

# Glossary

Applies To: SQL Server 2016

## Terms

Term	Definition
action	An end-user-initiated operation on a selected cube or portion of a cube.
ActiveX Data Objects	A data access interface that communicates with OLE DB-compliant data sources to connect to, retrieve, manipulate, and update data.
ActiveX Data Objects (Multidimensional)	A high-level, language-independent set of object-based data access interfaces optimized for multidimensional data applications.
adapter host	The root abstract class Adapter, which defines the handshake between the adapter and the StreamInsight server in the ENQUEUE interaction point. It provides all the required adapter services such as memory management, and exception handling.
ADO	A data access interface that communicates with OLE DB-compliant data sources to connect to, retrieve, manipulate, and update data.
ADO MD	A high-level, language-independent set of object-based data access interfaces optimized for multidimensional data applications.
ADOMD.NET	A .NET managed data provider that provides access to multidimensional data sources, such as Microsoft SQL Server Analysis Services.
aggregate function	A function that performs a calculation on multiple values and returns a single value.
aggregate query	A query (SQL statement) that summarizes information from multiple rows by including an aggregate function such as Sum or Avg.
aggregation	A table or structure containing pre-calculated data for an online analytical processing (OLAP) cube. Aggregations support the rapid and efficient querying of a multidimensional database.
aggregation prefix	A string that is combined with a system-defined ID to create a unique name for a partition's aggregation table.

Term	Definition
aggregation wrapper	A wrapper that encapsulates a COM object within another COM object.
alias	An alternative label for some object, such as a file or data collection.
alias type	A user-defined data type based on one of the SQL Server system data types that can specify a certain data type, length, and nullability.
alignment	A condition whereby an index is built on the same partition scheme as that of its corresponding table.
allocation unit	A set of pages that can be operated on as a whole. Pages belonging to an allocation unit are tracked by Index Allocation Map (IAM) pages. An allocation unit consists of the IAM page chain and all pages marked allocated in that IAM page chain. An allocation unit can contain at most a single IAM chain, and an IAM chain must belong to one, and only one, allocation unit.
AMO	A collection of .NET namespaces included with Analysis Services, used to provide administrative functionality for client applications.
Analysis Management Objects	A collection of .NET namespaces included with Analysis Services, used to provide administrative functionality for client applications.
analytical data	Data that provides the values that are associated with spatial data. For example, spatial data defines the locations of cities in an area whereas analytical data provides the population for each city.
ancestor	In a tree structure, the element of which a given element is a child. Equivalent to a parent element.
ancestor element	In a tree structure, the element of which a given element is a child. Equivalent to a parent element.
anchor cap	A line cap where the width of the cap is bigger than the width of the line.
anchor member	The first invocation of a recursive CTE consists of one or more CTE_query_definition joined by UNION ALL, UNION, EXCEPT or INTERSECT operators. Because these query definitions form the base result set of the CTE structure, they are referred to as anchor members.
animation manager	A core component of an animation application and the central programmatic interface for managing (creating, scheduling, and controlling) animations.
anonymous subscription	A type of pull subscription for which detailed information about the subscription and the Subscriber is not stored.
ANSI to OEM conversion	The conversion of characters that must occur when data is transferred from a database that stores character data using a specific code page to a client application on a computer that uses a different code page. Typically, Windows-based client computers use ANSI/ISO code pages, and some databases (for compatibility reasons) might use OEM code pages.
API server cursor	A server cursor that is built to support the cursor functions of an API, such as ODBC, OLE DB, ADO, and DB-Library.

<b>Term</b>	<b>Definition</b>
API support	A set of routines that an application uses to request and carry out lower-level services performed by a computer's operating system. These routines usually carry out maintenance tasks such as managing files and displaying information.
application role	A SQL Server role created to support the security needs of an application.
application time	The clock time supplied by applications which must communicate their application time to the StreamInsight server so that all temporal operators refer to the timestamp of the events and never to the system clock of the host machine.
apply branch	The set of operations applied to an event group.
arbitration port	A TCP/IP port used by the cache hosts to determine whether a cache host in the cluster has become unavailable. The port number that is used for arbitration can be different on each cache host.
ARIMA	A method for determining dependencies in observations taken sequentially in time, that also supports multiplicative seasonality.
article	A component in a publication. For example, a table, a column, or a row.
assembly	A managed application module containing class metadata and managed code as an object in SQL Server, against which CLR functions, stored procedures, triggers, user-defined aggregates, and user-defined types can be created in SQL Server.
associative array	An array composed of a collection of keys and a collection of values, where each key is associated with one value. The keys and values can be of any type.
atom feed	An XML structure that contains metadata about content, such as the language version and the date when the content was last modified, and is sent to subscribers by using the Atom Publishing Protocol (AtomPub).
atomic	Pertaining to an operation where all the transaction data modifications must be performed; either all of the transaction data modifications are performed or none are performed.
attribute	A single characteristic or additional piece of information (financial or non-financial) that exists in a database.
attribute hierarchy	A flat hierarchy (typically having an All level and a member level) containing a single attribute. It is created from one column in a dimension table, if supported by the cube.
attribute relationship	The hierarchy associated with an attribute containing a single level based on the corresponding column in a dimension table.
attribute type	The type of information contained by an attribute, such as quarters or months in a time dimension, which may enable specific treatment by the server and client applications.
auditing	The process an operating system uses to detect and record security-related events, such as an attempt to create, to access, or to delete objects such as files and directories. The records of such events are stored in a file known as a security log, whose contents are available only to those with the proper clearance.

Term	Definition
authentication	The process of verifying the identity of a user, computer, process, or other entity by validating the credentials provided by the entity. Common forms of credentials are digital signatures, smart cards, biometric data, and a combination of user names and passwords.
authenticator	A data structure used by one party to prove that another party knows a secret key. In the Kerberos authentication protocol, authenticators include timestamps, to prevent replay attacks, and are encrypted with the session key issued by the Key Distribution Center (KDC).
authorization	The process of granting a person, computer process, or device access to certain information, services or functionality. Authorization is derived from the identity of the person, computer process, or device requesting access, which is verified through authentication.
autocommit mode	The default transaction management mode for the Database Engine. The Database Engine automatically starts a transaction for each individual Transact-SQL statement. When the statement completes, the transaction is committed or rolled back based on the success or failure of the statement.
auto-consistency check	A feature that automatically runs a consistency check on protected data sources when it detects an inconsistent replica.
automatic failover	A failover that occurs automatically on the loss of the primary replica.
automatic recovery	Recovery that occurs every time SQL Server is restarted.
auto-protection	In DPM, a feature that automatically identifies and adds new data sources for protection.
autoregressive integrated moving average	A method for determining dependencies in observations taken sequentially in time, that also supports multiplicative seasonality.
axis	A set of tuples. Each tuple is a vector of members. A set of axes defines the coordinates of a multidimensional data set.
backup set	A collection of files, folders, and other data that have been backed up and stored in a file or on one or more tapes.
balanced hierarchy	A dimension hierarchy in which all leaf nodes are the same distance from the root node.
base backup	A data backup of a database or files upon which a differential backup is fully or partially based. The base backup is the most recent full or file backup of the database or files.
base data type	Any system-supplied data type, for example, char , varchar , binary , and varbinary . User-defined data types are derived from base data types.
base object	The object that a synonym references.

Term	Definition
base table	A table stored permanently in a database. Base tables are referenced by views, cursors, SQL statements, and stored procedures.
basic marker map	A map that displays a marker at each location (for example, cities) and varies marker color, size, and type.
batch	A set of requests or transactions that have been grouped together.
batch job	A set of computer processes that can be run without user interaction.
batch processing	The execution of a batch file.
blittable type	A data type that has a unique characteristic and an identical presentation in memory for both managed and unmanaged environments. It can be directly shared.
BLOB	A discrete packet of binary data that has an exceptionally large size, such as pictures or audio tracks stored as digital data, or any variable or table column large enough to hold such values. The designation "binary large object" typically refers to a packet of data that is stored in a database and is treated as a sequence of uninterpreted bytes.
block	A Transact-SQL statement enclosed by BEGIN and END.
block cursor	A cursor with a rowset size greater than 1.
blocking transaction	A transaction that causes another transaction to fail.
Boolean expression	An expression that yields a Boolean value (true or false). Such expressions can involve comparisons (testing values for equality or, for non-Boolean values, the < [less than] or > [greater than] relation) and logical combination (using Boolean operators such as AND, OR, and XOR) of Boolean expressions.
Boolean operator	An operator designed to work with Boolean values. The four most common Boolean operators in programming use are AND (logical conjunction), OR (logical inclusion), XOR (exclusive OR), and NOT (logical negation).
bound stream	An event stream that contains all the information needed to produce events. Either the information is an already instantiated data source, or the information is sufficient for the StreamInsight server to start the data source.
bounding box	The smallest rectangular area that will surround a path, shape, or group of objects.
browse mode	A function that lets you scan database rows and update their values one row at a time.
B-tree	A tree structure for storing database indexes.
bubble map	A geographical map that displays a circle over specific locations, where the radius of the circle is proportional to a numeric value.

<b>Term</b>	<b>Definition</b>
buffer pool	A block of memory reserved for index and table data pages.
buffer size	The size of the area of memory reserved for temporary storage of data.
built-in functions	A group of predefined functions provided as part of the Transact-SQL and Multidimensional Expressions languages.
BUILTIN\Administrators	User account (local administrators)
bulk copy	An action of copying a large set of data.
bulk export	To copy a large set of data rows out of a SQL Server table into a data file.
bulk import	To load a large amount of data, usually in batches, from a data file or repository to another data repository.
bulk load	An action of inserting a large set of rows into a table.
bulk log backup	A backup that includes log and data pages changed by bulk operations. Point-in-time recovery is not allowed.
bulk rowset provider	A provider used for the OPENROWSET instruction to read data from a file. In SQL Server 2005, OPENROWSET can read from a data file without loading the data into a target table. This enables you to use OPENROWSET with a simple SELECT statement.
bulk-logged recovery model	A database recovery mode that minimally logs bulk operations, such as index creation and bulk imports, while fully logging other transactions. Bulk-logged recovery increases performance for bulk operations, and is intended to be used as an adjunct to the full recovery model.
Business Intelligence Development Studio	A project development and management tool for business intelligence solution developers. It can be used to design end-to-end business intelligence solutions that integrate projects from Microsoft SQL Server Analysis Services (SSAS), Microsoft SQL Server Integration Services (SSIS), and Microsoft SQL Server Reporting Services (SSRS).
business logic	The part of an application program that performs the required data processing of the business. It refers to the routines that perform the data entry, update, query and report processing, and more specifically to the processing that takes place behind the scenes rather than the presentation logic required to display the data on the screen.
business logic handler	A merge replication feature that allows you to run custom code during the synchronization process.
business logic handler framework	The business logic handler framework allows you to write a managed code assembly that is called during the merge synchronization process.
business rules	The logical rules that are used to run a business
cache aging	The mechanism of caching that determines when a cache row is outdated and must be refreshed.

Term	Definition
cache client	A .NET application that uses the Windows Server AppFabric client APIs to communicate with and store data to a Windows Server AppFabric distributed cache system.
cache cluster	The instantiation of the distributed cache service, made up of one or more instances of the cache host service working together to store and distribute data. Data is stored in memory to minimize response times for data requests. This clustering technology differs from Windows Clustering.
cache invalidation	The process of flagging an object in the cache so that it will no longer be used by any cache clients. This occurs when an object remains in cache longer than the cache time-out value (when it expires).
cache item	An object that is stored in the cache and additional information associated with that object, such as tags and version. It can be extracted from the cache cluster using the GetCacheItem client API.
cache notification	An asynchronous notification that can be triggered by a variety of cache operations on the cache cluster. Cache notifications can be used to invoke application methods or automatically invalidate locally cached objects.
cache operation	An event that occurs on regions or cached items that can trigger a cache notification.
cache port	A TCP/IP port used by cache hosts to transmit data to and from the cache clients. The port number used for the cache port can be different on each cache host. These settings are maintained in the cluster configuration settings.
cache region	A container of data, within a cache, that co-locates all cached objects on a single cache host. Cache Regions enable the ability to search all cached objects in the region by using descriptive strings, called tags.
cache service	The distributed, in-memory caching solution that enables users to build highly scalable and responsive applications by bringing data closer to end users.
cache tag	One or more optional string-based identifiers that can be associated with each cached object stored in a region. Regions allow you to retrieve cached objects based on one or more tags.
cache-aside programming pattern	A programming pattern in which if the data is not present in the cache, the application, not the distributed cache system, must reload data into the cache from the original data source.
cache-enabled application	An application that uses the Windows Server AppFabric cache client to store data in cache on the cache cluster.
calculated column	A type of column that displays the results of mathematical or logical operations or expressions instead of stored data.
calculated field	A field defined in a query that displays the result of an expression rather than displaying stored data. The value is recalculated each time a value in the expression changes.
calculated member	A member of a dimension whose value is calculated at run time by using an expression. Calculated member values can be derived from the values of other members.

Term	Definition
calculation condition	A Multidimensional Expressions (MDX) logical expression that is used to determine whether a calculation formula will be applied against a cell in a calculation subcube.
calculation formula	A Multidimensional Expressions (MDX) expression used to supply a value for cells in a calculation subcube, subject to the application of a calculation condition.
calculation pass	A stage of calculation in a multidimensional cube in which applicable calculations are evaluated.
calculation pass number	An ordinal position used to refer to a calculation pass.
calculation subcube	The set of multidimensional cube cells that is used to create a calculated cells definition. The set of cells is defined by a combination of MDX set expressions.
callback	The process used to authenticate users calling in to a network. During callback, the network validates the caller's username and password, hangs up, and then returns the call, usually to a preauthorized number. This process prevents unauthorized access to an account even if an individual's logon ID and password have been stolen.
call-level interface	The interface supported by ODBC for use by an application.
candidate key	A column or set of columns that have a unique value for each row in a table.
cap	For paths that contain unconnected ends, such as lines, the end of a stroke. You can change the way the stroke looks at each end by applying one of four end cap styles: flat cap, round cap, square cap, and triangle cap.
cardinality	The number of entities that can exist on each side of a relationship.
carousel view	In Power Pivot Gallery, a specialized view where the preview area is centered and the thumbnails that immediately precede and follow the current thumbnail are adjacent to the preview area.
cascading delete	For relationships that enforce referential integrity between tables, the deletion of all related records in the related table or tables when a record in the primary table is deleted.
cascading update	For relationships that enforce referential integrity between tables, the updating of all related records in the related table or tables when a record in the primary table is changed.
case	An abstract view of data characterized by attributes and relations to other cases.
case key	The element of a case by which the case is referenced within a case set.
catalog views	Built-in views that form the system catalog for SQL Server.
catastrophic error	An error that causes the system or a program to fail abruptly with no hope of recovery. An example of a fatal error is an uncaught exception that cannot be handled.
CD sleeve	A case for holding CDs.



Term	Definition
CD-ROM	A form of storage characterized by high capacity (roughly 650 MB) and the use of laser optics instead of magnetic means for reading data.
cell	In a cube, the set of properties, including a value, specified by the intersection when one member is selected from each dimension.
cellset	In ADO MD, an object that contains a collection of cells selected from cubes or other cellsets by a multidimensional query.
centralized registration model	A registration model that removes all certificate subscriber participation from the management policy. For the workflow, a user designated as the originator will initiate the request and an enrollment agent will execute the request.
CEP	The continuous and incremental processing of event streams from multiple sources based on declarative query and pattern specifications with near-zero latency.
CEP engine	The core engine and adapter framework components of Microsoft StreamInsight. The StreamInsight server can be used to process and analyze the event streams associated with a complex event processing application.
CERN	A physics research center located in Geneva, Switzerland, where the original development of the World Wide Web took place under the leadership of Tim Berners-Lee in 1989 as a method to facilitate communication among members of the scientific community.
certificate	A digital document that is commonly used for authentication and to help secure information on a network. A certificate binds a public key to an entity that holds the corresponding private key. Certificates are digitally signed by the certification authority that issues them, and they can be issued for a user, a computer, or a service.
certificate enrollment	The process of requesting, receiving, and installing a certificate.
certificate issuer	The certification authority which issued the certificate to the subject.
Certificate Lifecycle Manager Client	A suite of Certificate Lifecycle Manager (CLM) client tools that assist end users with managing their smart cards. The tools include the Smart Card Self Service Control, the Smart Card Personalization Control, and the Certificate Profile Update Control. See Smart Card Self-Service Control, Smart Card Personalization Control, Certificate Profile Update Control.
certificate manager	A Certificate Lifecycle Manager (CLM) user that has the appropriate CLM permissions to either administer other CLM users or to administer the CLM application itself.
certificate manager Web portal	A Web application running on the Certificate Lifecycle Manager (CLM) server. This portal allows certificate administrators to administer other users' certificates and smart cards. The certificate subscriber and certificate manager Web portals are both accessed through the same universal resource locator (URL); however, the content displayed is based on a user's roles and permissions.

Term	Definition
Certificate Profile Update Control	An ActiveX control that automates the update of Certificate Lifecycle Manager (CLM) profiles on client computers.
certificate revocation	The process of revoking a digital certificate.
certificate subscriber	A user that needs certificates with or without smart cards. Certificate subscribers can access a small number of functions that can only be performed for the user's own certificates.
certificate subscriber Web portal	A Web application running on the Certificate Lifecycle Manager (CLM) server. This component of the CLM server interacts directly with users in a self-service mode. The specific functionality is based upon Active Directory group memberships and permissions. The certificate subscriber and certificate manager Web portals are both accessed through the same universal resource locator (URL); however, the content displayed is based on a user's roles and permissions.
certificate template	A Windows construct that specifies the format and content of certificates based on their intended usage. When requesting a certificate from a Windows enterprise certification authority (CA), certificate requestors can select from a variety of certificate types that are based on certificate templates.
change applier	An object that performs conflict detection, conflict handling, and change application for a batch of changes.
change propagation	The process of applying changes from one replica to another.
change script	A text file that contains SQL statements for all changes made to a database, in the order in which they were made, during an editing session.
change unit	The minimal unit of change tracking in a store. In change propagation, only the units that are changed must be sent; whereas, in conflict detection, independent changes to the same unit are considered a conflict.
changing dimension	A dimension that has a flexible member structure, and is designed to support frequent changes to structure and data.
character encoding	A one-to-one mapping between a set of characters and a set of numbers.
character set	A grouping of alphabetic, numeric, and other characters that have some relationship in common. For example, the standard ASCII character set includes letters, numbers, symbols, and control codes that make up the ASCII coding scheme.
chart data region	A report item on a report layout that displays data in a graphical format.
checkpoint	An event in which the Database Engine writes dirty buffer pages to disk. Each checkpoint writes to disk all the pages that were dirty at the last checkpoint and still have not been written to disk.

