two types of indexes: clustered and nonclustered. Indexes are also used to enforce uniqueness.

INSERT A DML statement used to add new rows to a table.

INSTEAD OF trigger A trigger configured to fire instead of the action that caused it to fire. INSTEAD OF triggers are commonly used for updateable views and are only available with DML statements.

INTERSECT The INTERSECT operator can be used within a SELECT statement. It returns any distinct values from a query that are returned by the queries on both sides of the INTERSECT operator.

job SQL Server Agent workflows. A job can have one or more tasks within the workflow.

join Used to combine the contents of two or more tables. The most common join is an inner join. Other joins are left, right, full, and cross.

linked server A definition that specifies an external OLE DB database source, such as another SQL Server, or an Oracle server. Once defined, the linked server can be used for distributed queries using only the four-part name in the query.

lock A lock is an access restriction placed on part of a database to prevent other users or processes from viewing or modifying data as it is being viewed or modified by one process. Locks can be placed on rows, pages, extents, tables, or databases.

log shipping A high availability strategy where a copy of a database is created on another server or instance. The transaction log is periodically copied or shipped over to the standby server to keep it up-to-date.

lookup table In SQL Server databases, a lookup table is a table with relatively static data that can be used as a source for verifying data. For example, a table that includes the 50 states with their full spelling and two-letter abbreviations could be used as a lookup table.

maintenance plan A workflow of one or more maintenance tasks. Maintenance plans can be created manually or with the Maintenance Plan Wizard. Maintenance plans create jobs that are managed by the SQL Server Agent.

manifest file The installation file that is part of a deployment utility created by SSIS. Launching the manifest file will start the installation of the SSIS package and a server.

many-to-many table relationship A many-to-many relationship is between *three* tables where many rows in one table can link to many rows in a related table. A many-to-many relationship can't be created directly, but instead requires a third table (often referred to as a junction table) that connects the two tables. When looking at table diagrams, the one is typically identified with a key icon, and the many is typically identified with an infinity icon.

master database The primary system database that controls all the system-level operations of SQL Server. It records instance-wide metadata, such as logon accounts, endpoints, linked servers, and system configuration settings.

merge replication A replication strategy used when multiple subscribers are also acting as publishers. In other words, the data is updated from multiple sources.

Model database The template used when creating new databases. Any new database is created from a copy of the Model database and then modified from there.

msdb database Used by SQL Server Agent to store information on jobs, alerts, and operators. It also holds all historical data.

nonclustered index A nonclustered index is added to optimize queries. Tables can have multiple nonclustered indexes. A nonclustered index is similar to the index in the back of a book. By finding it in the index, you know specifically where to look in the book for the information.

nondeterministic Typically refers to functions such as nondeterministic functions. A nondeterministic function returns different results when called with the same input values. As an example, GETDATE() would return different results at different times. Indexed views can not include nondeterministic functions.

normalized A normalized database is where the tables in the database are reduced to their simplest terms. The concept came from a paper written by Dr. E. F. Codd in 1970 where the first three normal forms were defined.

ODBC Open Database Connectivity (ODBC) is an application programming interface (API) that supports access to any data source as long as an ODBC driver is available.

OLAP Online analytical processing (OLAP) refers to the process of creating multiple dimensions (or cubes) of a database and then performing the multidimensional analysis of business data. OLAP is part of the broader category of business intelligence and is supported by SQL Server Analysis Services (SSAS) in SQL Server 2005. This book's focus has been primarily on OLTP, not OLAP.

OLE DB Object Linking and Embedding Database (OLE DB) is an application programming interface (API) used for accessing data. OLE DB supports accessing data in many formats besides database (such as spreadsheets and text files).

OLTP Online transaction processing (OLTP). An OLTP database is optimized for changing data, as opposed to an OLAP database, which is optimized for data that is relatively static. Data is changed via DML statements (INSERT, UPDATE, DELETE).

one-to-many table relationship A one-to-many relationship is between two tables where a single row in one table can link to many rows in the related table. When looking at table diagrams, the one is typically identified with a key icon, while the many is typically identified with an infinity icon.

one-to-one table relationship A one-to-one table relationship is between two tables where a single row in one table is linked to a single row in another table. When looking at table diagrams, the one is typically identified with a key icon, and the many is typically identified with an infinity icon.