Term	Definition
checksum	A calculated value that is used to test data for the presence of errors that can occur when data is transmitted or when it is written to disk. The checksum is calculated for a given chunk of data by sequentially combining all the bytes of data with a series of arithmetic or logical operations. After the data is transmitted or stored, a new checksum is calculated in the same way using the (possibly faulty) transmitted or stored data. If the two checksums do not match, an error has occurred, and the data should be transmitted or stored again. Checksums cannot detect all errors, and they cannot be used to correct erroneous data.
child	In a tree structure, the relationship of a node to its immediate predecessor.
chronicle	A table that stores state information for a single application. An example is an event chronicle, which can store event data for use with scheduled subscriptions.
chunk	A specified amount of data.
CIDER Shell	A UI container with tabbed interface for hosting the TSQLEditor component and the related output from TSQL query execution.
claims identity	A unique identifier that represents a specific user, application, computer, or other entity, enabling it to gain access to multiple resources, such as applications and network resources, without entering credentials multiple times. It also enables resources to validate requests from an entity.
clean shutdown	A system shutdown that occurs without errors.
clear text	Data in its unencrypted or decrypted form.
cleartext	Data in its unencrypted or decrypted form.
CLI	The interface supported by ODBC for use by an application.
clickstream analysis	Clickstream data are information that users generate as they move from page to page and click on items within a Web site, usually stored in log files. Web site designers can use clickstream data to improve users' experiences with a site.
clickthrough report	A report that displays related report model data when you click data within a rendered Report Builder report.
client	A service, application, or device that wants to integrate into the Microsoft Sync Framework architecture.
client	A computer or program that connects to or requests the services of another computer or program.
client code generation	The action of generating code for the client project based on operations and entities exposed in the middle tier. A RIA Services link must exist between the client and server projects.
client cursor	A cursor that is implemented on the client. The entire result set is first transferred to the client, and the client API software implements the cursor functionality from this cached result set.
Client Statistics pane	One of the tabs that hosts the output of the client statistics information.

Term	Definition
client subscription	A subscription to a merge publication that uses the priority value of the Publisher for conflict detection and resolution.
client type	Information that determines how a cache client functions and impacts the performance of your application. There are two client types: a simple client type and a routing client type.
CLM Audit	A Certificate Lifecycle Manager (CLM) extended permission in Active Directory that allows the generation and display of CLM policy templates, defining management policies within a profile template, and generating CLM reports.
CLM credentials	User account information that can be used to authenticate a user to Certificate Lifecycle Manager (CLM). These credentials can be in the form of domain credentials or one-time passwords.
CLM Enroll	A Certificate Lifecycle Manager (CLM) extended permission in Active Directory that allows the user to specify the workflow and the data to be collected while issuing certificates using a template. This extended permission only applies to profile templates.
CLM Enrollment Agent	A Certificate Lifecycle Manager (CLM) extended permission in Active Directory that allows a user or group to perform certificate requests on behalf of another user. The issued certificate's subject will contain the target user's name, rather than the requestor's name.
CLM Recover	A Certificate Lifecycle Manager (CLM) extended permission in Active Directory that allows the initiation of encryption key recovery from the certification authority database.
CLM Renew	A Certificate Lifecycle Manager (CLM) extended permission in Active Directory that allows the initiation, running, or completion of an enrollment request. The renew request replaces a user's certificate that is near its expiration date with a new certificate that has a new validity period.
CLM reports	Audit information pertaining to credential management activities within Certificate Lifecycle Manager (CLM).
CLM Request Enroll	A Certificate Lifecycle Manager (CLM) extended permission in Active Directory that allows the initiation, running, or completion of an enrollment request.
CLM Request Recover	A Certificate Lifecycle Manager (CLM) extended permission in Active Directory that allows the initiation of encryption key recovery from the certification authority database.
CLM Request Renew	A Certificate Lifecycle Manager (CLM) extended permission in Active Directory that allows the initiation, running, or completion of an enrollment request. The renew request replaces a user's certificate that is near its expiration date with a new certificate that has a new validity period.
CLM Request Revoke	A Certificate Lifecycle Manager (CLM) extended permission in Active Directory that allows the revocation of a certificate before the expiration of the certificate's validity period. An example of when this is necessary is if a user's computer or smart card is compromised (stolen).
CLM Request Unblock Smart Card	A Certificate Lifecycle Manager (CLM) extended permission in Active Directory that enables a smart card's User Personal Identification Number (PIN) to be reset, allowing access to the key material on a smart card and for that material to be re-established.

Term	Definition
CLM Revoke	A Certificate Lifecycle Manager (CLM) extended permission in Active Directory that allows the revocation of a certificate before the expiration of the certificate's validity period. An example of when this is necessary is if a user's computer or smart card is compromised (stolen).
clock vector	A collection of clock vector elements that represents updates to a replica. Any change that occurs between 0 and the tick count is contained in the vector.
clock vector element	A pair of values, consisting of a replica key and a tick count, that represents a change to a replica.
CLR function	A function created against a SQL Server assembly whose implementation is defined in an assembly created in the .NET Framework common language runtime (CLR).
CLR stored procedure	A stored procedure created against a SQL Server assembly whose implementation is defined in an assembly created in the .NET Framework common language runtime (CLR).
CLR trigger	A trigger created against a SQL Server assembly whose implementation is defined in an assembly created in the .NET Framework common language runtime (CLR).
CLR user-defined type	A user-defined data type created against a SQL Server assembly whose implementation is defined in an assembly created in the .NET Framework common language runtime (CLR).
cluster configuration storage location	The shared location (or shared storage location) where cluster configuration information is persisted. It can be a shared file or a database.
cluster disk resource	A disk on a cluster storage device.
cluster node	An individual computer in a server cluster.
cluster port	A TCP/IP port used by the cache hosts to manage the cache cluster. The port number used for the cluster ports can be different on each cache host. These settings are maintained in the cluster configuration settings.
cluster repair	A repair operation in which all missing or corrupt files are replaced, all missing or corrupt registry keys are replaced and all missing or invalid configuration values are set to default values.
clustered index	A B-tree-based index in which the logical order of the key values determines the physical order of the corresponding rows in a table.
clustered server	A server that belongs to a server cluster.
clustering	A data mining technique that analyzes data to group records together according to their location within the multidimensional attribute space.
coarse-grained lock	A lock that applies to a large amount of code or data.

<b>Term</b>	<b>Definition</b>
code access security	A mechanism provided by the common language runtime whereby managed code is granted permissions by security policy and these permissions are enforced, helping to limit the operations that the code will be allowed to perform.
code element	The minimum bit combination that can represent a unit of encoded text for processing or exchange.
code page	A table that relates the character codes (code point values) used by a program to keys on the keyboard or to characters on the display. This provides support for character sets and keyboard layouts for different countries or regions.
code point	The minimum bit combination that can represent a unit of encoded text for processing or exchange.
cold latency	The time that elapses when the workflow is being used for the first time and the XOML or XAML needs to be compiled.
cold standby	A second data center that can provide availability within hours or days.
collation	A set of rules that determines how data is compared, ordered, and presented.
collection	An object that contains a set of related objects. An object's position in the collection can change whenever a change occurs in the collection; therefore, the position of any specific object in a collection may vary.
collection item	An instance of a collector type that is created with a specific set of input properties and collection frequency, and that is used to gather specific types of data.
collection mode	The frequency at which data is collected and uploaded to the management data warehouse.
collection set	A group of collection items with which a user can interact through the user interface.
collector type	A logical wrapper around the SQL Server Integration Services packages that provide the actual mechanism for collecting data and uploading it to the management data warehouse.
collocate	To select a partitioned table that contains related data and join with this table on the partitioning column.
collocation	A condition whereby partitioned tables and indexes are partitioned according to equivalent partition functions.
color range	The range of colors available to a display device.
color rule	A rule that applies to fill colors for polygons, lines, and markers that represent points or polygon center points.
color scale	A scale that displays the results of color rules only.
column	The area in each row of a database table that stores the data value for some attribute of the object modeled by the table.

<b>Term</b>	<b>Definition</b>
column binding	The binding of an Analysis Services object to a column in a data source view.
column delimiter	A character which separates columns from each other in the CSV file being imported/exported.
column filter	A filter that restricts columns that are to be included as part of a snapshot, transactional, or merge publication.
Column Pattern Profile	A report containing a set of regular expressions that cover the specified percentage of values in a string column.
column set	An untyped XML representation that combines all the sparse columns of a table into a structured output.
column-level collation	Supporting multiple collations in a single instance.
column-level constraint	A constraint definition that is specified within a column definition when a table is created or altered.
Columns Grid	An editable grid structure in the Table Designer that lists the columns of a table and additional information about each column.
columnstore index	Stores each column in a separate set of disk pages rather than storing multiple rows per page.
COM	An object-based programming model designed to promote software interoperability; it allows two or more applications or components to easily cooperate with one another, even if they were written by different vendors, at different times, in different programming languages, or if they are running on different computers running different operating systems.
command buffer	An area in memory in which commands entered by the user are kept. A command buffer can enable the user to repeat commands without retyping them completely, edit past commands to change some argument or correct a mistake, undo commands, or obtain a list of past commands.
command prompt	An interface between the operating system and the user in which the user types command language strings of text that are passed to the command interpreter for execution.
command relationship	Provides instructions to hardware based on natural-language questions or commands.
commit	An operation that saves all changes to databases, cubes, or dimensions made since the start of a transaction.
Commit Preview	The title of a window that displays actions to be taken during the commit operation.
committed	Characteristic of a transaction that is logged and cannot be rolled back.

Term	Definition
commodity channel index formula	A formula that calculates the mean deviation of the daily average price of a commodity from the moving average. A value above 100 indicates that the commodity is overbought, and a value below -100 indicates that the commodity is oversold.
comparator	A device for comparing two items to determine whether they are equal. In electronics, for example, a comparator is a circuit that compares two input voltages and indicates which is higher.
compilation error	An error which occurs while compiling an application. These compilation errors typically occur because syntax was entered incorrectly.
compile time	The amount of time required to perform a compilation of a program. Compile time can range from a fraction of a second to many hours, depending on the size and complexity of the program, the speed of the compiler, and the performance of the hardware.
complete database restore	A restore of a full database backup, the most recent differential database backup (if any), and the log backups (if any) taken since the full database backup.
complex event processing	The continuous and incremental processing of event streams from multiple sources based on declarative query and pattern specifications with near-zero latency.
Component Object Model	An object-based programming model designed to promote software interoperability; it allows two or more applications or components to easily cooperate with one another, even if they were written by different vendors, at different times, in different programming languages, or if they are running on different computers running different operating systems.
composable	Pertaining to the ability to form complex queries by using query components (objects or operators) as reusable building blocks. This is done by linking query components together or encapsulating query components within each other.
composed environment	A virtual environment that was created from virtual machines. Those virtual machines were created outside of Microsoft Test Manager and are already deployed on a host group.
composite index	An index that uses more than one column in a table to index data.
composite key	A key whose definition consists of two or more fields in a file, columns in a table, or attributes in a relation.
compositional hierarchy	A set of entities that are conceptually part of a hierarchy, such as a parent entity and a child entity. Data operations require that the entities be treated as a single unit.
computed column	A virtual column in a table whose value is computed at run time.
computed field	A value in a formatted notification that has been computed by using a Transact-SQL expression.



Term	Definition
COM-structured storage file	A component object model (COM) compound file used by Data Transformation Services (DTS) to store the version history of a saved DTS package.
concatenation	The process of combining two or more character strings or expressions into a single character string or expression, or combining two or more binary strings or expressions into a single binary string or expression.
concurrency	A process that allows multiple users to access and change shared data at the same time. The Entity Framework implements an optimistic concurrency model.
concurrency conflict	A conflict that occurs when the same item or change unit is changed on two different replicas that are later synchronized.
concurrency model	A way in which an application can be designed to account for concurrent operations that use the same cached data. Windows Server AppFabric supports optimistic and pessimistic concurrency models.
concurrent operation	A computer operation in which two or more processes (programs) have access to the microprocessor's time and are therefore carried out nearly simultaneously. Because a microprocessor can work with much smaller units of time than people can perceive, concurrent processes appear to be occurring simultaneously but in reality are not.
conditional expression	An expression that yields a Boolean value (true or false). Such expressions can involve comparisons (testing values for equality or, for non-Boolean values, the < [less than] or > [greater than] relation) and logical combination (using Boolean operators such as AND, OR, and XOR) of Boolean expressions.
conditional split	A restore of a full database backup, the most recent differential database backup (if any), and the log backups (if any) taken since the full database backup.
config file	A file that contains machine-readable operating specifications for a piece of hardware or software or that contains information on another file or on a specific user, such as the user's logon ID.
configuration	In reference to a single microcomputer, the sum of a system's internal and external components, including memory, disk drives, keyboard, video, and generally less critical add-on hardware, such as a mouse, modem, or printer. Software (the operating system and various device drivers), the user's choices established through configuration files such as the AUTOEXEC.BAT and CONFIG.SYS files on IBM PCs and compatibles, and sometimes hardware (switches and jumpers) are needed to "configure the configuration" to work correctly. Although system configuration can be changed, as by adding more memory or disk capacity, the basic structure of the system--its architecture--remains the same.
configuration file	A file that contains machine-readable operating specifications for a piece of hardware or software or that contains information on another file or on a specific user, such as the user's logon ID.
Configuration Tools	In SQL Server, a menu item which allows the user to enable, disable, start, or stop the features, services, and remote connectivity of the SQL Server installations.
conflict detection	The process of determining which operations were made by one replica without knowledge of the other, such as when two replicas make local updates to the same item.

Term	Definition
conflict resolution method	The method that is used to determine which change is written to the store in the event of a conflict. Typical conflict resolution methods are as follows: last writer wins, source wins, destination wins, custom, or deferred. For custom resolution, the resolving application reads the conflict from the conflict log and selects a resolution. For deferred resolution, the conflict is logged together with the conflicting change data and the made-with knowledge of the change.
conflict resolver	A special mechanism which handles resolving of conflict situations.
Connection Director	A connectivity technology where applications based on different data access technologies (.NET or native Win32) can share the same connection information. Connection information can be centrally managed for such client applications.
connection manager	A logical representation of a run-time connection to a data source.
connection string	A series of arguments that define the location of a resource and how to connect to it.
consistency unit	The minimal unit of data synchronization. Because all changes that have the same consistency unit are sent together, synchronization can never be interrupted with part of a consistency unit applied.
constant	A numeric or string value that is not calculated and, therefore, does not change.
constraint conflict	A conflict that violates constraints that are put on items or change units, such as the relationship of folders or the location of identically-named data within a file system.
constraint violation	A violation that occurs when the restriction criteria are not satisfied.
contained database	A SQL Server database that includes all of the user authentication, database settings, and metadata required to define and access the database, and has no configuration dependencies on the instance of the SQL Server Database Engine where the database is installed.
container	A control flow element that provides package structure.
content formatter	The part of the distributor that turns raw notification data into readable messages.
content key	The cryptographic key used to both encrypt and decrypt protected content during publishing and consumption.
contention	On a network, competition among stations for the opportunity to use a communications line or network resource.
Context pane	A tree view included in the Table Designer that lists objects related to a table.
context switch	The changing of the identity against which permissions to execute statements or perform actions are checked.

<b>Term</b>	<b>Definition</b>
continuation media	The series of removable backup media used after the initial medium becomes full, allowing continuation of the backup operation.
continuation tape	A tape that is used after the initial tape in a media family fills, allowing continuation of a media family.
contract	A Service Broker object that defines the message types that can be exchanged within a given conversation.
control flow	A group of connected control flow elements that perform tasks.
control-break report	A report that summarizes data in user-defined groups or breaks. A new group is triggered when different data is encountered.
control-of-flow language	Transact-SQL keywords that control the flow of execution of SQL statements and statement blocks in triggers, stored procedures, and batches.
conversation endpoint	The object which represents a party participating in the conversation.
conversation group	A group of related Service Broker conversations. Messages in the same conversation group can only be processed by one service program at a time.
conversation handle	An handle which uniquely defines a conversation.
convex hull	The smallest convex set that contains X in the Euclidean plane or Euclidean space.
coordinate system	In n-dimensional space, a set of n linearly independent vectors anchored to a point (called the origin). A group of coordinates specifies a point in space (or a vector from the origin) by indicating how far to travel along each vector to reach the point (or tip of the vector).
correlated subquery	A subquery that references a column in the outer statement. The inner query is run for each candidate row in the outer statement.
count window	A window with a variable window size that moves along a timeline with each distinct event start time.
countersign	To sign a document already signed by the other party.
CPU busy	A SQL Server statistic that reports the time, in milliseconds, that the central processing unit (CPU) spent on SQL Server work.
crawl	The process of scanning content to compile and maintain an index.
credentials	Information that includes identification and proof of identification that is used to gain access to local and network resources. Examples of credentials are user names and passwords, smart cards, and certificates.

Term	Definition
cross-database ownership chaining	An ownership chain that spans more than one database.
cross-validation	A method for evaluating the accuracy of a data mining model.
CTI event	A special punctuation event that indicates the completeness of the existing events.
cube	A set of data that is organized and summarized into a multidimensional structure that is defined by a set of dimensions and measures.
cube role	A collection of users and groups with the same access to a cube. A cube role is created when you assign a database role to a cube, and it applies only to that cube.
current time increment event	A special punctuation event that indicates the completeness of the existing events.
cursor	An entity that maps over a result set and establishes a position on a single row within the result set.
cursor degradation	The return of a different type of cursor than the user had declared.
cursor library	A part of the ODBC and DB-Library application programming interfaces (APIs) that implements client cursors
custom rollup	An aggregation calculation that is customized for a dimension level or member, and that overrides the aggregate functions of a cube's measures.
custom rule	In a role, a specification that limits the dimension members or cube cells that users in the role are permitted to access.
custom variable	A variable provided by package developers.
custom volume	A volume that is not in the DPM storage pool and is specified to store the replica and recovery points for a protection group member.
cyclic protection	A type of protection between two DPM servers where each server protects the data on the other.
DAC	A application that captures the SQL Server database and instance objects used by a client-server or 3-tier application.
DAC instance	A copy of a DAC deployed on an instance of the Database Engine. There can be multiple DAC instances on the same instance of the Database Engine.
DAC package	An XML manifest that contains all of the objects defined for the DAC; the package gets created when a developer builds a DAC project.

Term	Definition
DAC package file	The XML file that is the container of a DAC package.
DAC placement policy	A PBM policy that comprises a set of conditions, which serve as prerequisites on the target instance of SQL Server where the DAC can be deployed.
DAC project	A Visual Studio project used by database developers to create and develop a DAC. DAC projects get full support from Visual Studio and VSTS source code control, versioning, and development project management.
data adapter	An object used to submit data to and retrieve data from databases, Web services, and Extensible Markup Language (XML) files.
data backup	Any backup that includes the full image of one or more data files.
data binning	The process of grouping data into specific bins or groups according to defined criteria.
data block	In text, ntext, and image data, a data block is the unit of data transferred at one time between an application and an instance of SQL Server. The term is also applied to the units of storage for these data types.
data co-location	In DPM, a feature that enables protection of multiple data sources on a single volume or on the same tape. This allows you to store more data on each volume or tape.
data connection	A connection that specifies the name, type, location, and, optionally, other information about a database file or server.
Data Control Language	The subset of SQL statements used to control permissions on database objects.
data convergence	Data at the Publisher and the Subscriber that matches.
data corruption	A process wherein data in memory or on disk is unintentionally changed, with its meaning thereby altered or obliterated.
data definition	The attributes, properties, and objects in a database.
data definition language	A language that defines all attributes and properties of a database, especially record layouts, field definitions, key fields, file locations, and storage strategy.
data description language	A language that defines all attributes and properties of a database, especially record layouts, field definitions, key fields, file locations, and storage strategy.
data dictionary	A database containing data about all the databases in a database system. Data dictionaries store all the various schema and file specifications and their locations. They also contain information about which programs use which data and which users are interested in which reports.

<b>Term</b>	<b>Definition</b>
data element	A single unit of data.
data explosion	The exponential growth in size of a multidimensional structure, such as a cube, due to the storage of aggregated data.
data extension	A plug-in that processes data for a specific kind of data source. For example, Microsoft OLE DB Provider for DB2.
data feed	An XML data stream in Atom 1.0 format.
data flow	The movement of data through a group of connected elements that extract, transform, and load data.
data flow component	A component of SQL Server 2005 Integration Services that manipulates data.
data flow engine	An engine that executes the data flow in a package.
data flow task	The task that encapsulates the data flow engine that moves data between sources and destinations, providing the facility to transform, clean, and modify data as it is moved.
data integrity	The accuracy of data and its conformity to its expected value, especially after being transmitted or processed.
data manipulation language	The subset of SQL statements that is used to retrieve and manipulate data. DML statements typically start with SELECT INSERT UPDATE or DELETE.
data mart	A subset of the contents of a data warehouse that tends to contain data focused at the department level, or on a specific business area.
data member	A child member associated with a parent member in a parent-child hierarchy. A data member contains the data value for its parent member, rather than the aggregated value for the parent's children.
data mining	The process of identifying commercially useful patterns or relationships in databases or other computer repositories through the use of advanced statistical tools.
data mining extension	In Analysis Services, a statement that performs mining tasks programmatically.
data mining model training	The process a data mining model uses to estimate model parameters by evaluating a set of known and predictable data.
data processing extension	A plug-in that processes data for a specific kind of data source (similar to a database driver).
Data Processor component	A component of the report server engine that processes data.

<b>Term</b>	<b>Definition</b>
Data Profile Viewer	A stand-alone utility that displays the profile output in both summary and detail format with optional drilldown capability.
data provider	A known data source specific to a target type that provides data to a collector type.
data pump	A component used in SQL Server 2000 Transformation Services (DTS) to import, export, and transform data between heterogeneous data stores.
data region	A report item that provides data manipulation and display functionality for iterative data from an underlying dataset.
data scrubbing	The process of building a data warehouse out of data coming from multiple online transaction processing (OLTP) systems.
data segment	The portion of memory or auxiliary storage that contains the data used by a program.
data set	A collection of related information made up of separate elements that can be treated as a unit in data handling.
data source view	A named selection of database objects--such as tables, views, relationships, and stored procedures, based on one or more data sources--that defines the schema referenced by OLAP and data mining objects in an Analysis Services databases. It can also be used to define sources, destinations, and lookup tables for DTS tasks, transformations, and data adapters.
data steward	The person responsible for maintaining a data element in a metadata registry.
data tap	Collecting data from a specified path in SQL Server Integration Services. Tapped data can be written to CSV files when the SSIS package executes, and customers can specify which data flow to tap.
data viewer	A graphical tool that displays data as it moves between two data flow components at run time.
data warehouse	The database that stores operations data for long periods of time. This data is then used by the Operations Manager reporting server to build reports. By default, this database is named OperationsManagerDW.
Data Warehouse	The database that stores operations data for long periods of time. This data is then used by the Operations Manager reporting server to build reports. By default, this database is named OperationsManagerDW.
database administrator	The person who manages a database. The administrator determines the content, internal structure, and access strategy for a database, defines security and integrity, and monitors performance.
database catalog	The part of a database that contains the definition of all the objects in the database, as well as the definition of the database.
database diagram	A graphical representation of any portion of a database schema. It can be either a whole or partial picture of the structure of the database. It includes tables, the columns they contain, and the relationships between the tables.
database engine	The program module or modules that provide access to a database management system (DBMS).

Term	Definition
Database Engine Tuning Advisor	A tool for tuning the physical database design that helps users to select and create an optimal set of indexes, indexed views, and partitioning.
Database Explorer	A simple database administration tool that lets the user perform database operations such as creating new tables, querying and modifying existing data, and other database development functions.
database file	One of the physical files that make up a database.
database language	The language used for accessing, querying, updating, and managing data in relational database systems.
database management system	A layer of software between the physical database and the user. The DBMS manages all access to the database.
database manager	A layer of software between the physical database and the user. The DBMS manages all access to the database.
database mirroring	Immediately reproducing every update to a read-write database (the principal database) onto a read-only mirror of that database (the mirror database) residing on a separate instance of the database engine (the mirror server). In production environments, the mirror server is on another machine. The mirror database is created by restoring a full backup of the principal database (without recovery).
Database Mirroring Monitor	A tool used to monitor any subset of the mirrored databases on a server instance.
database mirroring partner	One in a pair of server instances that act as role-switching partners for a mirrored database.
database mirroring partners	A pair of server instances that act as role-switching partners for a mirrored database.
database project	A collection of one or more data connections (a database and the information needed to access that database).
database reference	A path, expression or filename that resolves to a database.
database role	A collection of users and groups with the same access to an Analysis Services database.
database schema	The names of tables, fields, data types, and primary and foreign keys of a database.
database script	A collection of statements used to create database objects.



<b>Term</b>	<b>Definition</b>
database snapshot	A read-only, static view of a database at the moment of snapshot creation.
database structure	The names of tables, fields, data types, and primary and foreign keys of a database.
Database view	A read-only, static snapshot of a source database at the moment of the view's creation.
data-definition query	An SQL-specific query that contains data definition language (DDL) statements. These statements allow you to create or alter objects in the database.
data-driven subscription	A subscription that takes generated output for subscription values (for example, a list of employee e-mail addresses).
datareader	A stream of data that is returned by an ADO.NET query.
data-tier application	A application that captures the SQL Server database and instance objects used by a client-server or 3-tier application.
data-tier application instance	A copy of a DAC deployed on an instance of the Database Engine. There can be multiple DAC instances on the same instance of the Database Engine.
data-tier application package	An XML manifest that contains all of the objects defined for the DAC; the package gets created when a developer builds a DAC project.
date	A SQL Server system data type that stores a date value from January 1, 1 A.D., through December 31, 9999
DB	A collection of data formatted/arranged to allow for easy search and retrieval.
DBCS	A character set that can use more than one byte to represent a single character. A DBCS includes some characters that consist of 1 byte and some characters that consist of 2 bytes. Languages such as Chinese, Japanese, and Korean use DBCS.
DBMS	A layer of software between the physical database and the user. The DBMS manages all access to the database.
DDL	A language that defines all attributes and properties of a database, especially record layouts, field definitions, key fields, file locations, and storage strategy.
DDL trigger	A special kind of trigger that executes in response to Data Definition Language (DDL) statements.
deadlock	A situation when two users, each having a lock on one piece of data, attempt to acquire a lock on the other's piece.
decision support	Systems designed to support the complex analytic analysis required to discover business trends.

Term	Definition
decision tree	A treelike model of data produced by certain data mining methods. Decision trees can be used for prediction.
declaration	A binding of an identifier to the information that relates to it. For example, to make a declaration of a constant means to bind the name of the constant with its value. Declaration usually occurs in a program's source code; the actual binding can take place at compile time or run time.
Declarative Management Framework	A policy based system of SQL Server management.
Declarative Management Framework Facet	A set of logical pre-defined properties that model the behavior or characteristics for certain types of managed targets (such as a database, table, login, view,etc) in policy-based management.
declarative referential integrity	FOREIGN KEY constraints defined as part of a table definition that enforce proper relationships between tables.
dedicated administrator connection	A dedicated connection that allows an administrator to connect to a server when the Database Engine will not respond to regular connections.
default	A value that is automatically used by a program when the user does not specify an alternative. Defaults are built into a program when a value or option must be assumed for the program to function.
default database	The database the user is connected to immediately after logging in to SQL Server.
default instance	The instance of SQL Server that uses the same name as the computer name on which it is installed.
default language	The human language that SQL Server uses for errors and messages if a user does not specify a language.
default member	The dimension member used in a query when no member is specified for the dimension.
default result set	The default mode that SQL Server uses to return a result set back to a client.
defection	The removal of a server from multiserver operations.
deferred transaction	A transaction that is not committed when the roll forward phase of recovery completes and that cannot be rolled back during database startup because data needed by roll back is offline. This data can reside in either a page or a file.
degenerate dimension	A relationship between a dimension and a measure group in which the dimension main table is the same as the measure group table.

<b>Term</b>	<b>Definition</b>
DEK	A bit string that is used in conjunction with an encryption algorithm to encrypt and decrypt data.
delegated registration model	A registration model in which a person other than the certificate subscriber initiates the certificate transaction. The certificate subscriber then completes the transaction by providing a supplied one-time password.
DELETE clause	A part of a DML Statement that contains the DELETE keyword and associated parameters.
delete level	In Data Transformation Services, the amount and kind of data to remove from a data warehouse.
delimited identifier	An object in a database that requires the use of special characters (delimiters) because the object name does not comply with the formatting rules of regular identifiers.
delivery channel	A pipeline between a distributor and a delivery service.
delivery channel type	The protocol for a delivery channel, such as Simple Mail Transfer Protocol (SMTP) or File.
delivery extension	A plug-in that delivers reports to a specific target (for example, e-mail delivery).
delivery protocol	The set of communication rules used to route notification messages to external delivery systems.
denormalize	To introduce redundancy into a table to incorporate data from a related table.
deploy	To build a DAC instance, either directly from a DAC package or from a DAC previously imported to the SQL Server Utility.
deployed environment	A group of virtual machines located on a team project host group and controlled by Microsoft Test Manager. A deployed environment can be running or stopped.
dequeue	To remove from a queue.
derived column	A transformation that creates new column values by applying expressions to transformation input columns.
deserializati on	The process of converting an object from a serial storage format to binary format in the form of an object that applications can use. This happens when the object is retrieved from the cache cluster with the Get client APIs.
destination	The SSIS data flow component that loads data into data stores or creates in-memory datasets.
destination	A synchronization provider that provide its current knowledge, accept a list of changes from the source provider, detect any conflicts between that list and its own items, and apply changes to its data store.
destination adapter	A data flow component that loads data into a data store.

Term	Definition
destination provider	A synchronization provider that provide its current knowledge, accept a list of changes from the source provider, detect any conflicts between that list and its own items, and apply changes to its data store.
detect	To find something.
device type	A value from a developer-defined list that specifies the types of devices that a given application will support.
diacritic	A mark placed over, under, or through a character, usually to indicate a change in phonetic value from the unmarked state.
diacritical mark	A mark placed over, under, or through a character, usually to indicate a change in phonetic value from the unmarked state.
dialect	The syntax and general rules used to parse a string or a query statement.
diamond-shape relationship	A chain of attribute relationships that splits and rejoins but that contains no redundant relationships. For example, Day->Month->Year and Day->Quarter->Year have the same start and end points, but do not have any common relationships.
differencer	An interface to a tool that creates a DifferencingService object.
differential backup	A backup containing only changes made to the database since the preceding data backup on which the differential backup is based.
differential base	The most recent full backup of all the data in a database or in a subset of the files or filegroups of a database.
digest delivery	A method of sending notifications that combines multiple notifications within a batch and sends the resulting message to a subscriber.
digital certificate	A digital document that is commonly used for authentication and to help secure information on a network. A certificate binds a public key to an entity that holds the corresponding private key. Certificates are digitally signed by the certification authority that issues them, and they can be issued for a user, a computer, or a service.
dimension	A structural attribute of a cube that organizes data into levels. For example, a Geography dimension might include the members Country, Region, State or Province, and City.
dimension expression	A valid Multidimensional Expressions (MDX) expression that returns a dimension.
dimension granularity	The lowest level available to a particular dimension in relation to a particular measure group. The "natural" or physical grain is always that of the key that joins the main dimension table to the primary fact table.
dimension hierarchy	A logical tree structure that organizes the members of a dimension such that each member has one parent member and zero or more child members.
dimension level	The name of a set of members in a dimension hierarchy such that all members of the set are at the same distance from the root of the hierarchy. For example, a time hierarchy may contain the levels Year, Month, and Day.

Term	Definition
dimension member	A single position or item in a dimension. Dimension members can be user-defined or predefined and can have properties associated with them.
dimension member property	A characteristic of a dimension member. Dimension member properties can be alphanumeric, Boolean, or Date/Time data types, and can be user-defined or predefined.
dimension table	A table in a data warehouse whose entries describe data in a fact table.
direct connect	The state of being connected to a back-end database, so that any changes you make to a database diagram automatically update your database when you save the diagram or selected items in it.
direct response mode	The default mode in which SQL Server statistics are gathered separately from the SQL Server Statistics display.
dirty page	A buffer page that contains modifications that have not been written to disk.
dirty read	A read that contains uncommitted data.
disabled index	Any index that has been marked as disabled. A disabled index is unavailable for use by the database engine. The index definition of a disabled index remains in the system catalog with no underlying index data.
discrete signal	A time series consisting of a sequence of quantities, that is a time series that is a function over a domain of discrete integers.
discretize	To put values of a continuous set of data into groups so that there are a discrete number of possible states.
discretized column	A column that represents finite, counted data
distinct count measure	A measure commonly used to determine for each member of a dimension how many distinct, lowest-level members of another dimension share rows in the fact table.
distributed partitioned view	A view that joins horizontally partitioned data from a set of member tables across more than one server, making the data appear as if from one table.
distributed query	A single query that accesses data from multiple data sources.
distributed transaction	A transaction that spans multiple data sources.
distribution cleanup agent	A scheduled job that runs under SQL Server Agent. After all Subscribers have received a transaction, the agent removes the transaction from the distribution database. It also cleans up snapshot files from the file system after entries corresponding to those files have been removed from the distribution database.

Term	Definition
distribution database	A database on the Distributor that stores data for replication including transactions, snapshot jobs, synchronization status, and replication history information.
distribution retention period	In transactional replication, the amount of time transactions are stored in the distribution database.
distributor	A database instance that acts as a store for replication-specific data associated with one or more Publishers.
DMF	One of a set of built-in functions that returns server state information about values, objects, and settings in SQL Server.
DML	The subset of SQL statements that is used to retrieve and manipulate data. DML statements typically start with SELECT INSERT UPDATE or DELETE.
DML trigger	A stored procedure that executes when data in a specified table is modified.
DMV	A set of built-in views that return server state information about values, objects, and settings in SQL Server.
DMX	In Analysis Services, a statement that performs mining tasks programmatically.
domain	The set of possible values that you can specify for an independent variable in a function, or for a database attribute.
domain	A collection of computers in a networked environment that share a common database, directory database, or tree. A domain is administered as a unit with common rules and procedures, which can include security policies, and each domain has a unique name.
domain context	A client-side representation of a domain service.
domain integrity	The validity of entries for a specific column of data.
domain operation	A method on a domain service that is exposed to a client application. It enables client applications to perform an action on the entity such as, query, update, insert, or delete records.
domain service	A service that encapsulates the business logic of an application. It exposes a set of related domain operations in a service layer.
dormant session	Session in pre-login state. Sessions can be initiated or ended to modify their state, but they generally remain in either a "sleep/idle" state, such as when the session has been initiated and is open at the server for client use; or a "dormant" state, such as when the session has been ended and the session is not currently available at the server for client use.
double byte character set	A character set that can use more than one byte to represent a single character. A DBCS includes some characters that consist of 1 byte and some characters that consist of 2 bytes. Languages such as Chinese, Japanese, and Korean use DBCS.

Term	Definition
double-byte character set	A character set that can use more than one byte to represent a single character. A DBCS includes some characters that consist of 1 byte and some characters that consist of 2 bytes. Languages such as Chinese, Japanese, and Korean use DBCS.
down	Not functioning, in reference to computers, printers, communications lines on networks, and other such hardware.
download-only article	An article in a merge publication that can be updated only at the Publisher or at a Subscriber that uses a server subscription.
DPM Client	The Data Protection Manager (DPM) Client enables the user to protect and recover their data as per the company protection policy configured by their backup administrator.
DPM engine	A policy-driven engine that DPM uses to protect and recover data.
DPM Management Shell	The command shell, based on Windows PowerShell (Powershell.exe), that makes available the cmdlets that perform functions in Data Protection Manager.
DPM Online	A feature that provides an online remote backup, both for securely storing data offsite for long durations and for disaster recovery (DR).
DPM Online account	A user account that DPM uses to start the DPM Online service.
DPM Online cache	A cache volume required by DPM Online on which the DPM server stores information for faster backup and recovery from DPM Online.
DPM Online protection	The process of using DPM online to protect data from loss or corruption by creating and maintaining replicas and recovery points of the data online.
DPM Online protection group	A collection of data sources that share the same DPM Online protection configuration.
DPM Online recovery	The process by which an administrator recovers previous versions of protected data from online recovery points on an internet-connected DPM server.
DPM Online recovery point	The date and time of a previous version of a data source that is available from DPM Online.
DPM Online replica	A complete copy of an online-protected data source on DPM Online. Each member of an online protection group on the DPM server is associated with a DPM Online replica.
DPM role	The grouping of users, objects, and permissions that is used by DPM administrators to manage DPM features that are used by end users.

Term	Definition
DPM Self-Service Recovery Configuration Tool	A tool that enables DPM administrators to authorize end users to perform self-service recovery of data by creating and managing DPM roles (grouping of users, objects, and permissions).
DPM Self-Service Recovery Tool	A tool that is used by end users to recover backups from DPM, without any action required from the DPM administrator.
DPM Self-Service Tool for SQL Server	A tool for SQL Server that enables backup administrators to authorize end users to recover backups of SQL Server databases from DPM, without further action from the backup administrator.
DPM SRT	Software provided with DPM to facilitate a Windows Server 2003 bare metal recovery for the DPM server and the computers that DPM protects.
DPM SST for SQL Server	A tool for SQL Server that enables backup administrators to authorize end users to recover backups of SQL Server databases from DPM, without further action from the backup administrator.
DPM System Recovery Tool	Software provided with DPM to facilitate a Windows Server 2003 bare metal recovery for the DPM server and the computers that DPM protects.
DQS	A knowledge-based data-quality system that enables users to perform knowledge discovery and management, data cleansing, data matching, integration with reference data services, and integrated profiling.
DQS knowledge base	A repository of metadata that is used by Data Quality Services to improve the quality of data. This metadata is created either by the user interactively or by the Data Quality Services platform in an automated knowledge discovery process.
drillthrough	In Analysis Services, a technique to retrieve the detailed data from which the data in a cube cell was summarized.
drill-through report	A secondary report that is displayed when a user clicks an item in a report. Detailed data is displayed in the same report.
drop-down list	A list that can be opened to reveal all choices for a given field.
DSN	The collection of information used to connect an application to a particular ODBC database.
DSN-less connection	A type of data connection that is created based on information in a data source name (DSN), but is stored as part of a project or application.
dump	A duplicate of a program, a disk, or data, made either for archiving purposes or for safeguarding files.
dump device	A tape or disk drive containing a backup medium.