OPENDATASOURCE A command that can be used to query external servers. OPENROWSET provides similar functionality. If the queries are to be repeated, it is recommended to create linked servers to make the queries less complex.

operator In SQL Server Agent, an operator identifies who is notified. It can be configured to send e-mails, pages, or Net Send messages.

parameterized report A SQL Server Reporting Services (SSRS) report that accepts input values (parameters) is known as a parameterized report. Parameters are used to complete a query used for the report so the report can be selective based on user input.

password policy New to SQL Server 2005, a password policy applies password security policies to SQL logins. This feature is only fully supported on Windows Server 2003. The CHECK_POLICY option should be set to ON to enable this feature.

Performance Also called Performance Monitor and System Monitor, though technically it is called simply Performance. Performance includes System Monitor, which can be used to measure system objects and counters in real time. Common objects measured on any server include CPU, memory, disk, and NIC. Performance also includes Performance Logs and Alerts which can be used to create counter logs or traces and performance alerts.

persisted computed column A computed column of data that is physically stored in the table. For example, a Products table could have an On_Hand_Cost calculated from the On_Hand_Inventory and Cost columns. A computed column is calculated every time the data is queried, but is not stored. This has a performance cost if the column is heavily queried. A persisted computed column can store the new value each time the On_Hand_Cost or On_Hand_Inventory value is changed using a DML trigger at a lower performance cost.

physical design structure Physical design structures (PDSs) include items such as indexes, indexed VIEWS, and partitioning. These are referenced in the Database Engine Tuning Advisor, which is used to evaluate a database and can recommend the implementation of different types of physical design structures for better performance.

PIVOT The PIVOT operator is used within a SELECT statement. It is used to create cross-tab reports (similar to a spreadsheet) from normalized data.

predicate An expression that can be used in search conditions of WHERE clauses, HAVING clauses, and join conditions of FROM clauses. Predicates evaluate to true, false, or unknown.

PRIMARY KEY (PK) This is a special column, or columns, within a table that's used to enforce uniqueness. Typically, the first column is created as a PRIMARY KEY and often it is the name of the table appended with ID. For example, the Employees table would have a PK of EmployeeID. A table can have only one PK. A unique clustered index is created by default when the PK is created.

Publication In the Replication Publisher metaphor, the Publication identifies that the data is replicated. Publications can include multiple articles.

Publisher In the Replication Publisher metaphor, the Publisher is the database that is the source of the data.

query optimizer An optimization process running within SQL Server. Any queries submitted to SQL Server are first processed by the query optimizer. It determines the best way to run the query, including what indexes to use and what types of joins to use. The output is a query execution plan, sometimes called a query plan or just a plan.

query plan Once the query optimizer determines the best way to execute a query, it creates a query plan. This identifies all the elements of the query, including what indexes are used, what types of joins are employed, and more. The query execution plan can be observed in SSMS by pressing `CTRL+L` or by selecting Query | Display Estimated Execution Plan.

recovery models Identifies how the transaction log is used and what can be recovered in the case of database failure. The three recovery models in SQL Server 2005 are simple, bulk-logged, and full.

recursion Occurs when one process calls itself to run again. With triggers, it's the process of a trigger firing itself. Indirect recursion is where an update to Table1 fires a trigger that affects Table2 that fires a trigger that updates Table1 again. Direct recursion is where an update to Table1 fires a trigger that affects Table1 again that fires the trigger again.

replication A group of technologies within SQL Server 2005 that are used to copy and distribute data and database objects from one database to another. Data is then regularly synchronized to maintain consistency. Replication uses a publishing metaphor with Publishers (data source), Distributors (process responsible for replicating the data and/or objects), and Subscribers (data target).

Replication Monitor Tool used to monitor replication. Can observe real-time activity, troubleshoot problems, and analyze past replication activity.

Report Builder Report Builder is a tool designed to let end users create reports from a report model. A great strength with Report Builder is that users can modify their reports, effectively creating ad hoc reports without the need to ask developers to modify the report.

Report Designer Report Designer is the tool used within Business Intelligence Development Studio (BIDS) to create reports. This is not available to most end users. Wizards are available to make the process of creating reports easier.

Report Manager Report Manager is the primary interface users can use to view, search, and subscribe to reports. For example, report models can be deployed, and users can then use the Report Builder in Report Manager to create their own reports. It can also be used to administer Report Server remotely.

report models Report models are templates used to create reports with Report Builder. They include the data source definitions (such as which server and which database to connect to for the model) and data source VIEW definitions (such as which tables or VIEWS to include in the model). Reports can't be viewed from a report model. Instead, the report model must be used to create a report using Report Builder.

report snapshot A report that contains data captured at a specific point in time. Since report snapshots hold datasets instead of queries, report snapshots can be used to limit processing costs by running the snapshot during off-peak times.

resource database A read-only database that contains system objects included with SQL Server 2005.

role Roles exist at both the server level and the database level. By adding logins to roles, they are automatically granted the permissions of that role. For example, by adding a user to the sysadmin role, the user can do anything on the instance of the server. By adding a user to the db_owner role, the user can do anything in the database.

roll back The process of undoing uncommitted transactions. As a part of the recovery process, uncommitted transactions are rolled back to ensure the database is recovered in a consistent state.

roll forward The process of applying committed transactions. As a part of the recovery process, committed transactions are rolled forward to ensure the database is recovered in a consistent state with the changed data in the database.

SAN Storage area network (SAN) architecture is used to attach remote computer storage hardware in such a way that servers recognize it as local storage. SANs are typically very expensive, so they are only used in large applications.

scalar A scalar result is a single result. Some functions are referred to as scalar because they return a single answer. For example, MAX and MIN are considered scalar functions. While they look at a group of data, they return only a single answer.

schema In SQL Server 2000, the schema identified the model of the database. For example, the tables, VIEWS, data types, and so on would be identified as the schema. In SQL Server 2005, the term schema is used to identify a collection of database entities within a single namespace. Schemas are the owners of the objects, and one or more users can be the owners of schemas.

schemabinding An option that allows an object (such as a VIEW) to prevent referenced tables or VIEWS from being modified in such a way that would cause the original object to no longer work.

SD³+C Microsoft's security mantra. It started as SD³ (Secure by Design, Secure by Default, and Secure in Deployment) and then evolved to SD³ + C (Secure by Design, Secure by Default, and Secure in Deployment and Communications)

security principal SQL logins or Windows logins that are granted access to the server, and users or roles granted access to a database.

SELECT A DML statement used to retrieve data from tables or VIEWS.

Service Broker A new service in SQL Server 2005 that allows developers to build asynchronous applications by exchanging messages. Contracts identify messages used by a Service Broker service and create an agreement between two services. Messages are held in queues if they can't be sent right away. Conversations between two applications are referred to as dialogs.

Service Broker contract A contract is an agreement between two services. It defines the message types an application uses.

Service Broker conversation An exchange of messages between two services. Also referred to as a dialog.

Service Broker queues Where messages are stored until delivered or retrieved.

simple recovery model In the simple model, transactions logged into the transaction log are not available for restore purposes. Simple recovery is used for internal recovery purposes and to maintain database consistency. It is automatically truncated and doesn't need to be backed up or otherwise maintained.

snapshot replication Replication of data taken at a moment of time. With snapshot replication, the entire data set is replicated at the same time.

SOAP Simple Object Access Protocol (SOAP) is a protocol used to exchange XML messages, typically using HTTP.

SQL Pronounced as *es-que-el* or *sequel*, SQL is commonly understood to be an abbreviation of Structured Query Language—a language used to retrieve, update, insert, and delete data in tables and VIEWS. SQL Server 2005 uses a version of SQL referred to as Transact-SQL.

SQL injection attack An Internet attack against a database accessible via a web page. Automated programs are available to launch attacks, and successful SQL injection attacks can obtain the entire layout of a database and all the data.

SQL login A login created to provide access to SQL Server. Unlike Windows logins which are associated with local or domain users or groups, a SQL Server login is not associated with any outside user or group. SQL logins are needed for users accessing SQL Server from non-Microsoft systems.

SQL Profiler Used as the GUI equivalent of the command line SQL Trace. Monitors SQL Server by capturing (or tracing) activity and saving it into a file or table. SQL Profiler can save traces used by DTA, and saved traces can also be imported into System Monitor with corresponding counter logs. It is effective in identifying slow-running queries and deadlocks.

SQL Server Agent A management tool used to create jobs, fire alerts, and notify operators. All SQL Server Agent data is stored in the msdb system database.

SQL Server Configuration Manager The tool for providing basic configuration management for SQL Server services and protocols. Services can be configured, enabled, and disabled in this tool. Protocols such as Shared Memory, TCP/IP, and Named Pipes can be enabled or disabled.