Term	Definition
dynamic cursor	A cursor that can reflect data modifications made to the underlying data while the cursor is open.
dynamic filter	A row filter available with merge replication that allows you to restrict the data replicated to a Subscriber based on a system function or user-defined function (for example: SUSER_SNAME()).
dynamic locking	The process used by SQL Server to determine the most cost-effective locks to use at any one time.
dynamic management function	One of a set of built-in functions that returns server state information about values, objects, and settings in SQL Server.
dynamic management view	A set of built-in views that return server state information about values, objects, and settings in SQL Server.
dynamic recovery	The process that detects and/or attempts to correct software failure or loss of data integrity within a relational database management system (RDBMS).
dynamic routing	Routing that adjusts automatically to the current conditions of a network. Dynamic routing typically uses one of several dynamic-routing protocols such as Routing Information Protocol (RIP) and Border Gateway Protocol (BGP). Compare static routing.
dynamic snapshot	In merge replication, a snapshot that includes only the data from a single partition.
eager loading	A pattern of loading where a specific set of related objects are loaded along with the objects that were explicitly requested in the query.
edge event	An event whose event payload is valid for a given interval; however, only the start time is known upon arrival to the CEP server. The valid end time of the event is provided later in a separate edge event.
effective policy	The set of enabled policies for a target.
e-mail	The exchange of text messages and computer files over a communications network, such as a local area network or the Internet.
encryption	The process of converting readable data (plaintext) into a coded form (ciphertext) to prevent it from being read by an unauthorized party.
encryption key	A bit string that is used in conjunction with an encryption algorithm to encrypt and decrypt data.
end cap	For paths that contain unconnected ends, such as lines, the end of a stroke. You can change the way the stroke looks at each end by applying one of four end cap styles: flat cap, round cap, square cap, and triangle cap.
endpoint	A synchronization provider and its associated replica.

Term	Definition
endpoint mapper	A service on a remote procedure call (RPC) server that maintains a database of dynamic endpoints and allows clients to map an interface/object UUID pair to a local dynamic endpoint.
enqueue	To place (an item) in a queue.
enroll	To add an instance of SQL Server to the set of SQL Server instances managed by a utility control point.
enrollment	The process of requesting, receiving, and installing a certificate.
Enrollment Agent	A user account used to request smart card certificates on behalf of another user account. A specific certificate template is applied to an Enrollment Agent.
enterprise license	A license that authorizes protection of both file and application resources on a single computer.
entity	In Reporting Services, a logical collection of model items, including source fields, roles, folders, and expressions, presented in familiar business terms.
entity integrity	A state in which every row of every table can be uniquely identified.
envelopes formula	A financial formula that calculates "envelopes" above and below a moving average using a specified percentage as the shift. The envelopes indicator is used to create signals for buying and selling. You can specify the percentage the formula uses to calculate the envelopes.
equijoin	A join in which the values in the columns being joined are compared for equality, and all columns are included in the results.
equirectangular projection	In a map report item, a very simple equidistant cylindrical projection in which the horizontal coordinate is the longitude and the vertical coordinate is the latitude.
error handling	The process of dealing with errors (or exceptions) as they arise during the running of a program. Some programming languages, such as C++, Ada, and Eiffel, have features that aid in error handling.
Error List	The name of a pane that shows T-SQL syntax or dependency errors.
error log	A file that lists errors that were encountered during an operation.
error state number	A number associated with SQL Server messages that helps Microsoft support engineers find the specific code location that issued the message.
ETL	The act of extracting data from various sources, transforming data to consistent types, and loading the transformed data for use by applications.
ETW-based log sink	A means of capturing trace events on the cache client or cache host with the Event Tracing for Windows (ETW) framework inside Windows.
event	The basic unit of data processed by the StreamInsight server. Each event contains a header that defines the event kind and the temporal properties of the event. An event (except the CTI event) typically contains an event payload, which is a .NET data structure that contains the data associated with the event.

Term	Definition
event	Any significant occurrence in the system or an application that requires a user to be notified or an entry to be added to a log.
event category	In SQL Trace, a grouping of similar and logically related event classes.
event chronicle	A table that stores event state information.
event chronicle rule	One or more Transact-SQL statements that manage the data in the event chronicle.
event class	In SQL Trace, a collection of properties that define an event.
event classification	A means of differentiating types of events that occur on the cache client and cache host. The Windows Server AppFabric log sinks follow the classification established with the System.Diagnostics.TraceLevel enumeration.
event collection stored procedures	System-generated stored procedures that an application can call to submit events to the event table in the application database.
event handler	A software routine that executes in response to an event.
event header	The portion of an event that defines the temporal properties of the event and the event kind. Temporal properties include a valid start time and end time associated with the event.
event kind	Event metadata that defines the event type.
event model	The event metadata that defines the temporal characteristics (shape) of the event.
event notification	A special kind of trigger that sends information about database events to a service broker.
event payload	The data portion of an event in which the data fields are defined as CLR (common language runtime) types. An event payload is a typed structure.
event provider	A provider that monitors a source of events and notifies the event table when events occur.
event source	The point of origin of an event.
event table	A table in the application database that stores event data.
Event Tracing for Windows (ETW)-based log sink	A means of capturing trace events on the cache client or cache host with the Event Tracing for Windows (ETW) framework inside Windows.

Term	Definition
Everyone	A type of user account.
eviction	The physical removal of a cached object from the memory of the cache host or hosts that it is stored on. This is typically done to keep the memory usage of the cache host service in check.
exclusive lock	A lock that prevents any other transaction from acquiring a lock on a resource until the original lock on the resource is released at the end of the transaction.
execute	To perform an instruction.
Execution Plan pane	One of the tabs that hosts the output for an estimated or actual execution plan that SQL Server uses.
execution tree	The path of data in the data flow of a SQL Server 2005 Integration Services package from sources through transformations to destinations.
exit module	A Certificate Services component that performs post-processing after a certificate is issued, such as the publication of an issued certificate to Active Directory.
expiration	The point at which an object has exceeded the cache time-out value. When an object expires, it is evicted.
explicit cap	An explicit hierarchy used as the top level of a derived hierarchy structure.
explicit hierarchy	In Master Data Services, a hierarchy that uses consolidated members to group other consolidated and leaf members.
explicit loading	A pattern of loading where related objects are not loaded until explicitly requested by using the Load method on a navigation property.
explicit transaction	A group of SQL statements enclosed within transaction delimiters that define both the start and end of the transaction.
exploded pie	A pie chart that displays the contribution of each value to a total while emphasizing individual values, by showing each slice of the pie as "pulled out," or separate, from the whole.
exploded pie chart	A pie chart that displays the contribution of each value to a total while emphasizing individual values, by showing each slice of the pie as "pulled out," or separate, from the whole.
exponential moving average	A moving average of data that gives more weight to the more recent data in the period and less weight to the older data in the period. The formula applies weighting factors which decrease exponentially. The weighting for each older data point decreases exponentially, giving much more importance to recent observations while still not discarding older observations entirely.
export format	UI text for subscriptions and HTML viewer. Corresponds to rendering extensions.
expression	Any combination of operators, constants, literal values, functions, and names of fields (columns), controls, and properties that evaluates to a single value.

Term	Definition
expression host assembly	All expressions found within a report are that are compiled into an assembly. The expression host assembly is stored as a part of the compiled report.
extended permission	A permission that is specific to an object added to the standard Active Directory object schema. The permission associated with the new object extends the existing default permission set.
Extended property	User-defined text (descriptive or instructional including input masks and formatting rules) specific to a database or database object. The text is stored in the database as a property of the database or object.
Extended Protection for Authentication	A security feature that helps protect against man-in-the-middle (MITM) attacks.
extended stored procedure	A function in a dynamic link library (DLL) that is coded using the SQL Server Extended Stored Procedure API. The function can then be invoked from Transact-SQL using the same statements that are used to execute Transact-SQL stored procedures.
Extensible Stylesheet Language	An XML vocabulary that is used to transform XML data to another form, such as HTML, by means of a style sheet that defines presentation rules.
Extensible Stylesheet Language Transformation	A declarative, XML-based language that is used to present or transform XML data.
Extensible Stylesheet Language Transformations	A declarative, XML-based language that is used to present or transform XML data.
extent	On a disk or other direct-access storage device, a continuous block of storage space reserved by the operating system for a particular file or program.
external delivery system	A system, such as Microsoft Exchange Server, that delivers formatted notifications to destination devices.
extract	To build a DAC package file that contains the definitions of all the objects in an existing database, as well as instance objects that are associated with the database.
extraction, transformation, and loading	The act of extracting data from various sources, transforming data to consistent types, and loading the transformed data for use by applications.

<b>Term</b>	<b>Definition</b>
facet	A set of logical pre-defined properties that model the behavior or characteristics for certain types of managed targets (such as a database, table, login, view,etc) in policy-based management.
Facet	A set of logical pre-defined properties that model the behavior or characteristics for certain types of managed targets (such as a database, table, login, view,etc) in policy-based management.
facet property	A predefined property that applies to a specific facet in Policy-Based Management.
fact	A row in a fact table in a data warehouse. A fact contains values that define a data event such as a sales transaction.
fact dimension	A relationship between a dimension and a measure group in which the dimension main table is the same as the measure group table.
fact table	A central table in a data warehouse schema that contains numerical measures and keys relating facts to dimension tables.
factory method	A method, usually defined as static, whose purpose is to return an instance of a class.
fail over	To switch processing from a failed component to its backup component.
failed transaction	A transaction that encountered an error and was not able to complete.
failover cluster	A group of servers that are in one location and that are networked together for the purpose of providing live backup in case one of the servers fails.
failover clustering	A high availability process in which an instance of an application or a service, running over one machine, can fail-over onto another machine in the failover cluster in the case the first one fails.
failover partner	The server used if the connection to the partner server fails.
fail-safe operator	A user who receives the alert if the designated operator cannot be reached.
failure notification	A type of cache notification triggered when the cache client misses one or more cache notifications.
fatal error	An error that causes the system or a program to fail abruptly with no hope of recovery. An example of a fatal error is an uncaught exception that cannot be handled.
federated database servers	A set of linked servers that shares the processing load of data by hosting partitions of a distributed partitioned view.
feed consumer	A software component that extracts items from a FeedSync feed and applies them to a destination replica by using a synchronization provider.

Term	Definition
fiber	A Windows NT lightweight thread scheduled within a single OS thread.
fiber mode	A situation where an instance of SQL Server allocates one Windows thread per SQL scheduler, and then allocates one fiber per worker thread, up to the value set in the max worker threads option.
field	An area in a window or record that stores a single data value.
field length	In bulk copy, the maximum number of characters needed to represent a data item in a bulk copy character format data file.
field marshaller	A SQL Server feature that handles marshaling for fields.
field terminator	In bulk copy, one or more characters marking the end of a field or row, separating one field or row in the data file from the next.
file backup	A backup of all the data in one or more files or filegroups.
file differential backup	A backup of one or more files containing only changes made to each file since its most recent file backup. A file differential backup requires a full file backup as a base.
file DSN	file Data Source Names. File-based data sources shared among all users with the same drivers installed. These data sources are not dedicated to a user or local to a computer.
file mapping	The association of a file's contents with a portion of the virtual address space of a process.
file restore	An operation that restores one or more files of a database.
file rollover	The process when a program closes a file, based on a certain event, and creates a new file.
filegroup	A named collection of one or more data files that forms a single unit of data allocation or for administration of a database.
file-mapping object	An object that maintains the association of a file's contents with a portion of the virtual address space of a process.
filestream	A sequence of bytes used to hold file data.
fill factor	An attribute of an index that defines how full the SQL Server Database Engine should make each page of the index.
filter forgotten knowledge	The knowledge that is used as the starting point for filter tracking. A filter-tracking replica can save storage space by removing ghosts and advancing the filter forgotten knowledge to contain the highest version of the ghosts that have been removed.
filter key	A 4-byte value that maps to a filter in a filter key map.

Term	Definition
filtered replica	A replica that stores item data only for items that are in a filter, such as a media storage replica that stores only songs that are rated as three stars or better.
filter-tracking replica	A replica that can identify which items are in a filter and which have moved in or out of the filter recently.
fine-grained lock	A lock that applies to a small amount of code or data.
fit	One of the criteria used for evaluating the success of a data mining algorithm. Fit is typically represented as a value between 0 and 1, and is calculated by taking the covariance between the predicted and actual values of evaluated cases and dividing by the standard deviations of the same predicted and actual values.
fixed database role	A predefined role that exists in each database. The scope of the role is limited to the database in which it is defined.
fixed server role	A predefined role that exists at the server level. The scope of the role is limited to the SQL Server instance in which it is defined.
FK	A key in a database table that comes from another table (also know as the "referenced table") and whose values match the primary key (PK) or unique key in the referenced table.
flat file	A file consisting of records of a single record type in which there is no embedded structure information that governs relationships between records.
flatten	To convert a nested structure into a flat structure.
flattened interface	An interface created to combine members of multiple interfaces.
flattened rowset	A multidimensional data set presented as a two-dimensional rowset in which unique combinations of elements of multiple dimensions are combined on an axis.
flexible ID	An identifier that is assigned to various synchronization entities, such as replicas. The identifier can be of fixed or variable length.
flexible identifier	An identifier that is assigned to various synchronization entities, such as replicas. The identifier can be of fixed or variable length.
fold count	A value that represents the number of partitions that will be created within the original data set.
folder hierarchy	A bounded namespace that uniquely identifies all reports, folders, shared data source items, and resources that are stored in and managed by a report server.
forced service	In a database mirroring session, a failover initiated by the database owner upon the failure of the principal server that transfers service to the mirror database while it is in an unknown state. Data may be lost.
foreign key	A key in a database table that comes from another table (also know as the "referenced table") and whose values match the primary key (PK) or unique key in the referenced table.



Term	Definition
foreign key association	An association between entities that is managed through foreign key properties.
foreign table	A table that contains a foreign key.
format file	A file containing meta information (such as data type and column size) that is used to interpret data when being read from or written to a data file.
forward-only cursor	A cursor that cannot be scrolled; rows can be read only in sequence from the first row to the last row.
free-form language	A language whose syntax is not constrained by the position of characters on a line. C and Pascal are free-form languages; FORTRAN is not.
full backup	A backup of an entire database.
full differential backup	A backup of all files in the database, containing only changes made to the database since the most recent full backup. A full differential backup requires a full backup as a base.
full outer join	A type of outer join in which all rows in all joined tables are included, whether they are matched or not. For example, a full outer join between titles and publishers shows all titles and all publishers, even those that have no match.
full recovery model	A database recovery mode that fully logs all transactions and retains all the log records until after they are backed up. The database can be recovered to the point of failure if the tail of the log is backed up after the failure. All forms of recovery are supported.
full-text catalog	A collection of full-text index components and other files that are organized in a specific directory structure and contain the data that is needed to perform queries.
full-text enabling	The process of allowing full-text querying to occur on the current database.
full-text query	As a SELECT statement, a query that searches for words, phrases, or multiple forms of a word or phrase in the character-based columns (of char, varchar, text, ntext, nchar, or nvarchar data types). The SELECT statement returns those rows meeting the search criteria.
full-text search	A search for one or more documents, records, or strings based on all of the actual text data rather than on an index containing a limited set of keywords.
Full-Text service	The SQL Server component that performs the full-text querying.
full-width character	In a double-byte character set, a character that is represented by 2 bytes and typically has a half-width variant.
function	A piece of code that operates as a single logical unit. A function is called by name, accepts optional input parameters, and returns a status and optional output parameters. Many programming languages support functions.

Term	Definition
fuzzy grouping	In Integration Services, a data cleaning methodology that examines values in a dataset and identifies groups of related data rows and the one data row that is the canonical representation of the group.
fuzzy matching	In Integration Services, a lookup methodology that uses an approximate matching algorithm to locate similar data values in a reference table.
GAC	A computer-wide code cache that stores assemblies specifically installed to be shared by many applications on the computer.
gap depth	A measure that specifies the distance between data series that are displayed along distinct rows, as a result of clustering.
garbage collection	A process for automatic recovery of heap memory. Blocks of memory that had been allocated but are no longer in use are freed, and blocks of memory still in use may be moved to consolidate the free memory into larger blocks.
garbage collector	The part of the operating system that performs garbage collection.
gated link	A protected link between two or more objects. During execution, permissions are not checked across the object relationship once it has been established and credentials don't have to be checked several times. This type of link is useful when it is not appropriate or manageable to give permissions to many dependent objects.
gather-write operation	A performance optimization where the Database Engine collects multiple modified data pages into a single write operation.
GC	A process for automatic recovery of heap memory. Blocks of memory that had been allocated but are no longer in use are freed, and blocks of memory still in use may be moved to consolidate the free memory into larger blocks.
generated code	Code that is automatically generated for the client project based on operations and entities exposed in the middle tier when a RIA Services link exists between the server and client projects.
generator	The component of Notification Services that matches events to subscriptions and produces notifications.
geographic data	A type of spatial data that stores ellipsoidal (round-earth) data, such as GPS latitude and longitude coordinates.
geometric data	A type of spatial data that supports planar, or Euclidean (flat-earth), data.
ghost	An item or change unit in a filtered replica that was in the filter and has moved out.
ghost record	Row in the leaf level of an index that has been marked for deletion, but has not yet been deleted by the database engine.
ghost row	Row in the leaf level of an index that has been marked for deletion, but has not yet been deleted by the database engine.

Term	Definition
global assembly cache	A computer-wide code cache that stores assemblies specifically installed to be shared by many applications on the computer.
global default	A default that is defined for a specific database and is shared by columns of different tables.
global ID	A unique identifier that is assigned to a data item. The identifier must be unique across all clients. A global identifier is a flexible identifier and so can be any format, but it is typically a GUID and an 8-byte prefix.
global identifier	A unique identifier that is assigned to a data item. The identifier must be unique across all clients. A global identifier is a flexible identifier and so can be any format, but it is typically a GUID and an 8-byte prefix.
global rule	A rule that is defined for a specific database and is shared by columns of different tables.
global subscription	A subscription to a merge publication with an assigned priority value used for conflict detection and resolution.
globalization	The process of designing and developing a software product to function in multiple locales. Globalization involves identifying the locales that must be supported, designing features that support those locales, and writing code that functions equally well in any of the supported locales.
granularity	A description, from "coarse" to "fine", of a computer activity or feature (such as screen resolution, searching and sorting, or time slice allocation) in terms of the size of the units it handles (pixels, sets of data, or time slices). The larger the pieces, the coarser the granularity.
granularity attribute	The single attribute is used to specify the level of granularity for a given dimension in relation to a given measure group.
graphical query designer	A query designer provided by the Reporting Services that allows the user to interactively build a query and view the results for data source types SQL Server, Oracle, OLE DB, and ODBC.
graphics primitive	A basic shape (a point, a line, circle, curve, or polygon) that a graphics adapter can manipulate as a discrete entity.
group	A collection of users, computers, contacts, and other groups that is used as security or as e-mail distribution collections. Distribution groups are used only for e-mail. Security groups are used both to grant access to resources and as e-mail distribution lists.
grouping	A set of data that is grouped together in a report.
half-width character	In a double-byte character set, a character that is represented by one byte and typically has a full-width variant.
hard disk	An inflexible platter coated with material in which data can be recorded magnetically with read/write heads.
hard page-break renderer	A rendering extension that maintains the report layout and formatting so that the resulting file is optimized for a consistent printing experience, or to view the report online in a book format.