SSAS SQL Server Analysis Services (SSAS) includes tools and features for OLAP that are used to design, deploy, and maintain cubes.

SSIS SQL Server Integration Services (SSIS) is used to build data integration packages, including ELT packages. It is the replacement for DTS in SQL Server 2000.

SSMS SQL Server Management Studio (SSMS) is the primary GUI to access, configure, manage, administer, and maintain SQL Server 2005 components. It combines the features of Enterprise Manager and Query Analyzer from SQL Server 2000.

SSRS SQL Server Reporting Services (SSRS) is a server-based reporting platform used to create and manage a wide variety of reports, including parameterized reports, report snapshots, and more. Reports can be deployed to an IIS server so they are easily accessible to users.

statistics Statistics are a small sample of a whole used to represent the whole. In SQL Server, statistics are maintained on indexes and used by the query optimizer to determine which indexes to use for a given query. Instead of scanning an entire index (the entire population), statistical data (a sampling) is maintained on the index.

Subscriber In the Replication Publisher metaphor, the Subscriber is the database that receives the data.

synonyms Synonyms are used as an alternative name for a schema-scoped object. Objects can be referenced using the single-part synonym name, or the two-part, three-part, or four-part name.

System Monitor An operating system tool that allows the measurement of counters and objects. Key operating system objects are the CPU, memory, disk drive, and NIC. When SQL Server is installed, additional objects are added that can be measured, such as SQL Agent:Statistics, SQL Server:Buffer Manager, SQL Server:Locks, SQL Server:SSIS Pipeline, and many more.

table A two-dimensional storage location of data. Databases within SQL Server 2005 are composed of multiple related tables.

tail-log The transaction log that is backed up from a possibly damaged database is called the tail-log. It holds all of the transactions since the last backup.

tempdb One of the system databases. The tempdb database is used for temporary storage of data, including temporary objects created in normal operation, and results from temporary sorts when creating or rebuilding indexes using the SORT_IN_TEMPDB option. Since it is recreated each time SQL Server is restarted, it should be set to the size needed instead of allowing frequent auto growth operations.

trace A trace is a collection of events and data. SQL Profiler is used to collect and monitor events. Creating a trace is sometimes referred to as capturing events.

transaction One or more T-SQL statements grouped together. In a transaction, either all operations must succeed or they all fail. Transactions are identified with the BEGIN TRAN statement, and ended with either a COMMIT TRAN or ROLLBACK TRAN statement.

transactional replication Replication that starts with a snapshot and then keeps the Subscribers up-to-date by using the transaction log. Transactions are recorded on the Publisher, distributed to the Subscribers, and then applied to keep the Subscribers up-to-date.

transaction log The file that holds a record of all transactions (INSERTS, UPDATES, DELETES) in a database. Data modifications are first recorded in the transaction log, and then periodically (at checkpoints) the data is written to the database. The transaction log functions differently depending on which recovery model is being used: simple, full, or bulk-logged.

TRIGGER A type of stored procedure that fires in response to action on a table. DML triggers are associated with INSERT, UPDATE, and DELETE statements. DDL triggers are associated with CREATE, ALTER, and DROP statements.

T-SQL Transact-SQL (T-SQL) is Microsoft's version of SQL. It is an extension of the SQL language defined by the International Standards Organization (ISO) and the American National Standards Institute (ANSI).

UNPIVOT The UNPIVOT operator is used within a SELECT statement to create a normalized data report from data that is stored as a spreadsheet.

UPDATE A DML statement used to modify data within a table or VIEW.

user-defined function Functions are routines that can accept parameters, perform an action, and return the result of that action. SQL Server includes many built-in functions, but user-defined functions can be created by any user with permission to do so. User-defined functions can be T-SQL functions or CLR-integrated functions.

VIEW A virtual table that represents data in one or more tables in an alternate way. Almost any SELECT statement can be converted to a VIEW.

witness server When using database mirroring, automatic failover can be configured by adding a witness server that monitors the status of the principal and mirror servers.

XML Extensible Markup Language (XML) is an offshoot of the Standard Generalized Markup Language (SGML), just as HTML is, and is designed to be relatively human-legible. Within SQL Server, XML is a data type that allows full XML documents to be stored unmodified within a SQL Server database.

XML methods SQL Server 2005 provides five different XML methods that can be used to query and manipulate XML data stored in an XML data type column. The five methods are Value, Nodes, Query, Modify, and Exist.