Term	Definition
hard-coding	The process of putting string or character literals in the main body of code, instead of in external resource files.
hard-coding	Basing numeric constants on the assumed length of a string; assumptions about language or culture-specific matters fixed in the code - e.g., string length, date formats, etc.
hardware security module	A secure device that provides cryptographic capabilities, typically by providing private keys used in Public-key cryptography.
hardware token	A secure device that provides cryptographic capabilities, typically by providing private keys used in Public-key cryptography.
hash partitioning	A way of partitioning a table or index by allowing SQL Server to apply an internal hash algorithm to spread rows across partitions based on the number of partitions specified and the values of one or more partitioning columns.
heat map	A type of map presentation where the intensity of color for each polygon corresponds to the related analytical data. For example, low values in a range appear as blue (cold) and high values as red (hot).
helpdesk	An individual or team of support professionals that provide technical assistance for an organization's network, hardware devices, and software.
heterogeneous data	Data stored in multiple formats.
hierarchy tree	A structure in which elements are related to each other hierarchically.
high availability	A Windows Server AppFabric feature that supports continuous availability of cached data by storing copies of that data on multiple cache hosts.
high availability	The ability of a system or device to be usable when it is needed. When expressed as a percentage, high availability is the actual service time divided by the required service time. Although high availability does not guarantee that a system will have no downtime, a network often is considered highly available if it achieves 99.999 percent network uptime.
high watermark	A memory consumption threshold on each cache host that specifies when objects are evicted out of memory, regardless of whether they have expired or not, until memory consumption goes back down to the low watermark.
high whisker	The highest value that is not an outlier on a box plot chart.
hint	An option or strategy specified for enforcement by the SQL Server query processor on SELECT, INSERT, UPDATE, or DELETE statements. The hint overrides any execution plan the query optimizer might select for a query.
history	A list of the user's actions within a program, such as commands entered in an operating system shell, menus passed through using Gopher, or links followed using a Web browser.

Term	Definition
holdability	Refers to the possibility of leaving result sets open ("on hold") that have been processed and are normally closed after this. For instance: "SQL Server supports holdability at the connection level only."
holdout	A percentage of training data that is reserved for use in measuring the accuracy of the structure of the data mining model.
holdout data	A percentage of training data that is reserved for use in measuring the accuracy of the structure of the data mining model.
holdout store	The data mining structure that is used to cache the holdout data. It contains references to the holdout data.
Home	Root folder in report server folder namespace.
home page	A document that serves as a starting point in a hypertext system. On the World Wide Web, an entry page for a set of Web pages and other files in a Web site. The home page is displayed by default when a visitor navigates to the site using a Web browser.
homogeneous data	Data that comes from multiple data sources that are all managed by the same software.
hop	In data communications, one segment of the path between routers on a geographically dispersed network.
hopping window	A type of window in which consecutive windows "hop" forward in time by a fixed period. The window is defined by two time spans: the period P and the window length L. For every P time unit a new window of size L is created.
horizontal partitioning	To segment a single table into multiple tables based on selected rows.
horizontal split	A horizontal orientation of the CIDER shell.
hot standby	A standby server that can support rapid failover without a loss of data from committed transactions.
hot standby server	A standby server that can support rapid failover without a loss of data from committed transactions.
HSM	A secure device that provides cryptographic capabilities, typically by providing private keys used in Public-key cryptography.
HTML	An application of the Standard Generalized Markup Language that uses tags to mark elements, such as text and graphics, in a document to indicate how Web browsers should display these elements to the user and should respond to user actions.
HTML Viewer	UI element consisting of a report toolbar and other navigation elements used to work with a report.
hybrid OLAP	A storage mode that uses a combination of multidimensional data structures and relational database tables to store multidimensional data.

Term	Definition
Hypertext Markup Language	An application of the Standard Generalized Markup Language that uses tags to mark elements, such as text and graphics, in a document to indicate how Web browsers should display these elements to the user and should respond to user actions.
identifying field	A field or group of fields that identify an entity as a unique object.
identifying relationship	A relationship where the primary key of the principal entity is part of the primary key of the dependent entity. In this kind of relationship, the dependent entity cannot exist without the principal entity.
identity column	A column in a table that has been assigned the identity property.
identity property	A property that generates values that uniquely identify each row in a table.
ideograph	A character in an Asian writing system that represents a concept or an idea, but not a particular word or pronunciation.
ideographic character	A character in an Asian writing system that represents a concept or an idea, but not a particular word or pronunciation.
IEC	One of two international standards bodies responsible for developing international data communications standards. The International Electrotechnical Commission (IEC) works closely with the International Organization for Standardization (ISO) to define standards of computing. They jointly published the ISO/IEC SQL-92 standard for SQL.
immediate updating subscription	A subscription to a transactional publication for which the user is able to make data modifications at the Subscriber. The data modifications are then immediately propagated to the Publisher using the two-phase commit protocol (2PC).
implicit cursor conversion	The return of a different type of cursor than the user had declared.
implicit transaction	A connection option in which each SQL statement executed by the connection is considered a separate transaction.
implied permission	Permission to perform an activity specific to a role.
inactive data source	A data source that has been backed up on the DPM server but is no longer being actively protected.
included column index	A nonclustered index containing both key and nonkey columns.

Term	Definition
incoming message	A message that has been sent across one or more messaging systems. It may have been sent only to you or to many other recipients. Incoming messages are placed in a receive folder designated to hold messages of a particular class. You can set up a different receive folder for each message class that you handle or use one folder for all of the classes.
incremental update	The set of operations that either adds new members to an existing cube or dimension, or adds new data to a partition.
independent association	An association between entities that is represented and tracked by an independent object.
Index Allocation Map	A page that maps the extents in a 4-GB part of a database file that is used by an allocation unit.
index page	A database page containing index rows.
indexed view	A view with a unique clustered index applied on it to improve the performance of some types of queries.
inference attack	A type of security threat in which a malicious user is able to derive the value of data without actually accessing it, for example by monitoring the response time to a query.
information model	An object-oriented schema that defines metadata constructs used to specify the structure and behavior of an application, process, component, or software artifact.
information technology	The formal name for a company's data processing department.
initial snapshot	Files that include schema and data, constraints, extended properties, indexes, triggers, and system tables that are necessary for replication.
initial synchronization	The first synchronization for a subscription, during which system tables and other objects that are required by replication, and the schema and data for each article, are copied to the Subscriber.
initial tape	In a media set using tape backup devices, the first tape in a media family.
inner join	An operation that retrieves rows from multiple source tables by comparing the values from columns shared between the source tables. An inner join excludes rows from a source table that have no matching rows in the other source tables.
in-person authentication	Physical authentication to complete a certificate request transaction. For example, an end user requesting his/her personal identification number (PIN) be unblocked will visit a certificate manager in person to provide in-person authentication with identification, such as an employee badge or drivers license.
InProc	A circumstance where the COM object's code is loaded from a DLL file and is located in the same process as the client.
input adapter	An adapter that accepts incoming event streams from external sources such as databases, files, ticker feeds, network ports, manufacturing devices and so on.

Term	Definition
input member	A member whose value is loaded directly from the data source instead of being calculated from other data.
input set	The set of data provided to a Multidimensional Expressions (MDX) value expression upon which the expression operates.
input source	Any table, view, or schema diagram used as an information source for a query.
input stream	A flow of information used in a program as a sequence of bytes that are associated with a particular task or destination. Input streams include series of characters read from the keyboard to memory and blocks of data read from disk files.
insensitive cursor	A cursor that does not reflect data modification made to the underlying data by other users while the cursor is open.
insert event	The event kind used to signify the arrival of an event into the stream. The insert event type consists of metadata that defines the valid lifetime of the event and the payload (data) fields of the event.
Insert Into query	A query that copies specific columns and rows from one table to another or to the same table.
Insert Values query	A query (SQL statement) that creates a new row and inserts values into specified columns.
instance	A copy of SQL Server running on a computer.
instance control provider	A provider that allows you to issue control commands against workflow instances in an instance store. For example, a SQL control provider lets you suspend, resume, or terminate instances stored in a SQL Server database. When you execute a cmdlet that controls a workflow instance in an instance store, the cmdlet internally uses the control provider for that instance store to send commands to the instance.
instance query provider	A provider that allows you to issue queries against an instance store. For example, a SQL query provider lets you query for workflow instances stored in a SQL Server database. When you execute a cmdlet that queries for instances against an instance store, the cmdlet internally uses the query provider to retrieve instances from that store.
instance store	A set of database tables that store workflow instance state and workflow instance metadata.
instance store provider	In Windows AppFabric, a provider that allows you to create instance store objects. For example, a SQL store provider allows clients to create SQL workflow instance store objects, which in turn allows clients to save and retrieve workflow instances to and from a persistence store.
integration	In computing, the combining of different activities, programs, or hardware components into a functional unit.
integrity	The accuracy of data and its conformity to its expected value, especially after being transmitted or processed.
integrity constraint	A property defined on a table that prevents data modifications that would create invalid data.



Term	Definition
intent lock	A lock that is placed on one level of a resource hierarchy to protect shared or exclusive locks on lower-level resources.
intent share	A lock that is placed on one level of a resource hierarchy to protect shared or exclusive locks on lower-level resources.
interactive structured query language	An interactive command prompt utility provided with SQL Server that lets users run Transact-SQL statements or batches from a server or workstation and view the results that are returned.
interface	A defined set of properties, methods, and collections that form a logical grouping of behaviors and data.
interface implication	If an interface implies another interface, then any class that implements the first interface must also implement the second interface. Interface implication is used in an information model to get some of the effects of multiple inheritance.
intermediate language	A computer language used as an intermediate step between the original source language, usually a high-level language, and the target language, usually machine code. Some high-level compilers use assembly language as an intermediate language.
International Electrotechnical Commission	One of two international standards bodies responsible for developing international data communications standards. The International Electrotechnical Commission (IEC) works closely with the International Organization for Standardization (ISO) to define standards of computing. They jointly published the ISO/IEC SQL-92 standard for SQL.
Internet Protocol security	A set of industry-standard, cryptography-based services and protocols that help to protect data over a network.
interprocess communication	The ability of one task or process to communicate with another in a multitasking operating system. Common methods include pipes, semaphores, shared memory, queues, signals, and mailboxes.
interval event	An event whose payload is valid for a given period of time. The metadata of the interval event requires that both the start and end time of the interval be provided in the event metadata. Interval events are valid only for this specific interval.
interval event model	The event model of an interval event.
invoke operation	A domain operation that is executed without tracking or deferred execution.
IP address	A binary number that uniquely identifies a host (computer) connected to the Internet to other Internet hosts, for the purposes of communication through the transfer of packets.
IPC	The ability of one task or process to communicate with another in a multitasking operating system. Common methods include pipes, semaphores, shared memory, queues, signals, and mailboxes.

Term	Definition
IPsec	A set of industry-standard, cryptography-based services and protocols that help to protect data over a network.
isolation level	The property of a transaction that controls the degree to which data is isolated for use by one process, and is guarded against interference from other processes.
ISQL	An interactive command prompt utility provided with SQL Server that lets users run Transact-SQL statements or batches from a server or workstation and view the results that are returned.
item	A unit of data or metadata that is being synchronized. A typical item of data might be a file or record, whereas a typical item of metadata might be a knowledge item.
item-level role assignment	A security policy that applies to an item in the report server folder namespace.
item-level role definition	A security template that defines a role used to control access to or interaction with an item in the report server folder namespace.
iterate	To execute one or more statements or instructions repeatedly. Statements or instructions so executed are said to be in a loop.
job	A specified series of operations, called steps, performed sequentially by a program to complete an action.
job history	Log that keeps a historical record of jobs.
join	To combine the contents of two or more tables and produce a result set that incorporates rows and columns from each table. Tables are typically joined using data that they have in common.
join column	A column referenced in a join condition.
join condition	A comparison clause that specifies how tables are related by their join columns.
join filter	A filter used in merge replication that extends the row filter of one table to a related table.
join operator	A comparison operator in a join condition that determines how the two sides of the condition are evaluated and which rows are returned.
join path	A series of joins indicating how two tables are related.
junction table	A table that establishes a relationship between other tables.
Kagi chart	A chart, mostly independent of time, used to track price movements and to make decisions on purchasing stock.
key	A string that identifies an object in the cache. This string must be unique within a region. Objects are associated with a key when they are added and then retrieved with the same key.

<b>Term</b>	<b>Definition</b>
key	In encryption, authentication, and digital signatures, a value used in combination with an algorithm to encrypt or decrypt information.
key	In an array, the field by which stored data is organized and accessed.
key	A column or group of columns that uniquely identifies a row (primary key), defines the relationship between two tables (foreign key), or is used to build an index.
key attribute	The attribute of a dimension that links the non-key attributes in the dimension to related measures.
key column	A column whose contents uniquely identify every row in a table.
key generator	A hardware or software component that is used to generate encryption key material.
key performance indicator	A predefined measure that is used to track performance of a strategic goal, objective, plan, initiative, or business process. A KPI is evaluated against a target. An explicit and measurable value taken directly from a data source. Key performance indicators (KPIs) are used to measure performance in a specific area, for example, revenue per customer.
key range lock	A lock that is used to lock ranges between records in a table to prevent phantom additions to, or deletions from, a set of records. Ensures serializable transactions.
key recovery	The process of recovering a user's private key.
Key Recovery Agent	A designated user that works with a certificate administrator to recover a user's private key. A specific certificate template is applied to a Key Recovery Agent.
keyset-driven cursor	A cursor that shows the effects of updates made to its member rows by other users while the cursor is open, but does not show the effects of inserts or deletes.
knowledge	The metadata about all the changes that a participant has seen and maintains.
KPI	A predefined measure that is used to track performance of a strategic goal, objective, plan, initiative, or business process. A KPI is evaluated against a target. An explicit and measurable value taken directly from a data source. Key performance indicators (KPIs) are used to measure performance in a specific area, for example, revenue per customer.
KRA	A designated user that works with a certificate administrator to recover a user's private key. A specific certificate template is applied to a Key Recovery Agent.
Language for non-Unicode programs	A Regional and Language Options setting that specifies the default code pages and associated bitmap font files for a specific computer that affects all of that computer's users. The default code pages and fonts enable a non-Unicode application written for one operating system language version to run correctly on another operating system language version.
language service parser	A component that is used to describe the functions and scope of the tokens in source code.

Term	Definition
language service scanner	A component that is used to identify types of tokens in source code. This information is used for syntax highlighting and for quickly identifying token types that can trigger other operations, for example, brace matching.
latch	A short-term synchronization object protecting actions that need not be locked for the life of a transaction. A latch is primarily used to protect a row that the storage engine is actively transferring from a base table or index to the relational engine.
latency	The delay that occurs while data is processed or delivered.
lazy loading	A pattern of data loading where related objects are not loaded until a navigation property is accessed.
lazy schema validation	An option that delays checking the remote schema to validate its metadata against a query until execution in order to increase performance.
lead byte	The byte value that is the first half of a double-byte character.
lead host	A cache host that has been designated to work with other lead hosts and to keep the cluster running at all times.
leaf	A node with no child objects represented in the tree.
leaf level	The bottom level of a clustered or nonclustered index, or the bottom level of a hierarchy.
leaf member	A member that has no descendents.
leaf node	A node with no child objects represented in the tree.
learned knowledge	The current knowledge of a source replica about a specific set of changes, and the logged conflicts of that replica.
least recently used	The type of eviction used by the cache cluster, where least recently used objects are evicted before the most recently used objects.
left outer join	A type of outer join in which all rows from the left-most table in the JOIN clause are included. When rows in the left table are not matched by rows in the right table, all result set columns that come from the right table are assigned a value of NULL.
level	The name of a set of members in a dimension hierarchy such that all members of the set are at the same distance from the root of the hierarchy. For example, a time hierarchy may contain the levels Year, Month, and Day.
lift chart	In Analysis Services, a chart that compares the accuracy of the predictions of each data mining model in the comparison set.
lightweight pooling	An option that provides a means of reducing the system overhead associated with the excessive context switching sometimes seen in symmetric multiprocessing (SMP) environments by performing the context switching inline, thus helping to reduce user/kernel ring transitions.
line layer	The layer in a map report that displays spatial data as lines, for example, lines that indicate paths or routes.

Term	Definition
linked dimension	A reference in a cube to a dimension in a different cube (that is, a cube with a different data source view that exists in the same or a different Analysis Services database). A linked dimension can only be related to measure groups in the source cube, and can only be edited in the source database.
linked measure group	A reference in a cube to a measure group in a different cube (that is, a cube with a different data source view that exists in the same or a different Analysis Services database). A linked measure group can only be edited in the source database.
linked server	A definition of an OLE DB data source used by SQL Server distributed queries. The linked server definition specifies the OLE DB provider required to access the data, and includes enough addressing information for the OLE DB provider to connect to the data.
linked table	An OLE DB rowset exposed by an OLE DB data source that has been defined as a linked server for use in SQL Server distributed queries.
linking table	A table that has associations with two other tables, and is used indirectly as an association between those two tables.
LINQ	A query syntax that defines a set of query operators that allow traversal, filter, and projection operations to be expressed in a direct, declarative way in any .NET-based programming language.
little endian	Pertaining to a processor memory architecture that stores numbers so that the least significant byte is placed first.
local cache	A feature that enables deserialized copies of cached objects to be saved in the memory of the same process that runs the cache-enabled application.
local Distributor	A server that is configured as both a Publisher and a Distributor for SQL Server Replication.
local partitioned view	A view that joins horizontally partitioned data from a set of member tables across a single server, making the data appear as if from one table.
local subscription	A subscription to a merge publication that uses the priority value of the Publisher for conflict detection and resolution.
locale	A collection of rules and data specific to a language and a geographic area. Locales include information on sorting rules, date and time formatting, numeric and monetary conventions, and character classification.
localization	The process of adapting a product and/or content (including text and non-text elements) to meet the language, cultural, and political expectations and/or requirements of a specific local market (locale).
lock	A restriction on access to a resource in a multiuser environment.
lock escalation	The process of converting many fine-grain locks into fewer coarse-grain locks, thereby reducing system overhead.

Term	Definition
log backup	A backup of transaction logs that includes all log records not backed up in previous log backups. Log backups are required under the full and bulk-logged recovery models and are unavailable under the simple recovery model.
log chain	A continuous sequence of transaction logs for a database. A new log chain begins with the first backup taken after the database is created, or when the database is switched from the simple to the full or bulk-logged recovery model. A log chain forks after a restore followed by a recovery, creating a new recovery branch.
log provider	A provider that logs package information at run time. Integration Services includes a variety of log providers that make it possible to capture events during package execution. Logs are created and stored in formats such as XML, text, database, or in the Windows event log.
Log Reader Agent	In Replication, the executable that monitors the transaction of each database configured for transactional replication, and copies the transactions marked for replication from the transaction into the distribution database.
log sequence number	A unique number assigned to each entry in a transaction log. LSNs are assigned sequentially according to the order in which entries are created.
log shipping	Copying, at regular intervals, log backup from a read-write database (the primary database) to one or more remote server instances (secondary servers). Each secondary server has a read-only database, called a secondary database, that was created by restoring a full backup of the primary database without recovery. The secondary server restores each copied log backup to the secondary database. The secondary servers are warm standbys for the primary server.
log shipping configuration	A single primary server, one or more secondary servers (each with a secondary database), and a monitor server.
log shipping job	A job performing one of the following log-shipping operations: backing up the transaction log of the primary database at the primary server (the backup job), copying the transaction log file to a secondary server (the copy job), or restoring the log backup to the secondary database on a secondary server (the restore job). The backup job resides on the primary server; the copy and restore jobs reside on each of the secondary servers. See also: primary database, primary server, secondary database, secondary server.
log sink	A tracing function of the cache client and cache host. Log sinks capture trace events from the cache client or cache host and can display them in a console, write them to a log file, or report them to the Event Tracing for Windows (ETW) framework inside Windows.
logic error	An error, such as a faulty algorithm, that causes a program to produce incorrect results but does not prevent the program from running. Consequently, a logic error is often very difficult to find.
logical name	A name used by SQL Server to identify a file.
logical record	A merge replication feature that allows you to define a relationship between related rows in different tables so that the rows are processed as a unit.
login ID	A string that is used to identify a user or entity to an operating system, directory service, or distributed system. For example, in Windows® integrated authentication, a login name uses the form "DOMAIN\username."

Term	Definition
login security mode	A security mode that determines the manner in which an instance of SQL Server validates a login request.
long parsing	In the SQL system there are two typical types of threads: - short thread: it is a process that use the resources for a short time and - long thread: it is a process that use the resources for a long time. long parsing: is the analysis of the threads that lived for a long time Note: the definition of short/long is based on the system calculation/statistic for each process.
Low Box	The lowest value of a box on a Box Plot chart.
low watermark	A memory consumption threshold on each cache host that specifies when expired objects are evicted out of memory.
low whisker	The lowest value that is not an outlier on a box plot chart.
LRU	The type of eviction used by the cache cluster, where least recently used objects are evicted before the most recently used objects.
LSN	A unique number assigned to each entry in a transaction log. LSNs are assigned sequentially according to the order in which entries are created.
luring attack	An attack in which the client is lured to voluntarily connect to the attacker.
made-with knowledge	In synchronization processes, the current knowledge of the source replica, to be used in conflict detection.
Manage Relationships	A UI element that enables a user to view, delete or create new relationships in a model.
managed code	Code that is executed by the common language runtime environment rather than directly by the operating system. Managed code applications gain common language runtime services such as automatic garbage collection, runtime type checking and security support, and so on. These services help provide uniform platform- and language-independent behavior of managed-code applications.
managed instance	An instance of SQL Server monitored by a utility control point.
management data warehouse	A relational database that is used to store data that is collected.
management policy	A definition of the workflows used for managing certificates within a Certificate Lifecycle Manager (CLM) profile template. A management policy defines who performs specific management tasks within the workflows, and provides management details for the entire lifecycle of the certificates within the profile template.
Management Studio	A suite of management tools included with Microsoft SQL Server for configuring, managing, and administering all components within Microsoft SQL Server.

Term	Definition
manual failover	In a database mirroring session, a failover initiated by the database owner, while the principal server is still running, that transfers service from the principal database to the mirror database while they are in a synchronized state.
many-to-many dimension	A relationship between a dimension and a measure group in which a single fact may be associated with many dimension members and a single dimension member may be associated with a many facts. To define this relationship between the dimension and the fact table, the dimension is joined to an intermediate fact table and the intermediate fact table is joined, in turn, to an intermediate dimension table that is joined to the fact table.
many-to-one relationship	A relationship between two tables in which one row in one table can relate to many rows in another table.
map	To associate data with a specified location in memory.
map control	A JavaScript control that contains the objects, methods, and events that you need to display maps powered by Bing Maps™ on your Web site.
map gallery	A gallery that contains maps from reports that are located in the map gallery folder for the report authoring environment.
map layer	A child element of the map, each map layer including elements for their map members and map member attributes.
map resolution	The accuracy at which the location and shape of map features can be depicted for a given map scale. In a large scale map (e.g. a map scale of 1:1) there is less reduction of features than those shown on a small scale map (e.g. 1:1,000,000).
map tile	One of a number of 256 x 256 pixel images that are combined to create a Bing map. A map tile contains a segment of a view of the earth in Mercator projection, with possible road and text overlays depending on the style of the Bing map.
map viewport	The area of the map to display in the map report item. For example, a map for the entire United States might be embedded in a report, but only the area for the northwestern states are displayed.
MAPI	A messaging architecture that enables multiple applications to interact with multiple messaging systems across a variety of hardware platforms. MAPI is built on the Component Object Model (COM) foundation.
mapper	A component that maps objects.
marker	A visual indicator that identifies a data point. In a map report, a marker is the visual indicator that identifies the location of each point on the point layer.
marker map	A map that displays a marker at each location (for example, cities) and varies marker color, size, and type.
market basket analysis	A standard data mining algorithm that analyzes a list of transactions to make predictions about which items are most frequently purchased together.



Term	Definition
master data	The critical data of a business, such as customer, product, location, employee, and asset. Master data fall generally into four groupings: people, things, places, and concepts and can be further categorized. For example, within people, there are customer, employee, and salesperson. Within things, there are product, part, store, and asset. Within concepts, there are things like contract, warrantee, and licenses. Finally, within places, there are office locations and geographic divisions.
master data management	The technology, tools, and processes required to create and maintain consistent and accurate lists of master data of an organization.
Master Data Manager	A component of the Master Data Services application for managing and accessing master data.
Master Data Services	A master data management application to consistently define and manage the critical data entities of an organization.
Master Data Services Configuration Manager	A SQL Server configuration manager used to create and configure Master Data Services databases, Web sites, and Web applications.
master database	The system database that records all the system-level information for an instance of SQL Server.
master file	The file installed with earlier versions of SQL Server used to store the master, model, and tempdb system databases and transaction logs and the pubs sample database and transaction log.
master merge	The process of combining shadow indexes with the current master index to form a new master index.
master server	A server that distributes jobs and receives events from multiple servers.
materialized view	A view in which the query result is cached as a concrete table that may be updated from the original base tables from time to time.
matrix data region	A report item on a report layout that displays data in a variable columnar format.
MBCS	A mixed-width character set, in which some characters consist of more than 1 byte. An MBCS is used in languages such as Japanese, Chinese, and Korean, where the 256 possible values of a single-byte character set are not sufficient to represent all possible characters.
MDS	A master data management application to consistently define and manage the critical data entities of an organization.
MDX	A language for querying and manipulating data in multidimensional objects (OLAP cubes).
measure	In a cube, a set of values that are usually numeric and are based on a column in the fact table of the cube. Measures are the central values that are aggregated and analyzed.
measure group	A collection of related measures in an Analysis Services cube. The measures are generally from the same fact table.

Term	Definition
media family	Data written by a backup operation to a backup device used by a media set. In a media set with only a single device, only one media family exists. In a striped media set, multiple media families exist. If the striped media set is unmirrored, each device corresponds to a family. A mirrored media set contains from two to four identical copies of each media family (called mirrors). Appending backups to a media set extends its media families.
media header	A label that provides information about the backup media.
media set	An ordered collection of backup media written to by one or more backup operations using a constant number of backup devices.
median	The middle value in a set of ordered numbers. The median value is determined by choosing the smallest value such that at least half of the values in the set are no greater than the chosen value. If the number of values within the set is odd, the median value corresponds to a single value. If the number of values within the set is even, the median value corresponds to the sum of the two middle values divided by two.
median price formula	A formula that calculates the average of the high and low prices.
median value	The middle value in a set of ordered numbers. The median value is determined by choosing the smallest value such that at least half of the values in the set are no greater than the chosen value. If the number of values within the set is odd, the median value corresponds to a single value. If the number of values within the set is even, the median value corresponds to the sum of the two middle values divided by two.
member	A single position or item in a dimension. Dimension members can be user-defined or predefined and can have properties associated with them.
member delegation	A modeling concept that describes how interface members are mapped from one interface to another.
member expression	A valid Multidimensional Expressions (MDX) expression that returns a member.
member property	A characteristic of a dimension member. Dimension member properties can be alphanumeric, Boolean, or Date/Time data types, and can be user-defined or predefined.
memo	A type of column that contains long strings, typically more than 255 characters.
memory broker	A software component that manages the distribution of memory resources in SQL Server.
memory clerk	A memory management component that allocates memory.
merge replication	A type of replication that allows sites to make autonomous changes to replicated data, and at a later time, merge changes and resolve conflicts when necessary.
merge tombstone	A marker created when a constraint conflict is resolved by merging the two items in conflict.

Term	Definition
message number	A number that identifies a SQL Server error message.
Message Queuing	A Microsoft technology that enables applications running at different times to communicate across heterogeneous networks and systems that may be temporarily offline.
message type	A definition of a Service Broker message. The message type specifies the name of the message and the type of validation Service Broker performs on incoming messages of that type.
Messages pane	One of the tabs that hosts the messages returned from SQL Server after a TSQL query has been executed.
Messaging Application Programming Interface	A messaging architecture that enables multiple applications to interact with multiple messaging systems across a variety of hardware platforms. MAPI is built on the Component Object Model (COM) foundation.
metadata	Information about the properties or structure of data that is not part of the values the data contains.
method	In object-oriented programming, a named code block that performs a task when called.
Microsoft Message Queuing	A Microsoft technology that enables applications running at different times to communicate across heterogeneous networks and systems that may be temporarily offline.
Microsoft Sequence Clustering algorithm	Algorithm that is a combination of sequence analysis and clustering, which identifies clusters of similarly ordered events in a sequence. The clusters can be used to predict the likely ordering of events in a sequence based on known characteristics.
Microsoft SharePoint Foundation Usage Data Import	The default timer job for SharePoint Foundation that imports usage log files into the logging database.
Microsoft SharePoint Server	A family of Microsoft enterprise portal servers, built on Windows SharePoint Services, used to aggregate SharePoint sites, information, and applications into a single portal.
Microsoft SQL Server	A family of Microsoft relational database management and analysis systems for e-commerce, line-of-business, and data warehousing solutions.
Microsoft SQL Server 2008 Express	A lightweight and embeddable version of Microsoft SQL Server 2008.

Term	Definition
Microsoft SQL Server 2008 Express with Advanced Services	A Microsoft relational database design and management system for e-commerce, line-of-business, and data warehousing solutions.
Microsoft SQL Server 2008 Express with Tools	A free, easy-to-use version of the SQL Server Express data platform that includes the graphical management tool: SQL Server Management Studio (SMSS) Express.
Microsoft SQL Server Books Online	A collection of electronic documentation that includes the complete documentation that ships with Microsoft SQL Server.
Microsoft SQL Server Business Intelligence	The Microsoft SQL Server-based data infrastructure and business intelligence platform consisting of Microsoft SQL Server Integration Services, Relational Engine, Master Data Services, Reporting Services, and Analysis Services.
Microsoft SQL Server Compact	A Microsoft relational database management and analysis system for e-commerce, line-of-business, and data warehousing solutions.
Microsoft SQL Server Compact 3.5 for Devices	A file that installs the SQL Server Compact 3.5 devices runtime components.
Microsoft SQL Server Notification Services	A Microsoft SQL Server add-in that provides a development framework and hosting server for building and deploying notification applications.
Microsoft SQL Server Playback Capture Wizard	A wizard that supports the capture of information about an event for later analysis.
Microsoft SQL Server Power Pivot for Microsoft Excel	A SQL Server add-in for Excel.
Microsoft SQL Server Power Pivot for Microsoft SharePoint	A Microsoft SQL Server technology that provides query processing and management control for Power Pivot workbooks published to SharePoint.

Term	Definition
Microsoft SQL Server Reporting Services	A server-based report generation environment for enterprise, Web-enabled reporting functionality so you can create reports that draw content from a variety of data sources, publish reports in various formats, and centrally manage security and subscriptions.
Microsoft SQL Server Reporting Services Report Builder	A report authoring tool that features a Microsoft Office-like authoring environment and features such as new sparkline, data bar, and indicator data visualizations, the ability to save report items as report parts, a wizard for creating maps, aggregates of aggregates, and enhanced support for expressions.
Microsoft SQL Server Service Broker	A technology that helps developers build scalable, secure database applications.
Microsoft SQL Server System CLR Types	A stand-alone package, part of SQL Server 2008 R2 Feature Pack, that contains the components implementing the geometry, geography, and hierarchy id types in SQL Server.
Microsoft Time Series algorithm	Algorithm that uses a linear regression decision tree approach to analyze time-related data, such as monthly sales data or yearly profits. The patterns it discovers can be used to predict values for future time steps.
middleware	Software that sits between two or more types of software and translates information between them. Middleware can cover a broad spectrum of software and generally sits between an application and an operating system, a network operating system, or a database management system.
mining model	An object that contains the definition of a data mining process and the results of the training activity. For example, a data mining model may specify the input, output, algorithm, and other properties of the process and hold the information gathered during the training activity, such as a decision tree.
mining structure	A data mining object that defines the data domain from which the mining models are built.
minor tick mark	A tick mark that corresponds to a minor scaling unit on an axis.
mirror database	In a database mirroring session, the copy of the database that is normally fully synchronized with the principal database.
mirror server	In a database mirroring configuration, the server instance where the mirror database resides. The mirror server is the mirroring partner whose copy of the database is currently the mirror database. The mirror server is a hot standby server.
mirrored media set	A media set that contains two to four identical copies (mirrors) of each media family. Restore operations require only one mirror per family, allowing a damaged media volume to be replaced by the corresponding volume from a mirror.

Term	Definition
mirroring	Immediately reproducing every update to a read-write database (the principal database) onto a read-only mirror of that database (the mirror database) residing on a separate instance of the database engine (the mirror server). In production environments, the mirror server is on another machine. The mirror database is created by restoring a full backup of the principal database (without recovery).
model database	A database that is installed with Microsoft SQL Server and that provides the template for new user databases. SQL Server creates a database by copying in the contents of the model database and then expanding the new database to the size requested.
model dependency	A relationship between two or more models in which one model is dependent on the information of another model.
modulo	An arithmetic operation whose result is the remainder of a division operation. For example, 17 modulo 3 = 2 because 17 divided by 3 yields a remainder of 2. Modulo operations are used in programming.
monitor server	In a log shipping configuration, a server instance on which every log shipping job in the configuration records its history and status. Each log shipping configuration has its own dedicated monitor server.
msdb	A database that stores scheduled jobs, alerts, and backup/restore history information.
MSMQ	A Microsoft technology that enables applications running at different times to communicate across heterogeneous networks and systems that may be temporarily offline.
multibase differential	A differential backup that includes files that were last backed up in distinct base backups.
multibyte character set	A mixed-width character set, in which some characters consist of more than 1 byte. An MBCS is used in languages such as Japanese, Chinese, and Korean, where the 256 possible values of a single-byte character set are not sufficient to represent all possible characters.
multicast delivery	A method for delivering notifications that formats a notification once and sends the resulting message to multiple subscribers.
multidimensional expression	A language for querying and manipulating data in multidimensional objects (OLAP cubes).
multidimensional OLAP	A storage mode that uses a proprietary multidimensional structure to store a partition's facts and aggregations or a dimension.
multidimensional structure	A database paradigm that treats data as cubes that contain dimensions and measures in cells.
multiplier	In arithmetic, the number that indicates how many times another number (the multiplicand) is multiplied.
multiserver administration	The process of automating administration across multiple instances of SQL Server.

Term	Definition
multithreaded server application	An application that creates multiple threads within a single process to service multiple user requests at the same time.
multiuser	Pertaining to any computer system that can be used by more than one person. Although a microcomputer shared by several people can be considered a multiuser system, the term is generally reserved for machines that can be accessed simultaneously by several people through communications facilities or via network terminals.
mutex	A programming technique that ensures that only one program or routine at a time can access some resource, such as a memory location, an I/O port, or a file, often through the use of semaphores, which are flags used in programs to coordinate the activities of more than one program or routine.
My Reports	A personalized workspace.
My Subscriptions	A page that lists all subscriptions that a user owns.
named cache	A configurable unit of in-memory storage that has policies associated with it and that is available across all cache hosts in a cache cluster.
named instance	An installation of SQL Server that is given a name to differentiate it from other named instances and from the default instance on the same computer.
named pipe	A portion of memory that can be used by one process to pass information to another process, so that the output of one is the input of the other. The second process can be local (on the same computer as the first) or remote (on a networked computer).
named update	A custom service operation that performs an action which is different than a simple query, update, insert, or delete operation.
named update method	A custom service operation that performs an action which is different than a simple query, update, insert, or delete operation.
naming convention	Any standard used more or less universally in the naming of objects, etc.
national language support API	Set of system functions in 32-bit Windows containing information that is based on language and cultural conventions.
native compiler	A compiler that produces machine code for the computer on which it is running, as opposed to a cross-compiler, which produces code for another type of computer. Most compilers are native compilers.
native format	A data format that maintains the native data types of a database. Native format is recommended when you bulk transfer data between multiple instances of Microsoft SQL Server using a data file that does not contain any extended/double-byte character set (DBCS) characters.

Term	Definition
natural hierarchy	A hierarchy in which at every level there is a one-to-many relationship between members in that level and members in the next lower level.
needle cap	One of the two appearance properties that can be applied to a radial gauge.
nested query	A SELECT statement that contains one or more subqueries.
Net-Library	A SQL Server communications component that isolates the SQL Server client software and the Database Engine from the network APIs.
network software	Software that enables groups of computers to communicate, including a component that facilitates connection to or participation in a network.
new line character	A control character that causes the cursor on a display or the printing mechanism on a printer to move to the beginning of the next line.
nickname	When used with merge replication system tables, a name for another Subscriber that is known to already have a specified generation of updated data.
niladic functions	Functions that do not have any input parameters.
NLS API	Set of system functions in 32-bit Windows containing information that is based on language and cultural conventions.
node	A synchronization provider and its associated replica.
noise word	A word such as 'the' or 'an' that is not useful for searches, or that a crawler should ignore when creating an index.
nonclustered index	A B-tree-based index in which the logical order of the index key values is different than the physical order of the corresponding rows in a table. The index contains row locators that point to the storage location of the table data.
non-contained database	A SQL Server database that stores database settings and metadata with the instance of SQL Server Database Engine where the database is installed, and requires logins in the master database for authentication.
Nonkey index column	Column in a nonclustered index that does not participate as a key column. Rather, the column is stored in the leaf-level of the index and is used in conjunction with the key columns to cover one or more queries.
nonleaf member	A member with one or more descendants.
nonnullable parameter	A parameter which cannot take a NULL value.
nonrepeatable read	An inconsistency that occurs when a transaction reads the same row more than once, and a separate transaction modifies that row between reads.



<b>Term</b>	<b>Definition</b>
normalization rule	A design rule that minimizes data redundancy and results in a database in which the Database Engine and application software can easily enforce integrity.
notification	A message or announcement sent to the user or administrator of a system. The recipient may be a human or an automated notification manager.
Notification Services	A Microsoft SQL Server add-in that provides a development framework and hosting server for building and deploying notification applications.
NSControl	The command prompt utility for administering Notification Services instances and applications.
NUL	A "device," recognized by the operating system, that can be addressed like a physical output device (such as a printer) but that discards any information sent to it.
null	Pertaining to a value that indicates missing or unknown data.
null key	A null value that is encountered in a key column.
null pointer	A pointer to nothing; usually a standardized memory address, such as 0. A null pointer usually marks the last of a linear sequence of pointers or indicates that a data search operation has come up empty.
nullability	The attribute of a column, parameter, or variable that specifies whether it allows null data values.
nullable property	A property which controls if a field can have a NULL value.
NUMA node	A group of processors with its own memory and possibly its own I/O channels.
numeric array	An array composed of a collection of keys and a collection of values, where each key is associated with one value. The values can be of any type, but the keys must be numeric.
numeric expression	Any expression that evaluates to a number. The expression can be any combination of variables, constants, functions, and operators.
object identifier	A number that identifies an object class or attribute. Object identifiers (OIDs) are organized into an industry-wide global hierarchy. An object identifier is represented as a dotted decimal string, such as 1.2.3.4, with each dot representing a new branch in the hierarchy. National/regional registration authorities issue root object identifiers to individuals or organizations, who manage the hierarchy below their root object identifier.
object lifetime	The span of time a cached object resides in cache and is available to be retrieved by cache clients. The object expires when its lifetime ends. Expired objects cannot be retrieved by cache clients, but remain in memory of the cache host until they are evicted. Specified as time to live (TTL).
object variable	A variable that contains a reference to an object.
ODS library	A set of C functions that makes an application a server. ODS library calls respond to requests from a client in a client/server network. Also manages the communication and data between the client and the server. ODS library follows the tabular data stream (TDS) protocol.

Term	Definition
Office File Validation	A security feature that validates files before allowing them to be loaded by the application, in order to protect against file format vulnerabilities.
offline restore	A restore during which the database is offline.
OLAP	A technology that uses multidimensional structures to provide rapid access to data for analysis. The source data for OLAP is commonly stored in data warehouses in a relational database.
OLAP cube	A set of data that is organized and summarized into a multidimensional structure that is defined by a set of dimensions and measures.
OLAP database	A relational database system capable of handling queries more complex than those handled by standard relational databases, through multidimensional access to data (viewing the data by several different criteria), intensive calculation capability, and specialized indexing techniques.
one-time password	Passwords produced by special password generating software or by a hardware token and that can be used only once.
online analytical processing	A technology that uses multidimensional structures to provide rapid access to data for analysis. The source data for OLAP is commonly stored in data warehouses in a relational database.
online restore	A restore in which one or more secondary filegroups, files belonging to secondary filegroups, or pages are restored while the database remains online. Online restore is available only in the SQL Server 2005 Enterprise Edition (including the Evaluation and Developer Editions).
Open Data Services library	A set of C functions that makes an application a server. ODS library calls respond to requests from a client in a client/server network. Also manages the communication and data between the client and the server. ODS library follows the tabular data stream (TDS) protocol.
operation code	The portion of a machine language or assembly language instruction that specifies the type of instruction and the structure of the data on which it operates.
operator	An atomic unit of a query as scheduled by the StreamInsight server to process the events on which the query is applied. Examples include SELECT, PROJECT, AGGREGATE, UNION, TOP K and JOIN. Operators are fully composable and have a specific number of inputs and outputs. See Other Terms: query, query template, composable.
operator	A sign or symbol that specifies the type of calculation to perform within an expression. There are mathematical, comparison, logical, and reference operators.
optimistic concurrency	A method of managing concurrency by using a cached object's version information. Because every update to an object changes its version number, using version information prevents the update from overwriting someone else's changes.
optimize synchronization	An option in merge replication that allows you to minimize network traffic when determining whether recent changes have caused a row to move into or out of a partition that is published to a Subscriber.
ordered set	A set of members returned in a specific order.

<b>Term</b>	<b>Definition</b>
origin object	An object in a repository that is the origin in a directional relationship.
OTP	Passwords produced by special password generating software or by a hardware token and that can be used only once.
outer join	A join that includes all the rows from the joined tables that meet the search conditions, even rows from one table for which there is no matching row in the other join table.
output adapter	An adapter that receives events processed by the server, transforms the events into a format expected by the output device (a database, text file, PDA, or other device), and emits the data to that device.
output column	In SQL Server Integration Services, a column that the source adds to the data flow and that is available as input column to the next data flow component in the data flow.
output stream	A flow of information that leaves a computer system and is associated with a particular task or destination.
overfitting	A problem in data mining when random variations in data are misclassified as important patterns. Overfitting often occurs when the data set is too small to represent the real world.
ownership chain	When an object references other objects and the calling and the called objects are owned by the same user. SQL Server uses the ownership chain to determine how to check permissions.
package	A collection of control flow and data flow elements that runs as a unit.
packet	A unit of information transmitted from one computer or device to another on a network.
Pad index	An option that specifies the space to leave open on each page in the intermediate levels of the index.
padding	In data storage, the addition of one or more bits, usually zeros, to a block of data in order to fill it, to force the actual data bits into a certain position, or to prevent the data from duplicating a bit pattern that has an established meaning, such as an embedded command
page	To return the results of a query in smaller subsets of data, thus making it possible for the user to navigate through the result set by viewing 'pages' of data.
page	In a virtual storage system, a fixed-length block of contiguous virtual addresses copied as a unit from memory to disk and back during paging operations.
page fault	The interrupt that occurs when software attempts to read from or write to a virtual memory location that is marked "not present."
page restore	An operation that restores one or more data pages. Page restore is intended for repairing isolated damaged pages.
pager	A pocket-sized wireless electronic device that uses radio signals to record incoming phone numbers or short text messages. Some pagers allow users to send messages as well.
paging system	A system that allows users to send and receive messages when they are out of range.

Term	Definition
PAL	The primary mechanism for securing the Publisher. It contains a list of logins, accounts, and groups that are granted access to the publication.
parallel execution	The apparently simultaneous execution of two or more routines or programs. Concurrent execution can be accomplished on a single process or by using time-sharing techniques, such as dividing programs into different tasks or threads of execution, or by using multiple processors.
parallel processing	A method of processing that can run only on a computer that contains two or more processors running simultaneously. Parallel processing differs from multiprocessing in the way a task is distributed over the available processors. In multiprocessing, a process might be divided up into sequential blocks, with one processor managing access to a database, another analyzing the data, and a third handling graphical output to the screen. Programmers working with systems that perform parallel processing must find ways to divide a task so that it is more or less evenly distributed among the processors available.
parameterized query	A query that accepts input values through parameters.
parameterized report	A published report that accepts input values through parameters.
parameterized row filter	A row filter available with merge replication that allows you to restrict the data replicated to a Subscriber based on a system function or user-defined function (for example: <code>SUSER_SNAME()</code> ).
partial backup	A backup of all the data in the primary filegroup, every read-write filegroup, and any optionally specified files. A partial backup of a read-only database contains only the primary filegroup.
partial database restore	A restore of only a portion of a database consisting of its primary filegroup and one or more secondary filegroups. The other filegroups remain permanently offline, though they can be restored later.
partial differential backup	A partial backup that is differential relative to a single, previous partial backup (the base backup). For a read-only database, a partial differential backup contains only the primary filegroup.
participant	A synchronization provider and its associated replica.
particle	A very small piece or part; an indivisible object.
partition function	A function that defines how the rows of a partitioned table or index are spread across a set of partitions based on the values of certain columns, called partitioning columns.
partition scheme	A database object that maps the partitions of a partition function to a set of filegroups.
partitioned index	An index built on a partition scheme, and whose data is horizontally divided into units which may be spread across more than one filegroup in a database.
partitioned snapshot	In merge replication, a snapshot that includes only the data from a single partition.

Term	Definition
partitioned table	A table built on a partition scheme, and whose data is horizontally divided into units which may be spread across more than one filegroup in a database.
partitioned table parallelism	The parallel execution strategy for queries that select from partitioned objects. As part of the execution strategy, the query processor determines the table partitions that are required for the query and the proportion of threads to allocate to each partition. In most cases, the query processor allocates an equal or almost equal number of threads to each partition, and then executes the query in parallel across the partitions.
partitioning	The process of replacing a table with multiple smaller tables.
partitioning column	The column of a table or index that a partition function uses to partition a table or index.
partner	In database mirroring, refers to the principal server or the mirror server.
pass order	The order of evaluation (from highest to lowest calculation pass number) and calculation (from lowest to highest calculation pass number) for calculated members, custom members, custom rollup formulas, and calculated cells in a multidimensional cube.
pass phrase	A sequence of words or other text used to gain access to a network, program, or data. A passphrase is generally longer for added security.
pass-through query	An SQL-specific query you use to send commands directly to an ODBC database server.
pass-through statement	A SELECT statement that is passed directly to the source database without modification or delay.
password policy	A collection of policy settings that define the password requirements for a Group Policy object (GPO).
password provider	A one-time-password generation and validation component for user authentication.
path	A data flow element that connects the output of one data flow component to the input of another data flow component.
PBM	A set of built-in functions that return server state information about values, objects, and settings in SQL Server. Policy Based Management allows a database administrator to declare the desired state of the system and checks the system for compliance with that state.
peer-to-peer replication	A type of transactional replication. In contrast to read-only transactional replication and transactional replication with updating subscriptions, the relationships between nodes in a Peer to Peer replication topology are peer relationships rather than hierarchical ones, with each node containing identical schema and data.
performance tools	Tools that you can use to evaluate the performance of a solution. Performance tools can have different purposes; some are designed to evaluate end-to-end performance while others focus on evaluating performance of a particular aspect of a solution.

Term	Definition
persisted computed column	A computed column of a table that is physically stored, and whose values are updated when any other columns that are part of its computation change. Applying the persisted property to a computed column allows for indexes to be created on it when the column is deterministic, but not precise.
persistence database	A type of persistence store that stores workflow instance state and workflow instance metadata in a SQL Server database.
perspective	A user-defined subset of a cube, whereas a view is a user-defined subset of tables and columns in a relational database.
pessimistic concurrency	A method of managing concurrency by using a lock technique to prevent other clients from updating the same object at the same time.
phantom	Pertaining to the insertion of a new row or the deletion of an existing row in a range of rows that were previously read by another task, where that task has not yet committed its transaction.
PHP	An open source scripting language that can be embedded in HTML documents to execute interactive functions on a Web server. It is generally used for Web development.
PHP:Hypertext Preprocessor	An open source scripting language that can be embedded in HTML documents to execute interactive functions on a Web server. It is generally used for Web development.
physical name	The name of the path where a file or mirrored file is located. The default is the path of the Master.dat file followed by the first eight characters of the file's logical name.
physical storage	The amount of RAM memory in a system, as distinguished from virtual memory.
Pickup folder	The directory from which messages are picked up.
piecemeal restore	A composite restore in which a database is restored in stages, with each stage corresponding to a restore sequence. The initial sequence restores the files in the primary filegroup, and, optionally, other files, to any point in time supported by the recovery model and brings the database online. Subsequent restore sequences bring remaining files to the point consistent with the database and bring them online.
PIN	A unique and secret identification code similar to a password that is assigned to an authorized user and used to gain access to personal information or assets via an electronic device.
PInvoke	The functionality provided by the common language runtime to enable managed code to call unmanaged native DLL entry points.
pivot	To rotate a table-valued expression by turning the unique values from one column in the expression into multiple columns in the output, and perform aggregations where they are required on any remaining column values that are wanted in the final output.
pivot currency	The currency against which exchange rates are entered in the rate measure group.

<b>Term</b>	<b>Definition</b>
PivotTable	An interactive technology in Microsoft Excel or Access that can show dynamic views of the same data from a list or a database.
PL/SQL	Oracle's data manipulation language that allows sequenced or grouped execution of SQL statements and is commonly used to manipulate data in an Oracle database. The syntax is similar to the Ada programming language.
placeholder item	An artificial address, instruction, or other datum fed into a computer only to fulfill prescribed conditions and not affecting operations for solving problems.
plaintext	Data in its unencrypted or decrypted form.
plan guide	A SQL Server module that attaches query hints to queries in deployed applications, without directly modifying the query.
planar	In computer graphics, lying within a plane.
platform invoke	The functionality provided by the common language runtime to enable managed code to call unmanaged native DLL entry points.
platform layer	The layer that includes the physical servers and services that support the services layer. The platform layer consists of many instances of SQL Server, each of which is managed by the SQL Azure fabric.
PMML	An XML-based language that enables sharing of defined predictive models between compliant vendor applications.
Point and Figure chart	A chart that plots day-to-day price movements without taking into consideration the passage of time.
point depth	The depth of data points displayed in a 3D chart area.
point event	An event occurrence as of a single point in time. Only the start time is required for the event. The CEP server infers the valid end time by adding a tick (the smallest unit of time in the underlying time data type) to the start time to set the valid time interval for the event. Point events are valid only for this single instant of time.
point event model	The event model of a point event.
point layer	The layer in a map report that displays spatial data as points, for examples, points that indicate cities or points of interest.
pointer	A needle, marker, or bar that indicates a single value displayed against a scale on a report.
point-in-time recovery	The process of recovering only the transactions within a log backup that were committed before a specific point in time, instead of recovering the whole backup.
poison message	A message containing information that an application cannot successfully process. A poison message is not a corrupt message, and may not be an invalid request.

Term	Definition
policy module	A Certificate Services component that determines whether a certificate request should be automatically approved, denied, or marked as pending.
policy-based management	A set of built-in functions that return server state information about values, objects, and settings in SQL Server. Policy Based Management allows a database administrator to declare the desired state of the system and checks the system for compliance with that state.
poller	A component or interface that monitors the status of other components. It performs this function by repeatedly polling the other component to provide its current status.
polling	The process of periodically determining the status of each device in a set so that the active program can process the events generated by each device, such as whether a mouse button was pressed or whether new data is available at a serial port. This can be contrasted with event-driven processing, in which the operating system alerts a program or routine to the occurrence of an event by means of an interrupt or message rather than having to check each device in turn.
polling query	A singleton query that returns a value Analysis Services can use to determine if changes have been made to a table or other relational object.
polygon layer	The layer in a map report that displays spatial data as areas, for example, areas that indicate geographical regions such as counties.
population	The process of scanning content to compile and maintain an index.
positioned update	An update, insert, or delete operation performed on a row at the current position of the cursor.
Power Pivot	A SQL Server add-in for Excel.
Power Pivot data	A SQL Server Analysis Services cube that is created and embedded through Microsoft SQL Server Power Pivot for Microsoft Excel.
Power Pivot for SharePoint	A Microsoft SQL Server technology that provides query processing and management control for Power Pivot workbooks published to SharePoint.
Power Pivot service	A middle tier service of the Analysis Services SharePoint integration feature that allocates requests, monitors server availability and health, and communicates with other services in the farm.
Power Pivot service application	A specific configuration of the Power Pivot service.
Power Pivot Web service	A Web service that performs request redirection for processing requests that are directed to a Power Pivot Engine service instance that is outside the farm.
Power Pivot workbook	An Excel 2010 workbook that contains Power Pivot data.



Term	Definition
PowerShell-based cache administration tool	The exclusive management tool for Windows Server AppFabric. With more than 130 standard command-line tools, this new administration-focused scripting language helps you achieve more control and productivity.
precedence constraint	A control flow element that connects tasks and containers into a sequenced workflow.
precomputed partition	A performance optimization that can be used with filtered merge publications.
predictable column	A data mining column that the algorithm will build a model around based on values of the input columns. Besides serving as an output column, a predictable column can also be used as input for other predictable columns within the same mining structure.
prediction	A data mining technique that analyzes existing data and uses the results to predict values of attributes for new records or missing attributes in existing records. For example, existing credit application data can be used to predict the credit risk for a new application.
Prediction Calculator	A new report that is based on logistic regression analysis and that presents each contributing factor together with a score calculated by the algorithm. The report is presented both as a worksheet that helps you enter data and make calculations of the probable outcomes, and as a printed report that does the same thing.
prefix characters	A set of 1 to 4 bytes that prefix each data field in a native-format bulk-copy data file.
prefix length	The number of prefix characters preceding each noncharacter field in a bcp native format data file.
prerequisite knowledge	The minimum knowledge that a destination provider is required to have to process a change or change batch.
presentation model	A data model that aggregates data from multiple entities in the data access layer. It is used to avoid directly exposing an entity to the client project.
primary database	A read-write database containing the active data of an application.
primary dimension table	In a snowflake schema in a data warehouse, a dimension table that is directly related to, and usually joined to, the fact table.
primary DPM server	A DPM server that protects file or application data sources.
primary key	One or more fields that uniquely identify each record in a table. In the same way that a license plate number identifies a car, the primary key uniquely identifies a record.
primary protection	A type of protection in which data on the protected server is directly protected by a primary DPM Server.

Term	Definition
primary server	In a log shipping configuration, the server instance where the primary database resides.
primary table	The "one" side of two related tables in a one-to-many relationship. A primary table should have a primary key and each record should be unique.
principal database	In database mirroring, a read-write database whose transaction log is continuously sent to the mirror server, which restores the log to the mirror database.
principal server	In database mirroring, the partner whose database is currently the principal database.
priority boost	An advanced option that specifies whether Microsoft® SQL Server™ should run at a higher Microsoft Windows NT® scheduling priority than other processes on the same computer.
private key	The secret half of a cryptographic key pair that is used with a public key algorithm. Private keys are typically used to decrypt a symmetric session key, digitally sign data, or decrypt data that has been encrypted with the corresponding public key.
proactive caching	A system that manages data obsolescence in a cube by which objects in MOLAP storage are automatically updated and processed in cache while queries are redirected to ROLAP storage.
procedure cache	The part of the SQL Server memory pool that is used to store execution plans for Transact-SQL batches, stored procedures, and triggers.
profile template	The core of all Certificate Lifecycle Manager (CLM) management activities. The profile template provides a single administrative unit that includes all information necessary to manage the multiple certificates that might be required by a user community throughout the certificate's lifecycle. A profile template also includes information regarding the final location of those certificates, which can be software-based (that is, stored on the local computer) or hardware-based (stored on a smart card). A profile template cannot include both software-based and smart card-based certificates.
profit chart	A diagram that displays the theoretical increase in profit that is associated with using various data models.
programmable	Capable of accepting instructions for performing a task or an operation. Being programmable is a characteristic of computers.
properties page	A dialog box that displays information about an object in the interface.
property	Attribute or characteristic of an object that is used to define its state, appearance, or value.
property mapping	A mapping between a variable and a property of a package element.
property page	A grouping of properties presented as a tabbed page of a property sheet.
protected computer	A computer that contains data sources that are protection group members.

Term	Definition
protected member	A data source within a protection group.
protocol	A standard set of formats and procedures that enable computers to exchange information.
provider	An in-process dynamic link library (DLL) that provides access to a database.
provider	A software component that allows a replica to synchronize its data with other replicas.
provider object	An object that's part of a data provider such as Oracle Provider for SQL Server.
proximity search	Full-text query searching for those occurrences where the specified words are close to one another.
proxy account	An account that is used to provide additional permissions for certain actions to users which do not have these permissions but have to execute these actions.
public key	The nonsecret half of a cryptographic key pair that is used with a public key algorithm. Public keys are typically used when encrypting a session key, verifying a digital signature, or encrypting data that can be decrypted with the corresponding private key.
publication	A collection of one or more articles from one database.
publication access list	The primary mechanism for securing the Publisher. It contains a list of logins, accounts, and groups that are granted access to the publication.
publication database	A database on the Publisher from which data and database objects are marked for replication and propagated to Subscribers.
publication retention period	In merge replication, the amount of time a subscription can remain unsynchronized.
published data	Data at the Publisher that has been replicated.
Publisher	A server that makes data available for replication to other servers. A Publisher also detects changed data and maintains information about all publications at the site.
publisher database	A server that makes data available for replication to other servers. A Publisher also detects changed data and maintains information about all publications at the site.
publishing server	A server running an instance of Analysis Services that stores the source cube for one or more linked cubes.
publishing table	The table at the Publisher in which data has been marked for replication and is part of a publication.
pull	The process of retrieving data from a network server.

Term	Definition
pull replication	Replication that is invoked at the target.
pull subscription	A subscription created and administered at the Subscriber. The Distribution Agent or Merge Agent for the subscription runs at the Subscriber.
pure log backup	A backup containing only the transaction log covering an interval without any bulk changes.
push	The process of sending data to a network server.
push replication	Replication that is invoked at the source.
push subscription	A subscription created and administered at the Publisher.
qualified component	A method of identifying a component in an MSI database indirectly by a pair "component category GUID, string qualifier" instead of identifying it directly by the component identifier.
qualifier	A modifier containing information that describes a class, instance, property, method, or parameter. Qualifiers are defined by the Common Information Model (CIM), by the CIM Object Manager, and by developers.
qualifier flavor	A flag that provides additional information about a qualifier, such as whether a derived class or instance can override the qualifier's original value.
quantum	A brief period of time when a given thread executes in a multitasking operating system. It performs the multitasking before it is rescheduled against other threads with the same priority. Previously known as a "time slice."
query	An instance of a query template that runs continuously in the StreamInsight server processing events received from instances of input adapters to which the query is bound and sending processed events to instances of output adapters to which it is bound.
query binder	An object that binds an existing StreamInsight query template to specific input and output adapters.
query binding	The process of binding instances of input adapters and instances of output adapters to an instance of a query template.
query designer	A tool that helps a user create the query command that specifies the data the user wants in a report dataset.
query governor	A configuration option that can be used to prevent system resources from being consumed by long-running queries.
query hint	A hint that specifies that the indicated hints should be used throughout the query. Query hints affect all operators in the statement.
query optimizer	A database engine component responsible for generating efficient execution plans for statements.

Term	Definition
query template	The fundamental unit of query composition. A query template defines the business logic required to continuously analyze and process events submitted to and emitted by the StreamInsight server.
question template	A structure that describes a set of questions that can be asked using a particular relationship or set of relationships.
Queue Reader Agent	The executable that reads the changes from subscribers in the queue and delivers the changes to the publisher.
quorum	In a database mirroring session with a witness server, a relationship in which the servers that can currently communicate with each other arbitrate who owns the role of principal server.
quoted identifier	An object in a database that requires the use of special characters (delimiters) because the object name does not comply with the formatting rules of regular identifiers.
RAD	A method of building computer systems in which the system is programmed and implemented in segments, rather than waiting until the entire project is completed for implementation. Developed by programmer James Martin, RAD uses such tools as CASE and visual programming.
ragged hierarchy	A hierarchy in which one or more levels do not contain members in one or more branches of the hierarchy.
range	A set of continuous item identifiers to which the same clock vector applies. A range is represented by a starting point, an ending point, and a clock vector that applies to all IDs that are in between.
range partition	A table partition that is defined by specific and customizable ranges of data.
range partitioning	A way of partitioning a table or index by specifying partitions to hold rows with ranges of values from a single partitioning column.
range query	A query that specifies a range of values as part of the search criteria, such as all rows from 10 through 100.
ranking function	Function that returns ranking information about each row in the window (partition) of a result set depending on the row's position within the window.
rapid application development	A method of building computer systems in which the system is programmed and implemented in segments, rather than waiting until the entire project is completed for implementation. Developed by programmer James Martin, RAD uses such tools as CASE and visual programming.
rate of change	The rate of price change compared with historical data. The rate of change is calculated against a period of days prior to the current price. The output is a percentage.
raw destination adapter	An SSIS destination that writes raw data to a file.
raw file	A native format for fast reading and writing of data.

Term	Definition
Raw File destination	An SSIS destination that writes raw data to a file.
Raw File source	An SSIS source that reads raw data from a file.
raw source adapter	An SSIS source that reads raw data from a file.
RDA	A service that provides a simple way for a smart device application to access (pull) and send (push) data to and from a remote SQL Server database table and a local SQL Server Mobile Edition database table. RDA can also be used to issue SQL commands on a server running SQL Server.
RDBMS	A database system that organizes data into related rows and columns as specified by a relational model.
RDL	A set of instructions that describe layout and query information for a report. RDL is composed of XML elements that conform to an XML grammar created for Reporting Services.
RDL Sandboxing	A feature that makes it possible to detect and restrict specific types of resource use by individual tenants in a scenario where multiple tenants share a single Web farm of report servers.
record set	A data structure made up of a group of database records. It can originate from a base table or from a query to the table.
recover	To put back into a stable condition. A computer user may be able to recover lost or damaged data by using a program to search for and salvage whatever information remains in storage. A database may be recovered by restoring its integrity after some problem has damaged it, such as abnormal termination of the database management program.
recovery	A phase of database startup that brings the database into a transaction-consistent state. Recovery can include rolling forward all the transactions in the log records (the redo phase) and rolling back uncommitted transactions (the undo phase), depending on how the database was shut down.
recovery branch	A range of LSNs that share the same recovery branch GUID. A new recovery branch originates when a database is created or when RESTORE WITH RECOVERY generates a recovery fork. A multiple-branch recovery path is possible that includes ranges of LSNs that cover two or more recovery fork points.
recovery fork point	The point (LSN,GUID) at which a new recovery branch is started, every time a RESTORE WITH RECOVERY is performed. Each recovery fork determines a parent-child relationship between the recovery branches. If you recover a database to an earlier point in time and begin using the database from that point, the recovery fork point starts a new recovery path.
recovery interval	The maximum amount of time that the Database Engine should require to recover a database.
recovery model	A database property that controls the basic behavior of backup and restore operations for a database. For instance, the recovery model controls how transactions are logged, whether the transaction log requires backing up, and what kinds of restore operations are available.

Term	Definition
recovery path	The sequence of data and log backups that have brought a database to a particular point in time (known as a recovery point). A recovery path is a specific set of transformations that have evolved the database over time, yet have maintained the consistency of the database. A recovery path describes a range of LSNs from a start point (LSN,GUID) to an end point (LSN,GUID). The range of LSNs in a recovery path can traverse one or more recovery branches from start to end.
recovery point	The point in the log chain at which rolling forward stops during a recovery.
rectangle	A report item that can be used as a container for multiple report items or as a graphical element on a report.
recursive hierarchy	A hierarchy of data in which all parent-child relationships are represented in the data.
recursive partitioning	The iterative process, used by data mining algorithm providers, of dividing data into groups until no more useful groups can be found.
redo	The phase during recovery that applies (rolls forward) logged changes to a database to bring the data forward in time.
redo phase	The phase during recovery that applies (rolls forward) logged changes to a database to bring the data forward in time.
redo set	The set of all files and pages being restored.
reference data	Data characterized by shared read operations and infrequent changes. Examples of reference data include flight schedules and product catalogs. Windows Server AppFabric offers the local cache feature for storing this type of data.
reference dimension	A relationship between a dimension and a measure group in which the dimension is coupled to the measure group through another dimension.
reference table	The source table to use in fuzzy lookups.
referenced key	A primary key or unique key referenced by a foreign key.
referencing key	A key in a database table that comes from another table (also know as the "referenced table") and whose values match the primary key (PK) or unique key in the referenced table.
reflexive relationship	A relationship from a column or combination of columns in a table to other columns in that same table.
region	A container of data, within a cache, that co-locates all cached objects on a single cache host. Cache Regions enable the ability to search all cached objects in the region by using descriptive strings, called tags.
region	A collection of 128 leaf level pages in logical order in a single file. Used to identify areas of a file that are fragmented.

Term	Definition
registration model	A defined method for submitting and approving enrollment requests. An enterprise generally chooses one registration model and modifies their management policies accordingly.
regression	The statistical process of predicting one or more continuous variables, such as profit or loss, based on other attributes in the dataset.
regressor	An input variable that has a linear relationship with the output variable.
relational database management system	A database system that organizes data into related rows and columns as specified by a relational model.
relational engine	A component of SQL Server that works with the storage engine. It is responsible for interpreting Transact-SQL search queries and mapping out the most efficient methods of searching the raw physical data provided by the storage engine and returning the results to the user.
relational OLAP	A storage mode that uses tables in a relational database to store multidimensional structures.
relational store	A data repository structured in accordance with the relational model.
relationship object	An object representing a pair of objects that assume a role in relation to each other.
relative date	A range of dates that is specified by using comparison operators and return data for a range of dates.
remote data	Data stored in an OLE DB data source that is separate from the current instance of SQL Server.
remote data access	A service that provides a simple way for a smart device application to access (pull) and send (push) data to and from a remote SQL Server database table and a local SQL Server Mobile Edition database table. RDA can also be used to issue SQL commands on a server running SQL Server.
remote Distributor	A server configured as a Distributor that is separate from the server configured as the Publisher.
remote login identification	The login identification assigned to a user for accessing remote procedures on a remote server.
remote partition	A partition whose data is stored on a server running an instance of Analysis Services, other than the one used to store the metadata of the partition.
remote service binding	A Service Broker object that specifies the local security credentials for a remote service.
remote stored procedure	A stored procedure located on one instance of SQL Server that is executed by a statement on another instance of SQL Server.



<b>Term</b>	<b>Definition</b>
remote table	A table stored in an OLE DB data source that is separate from the current instance of SQL Server.
rendered output	The output from a rendering extension.
rendered report	A fully processed report that contains both data and layout information, in a format suitable for viewing (such as HTML).
rendering extension	A plug-in that renders reports to a specific format (for example, an Excel rendering extension)
rendering object model	Report object model used by rendering extensions.
replica	A complete copy of protected data residing on a single volume on the DPM server. A replica is created for each protected data source after it is added to its protection group. With co-location, multiple data sources can have their replicas residing on the same replica volume.
replica	A particular repository of information to be synchronized.
replica ID	A value that uniquely identifies a replica.
replica key	A 4-byte value that maps to a replica ID in a replica key map.
replica tick count	A monotonically increasing number that is used to uniquely identify a change to an item in a replica.
replication	The process of copying content and/or configuration settings from one location, generally a server node, to another. Replication is done to ensure synchronization or fault tolerance.
Replication Management Objects	A managed code assembly that encapsulates replication functionalities for SQL Server.
Report Builder	A report authoring tool that features a Microsoft Office-like authoring environment and features such as new sparkline, data bar, and indicator data visualizations, the ability to save report items as report parts, a wizard for creating maps, aggregates of aggregates, and enhanced support for expressions.
report data pane	A data pane that displays a hierarchical view of the items that represent data in the user's report. The top level nodes represent built-in fields, parameters, images, and data source references.
report definition	The blueprint for a report before the report is processed or rendered. A report definition contains information about the query and layout for the report.
Report Definition Language	A set of instructions that describe layout and query information for a report. RDL is composed of XML elements that conform to an XML grammar created for Reporting Services.
Report Designer	A collection of design surfaces and graphical tools that are hosted within the Microsoft Visual Studio environment.

<b>Term</b>	<b>Definition</b>
report execution snapshot	A report snapshot that is cached. Report administrators create report execution snapshots if they want to run reports from static copies.
report history	A collection of previously run copies of a report.
report history snapshot	Report history that contains data captured at a specific point in time.
report intermediate format	Internal representation of a report.
report item	Entity on a report.
report layout template	A pre-designed table, matrix, or chart report template in Report Builder.
report link	URL to a report.
Report Manager	A Web-based report management tool
report model	A metadata description of business data used for creating ad hoc reports.
report part	A report item that has been published separately to a report server and that can be reused in other reports.
Report Processor component	A component that retrieves the report definition from the report server database and combines it with data from the data source for the report.
Report Project	A template in the report authoring environment.
Report Project Wizard	A wizard in the report authoring environment used to create reports.
report rendering	The action of combining the report layout with the data from the data source for the purpose of viewing the report.
report server	A location on the network where Report Builder is launched from and a report is saved, managed, and published.
report server administrator	A user with elevated privileges who can access all settings and content of a report server. A report server administrator is a user who is assigned to the Content Manager role, the System Administrator role, or both. All local administrators are automatically report server administrators, but additional users can become report server administrators for all or part of the report server namespace.

Term	Definition
report server database	A database that provides internal storage for a report server.
report server execution account	The account under which the Report Server Web service and Report Server Windows service run.
Report Server Web service	A Web service that hosts, processes, and delivers reports.
report snapshot	A static report that contains data captured at a specific point in time.
report-specific schedule	Schedule defined inline with a report. Report-specific schedules are defined in the context of an individual report, subscription, or report execution operation to determine cache expiration or snapshot updates.
repository	A database containing information models that, in conjunction with the executable software, manage the database.
repository engine	Object-oriented software that provides management support for and customer access to a repository database.
repository object	A COM object that represents a data construct stored in a repository type library.
Repository SQL schema	A set of standard tables used by the repository engine to manage all repository objects, relationships, and collections.
Repository Type Information Model	A core object model that represents repository type definitions for Metadata Services.
republisher	A Subscriber that publishes data that it has received from a Publisher.
reserved character	A keyboard character that has a special meaning to a program and, as a result, normally cannot be used in assigning names to files, documents, and other user-generated tools, such as macros. Characters commonly reserved for special uses include the asterisk (*), forward slash (/), backslash (\), question mark (?), and vertical bar ( ).
resolution strategy	A set of criteria that the repository engine evaluates sequentially when selecting an object, where multiple versions exist and version information is unspecified in the calling program.
resource	A special variable that holds a reference to a database connection or statement.
resource	Any item in a report server database that is not a report, folder, or shared data source item.
resource data	A type of data that is characterized by shared, concurrently read and written into operations, and accessed by many transactions. Examples of resource data include user accounts and auction items.

Term	Definition
resource governor	A feature in SQL Server 2008 that enables the user to manage SQL Server workload and resources by specifying limits on resource consumption by incoming requests.
restore	A multi-phase process that copies all the data and log pages from a specified backup to a specified database (the data-copy phase) and rolls forward all the transactions that are logged in the backup (the redo phase). At this point, by default, a restore rolls back any incomplete transactions (the undo phase), which completes the recovery of the database and makes it available to users.
restore sequence	A sequence of one or more restore commands that, typically, initializes the contents of the database, files, and/or pages being restored (the data-copy phase), rolls forward logged transactions (the redo phase), and rolls back uncommitted transactions (the undo phase).
result set	The set of records that results from running a query or applying a filter.
results	The set of records that results from running a query or applying a filter.
results pane	A pane that displays details about an item selected in another portion of the user interface. For example, in Microsoft Management Console (MMC), the details pane is the right pane that displays details for the selected item in the console tree.
retract event	An internal event kind used to modify an existing insert event by modifying the end time of the event.
reusable bookmark	A bookmark that can be consumed from a rowset for a given table and used on a different rowset of the same table to position on a corresponding row.
revocation delay	The period of time between when the credential revocation request is placed and when the credentials are actually revoked.
RIA	A web application that provides a user interface which is more similar to a desktop application than typical web pages. It is able to process user actions without posting the whole web page to a web server.
RIA Services link	A project-to-project link reference that facilitates generating presentation tier code from middle tier code.
rich Internet application	A web application that provides a user interface which is more similar to a desktop application than typical web pages. It is able to process user actions without posting the whole web page to a web server.
right outer join	An outer join in which all the records from the right side of the RIGHT JOIN operation in the query's SQL statement are added to the query's results, even if there are no matching values in the joined field from the table on the left.
ring index	An index that indicates the number of rings in a polygon instance.
RMO	A managed code assembly that encapsulates replication functionalities for SQL Server.
role assignment	The assignment of a specific role that determines whether a user or group can access a specific item and perform an operation on it.
role definition	The collection of task permissions associated with a role.

<b>Term</b>	<b>Definition</b>
role switching	In a database mirroring session, the taking over of the principal role by the mirror.
role-playing dimension	A single database dimension joined to the fact table on a different foreign keys to produce multiple cube dimensions.
roll back	To reverse changes.
roll forward	To apply logged changes to the data in a roll forward set to bring the data forward in time.
roll forward set	The set of all data restored by a restore sequence. A roll forward set is defined by restoring a series of one or more data backups.
roll up	To collect subsets of data from multiple locations in one location.
rollover file	A file created when the file rollover option causes SQL Server to close the current file and create a new file when the maximum file size is reached.
route	A Service Broker object that specifies the network address for a remote service.
routing client	A type of cache client that includes a routing table that is maintained by lead hosts in the cluster and enables the client to obtain cached data directly from the cache host on which the data resides.
routing table	A data structure used by the routing client to track the connectivity information of all cache hosts in the cache cluster. It is maintained by lead hosts in the cluster. It allows a routing client to obtain cached data directly from the cache host on which the data resides.
row aggregate function	A function that generates summary values, which appear as additional rows in the query results.
row lock	A lock on a single row in a table.
row versioning	In Online index operations, a feature that isolates the index operation from the effects of modifications that are made by other transactions.
row-overflow data	varchar, nvarchar, varbinary, or sql_variant data stored off the main data page of a table or index as a result of the combined widths of these columns exceeding the 8,060-byte row limit in a table.
rowset	A set of rows in which each row has one or more columns of data.
rs utility	Report scripting tool.
rs.exe	Report scripting tool.
rsconfig utility	Server connection management tool.
rsconfig.exe	Server connection management tool.

Term	Definition
rule firing	The process of running one of the application rules (event chronicle rules, subscription event rules, and subscription scheduled rules) defined in the application definition file.
runaway query	A query with an excessive running time, that can lead to a blocking problem. Runaway queries usually do not use use a query or lock time out.
run-time error	A software error that occurs while a program is being executed, as detected by a compiler or other supervisory program.
safe code	Code that is executed by the common language runtime environment rather than directly by the operating system. Managed code applications gain common language runtime services such as automatic garbage collection, runtime type checking and security support, and so on. These services help provide uniform platform- and language-independent behavior of managed-code applications.
sampling	A statistical process that yields some inferential knowledge about a population or data set of interest as a whole by observing or analyzing a portion of the population or data set.
save process	The process of writing data to disk.
savepoint	A location to which a transaction can return if part of the transaction is conditionally canceled or encounters an error, hence offering a mechanism to roll back portions of transactions.
SBCS	A character encoding in which each character is represented by 1 byte. Single byte character sets are mathematically limited to 256 characters.
scalar	A factor, coefficient, or variable consisting of a single value (as opposed to a record, an array, or some other complex data structure).
scalar-valued function	A function that returns a single value, such as a string, integer, or bit value.
scale break line	A line drawn across a chart area to indicate a significant gap between a high and low range of values on the chart.
scale-out deployment	A deployment model in which an installation configuration has multiple report server instances sharing a single report server database.
Scheduling and Delivery Processor	A component of the report server engine that handles scheduling and delivery. Works with SQL Agent.
schema rowset	A specially defined rowset that returns metadata about objects or functionality on an instance of SQL Server or Analysis Services. For example, the OLE DB schema rowset DBSCHEMA_COLUMNS describes columns in a table, while the Analysis Services schema rowset MDSHEMA_MEASURES describes the measures in a cube.
schema snapshot	A snapshot that includes schema for published tables and objects required by replication (triggers, metadata tables, and so on), but not user data.
schema-aware	Pertaining to a processing method based on a schema that defines elements, attributes and types that will be used to validate the input and output documents.

Term	Definition
scope	The extent to which an identifier, such as an object or property, can be referenced within a program. Scope can be global to the application or local to the active document.
scope	The set of data that is being synchronized.
script memory	The local memory (the client-side RAM) that is used by a PHP script.
script pane	The text editor portion of the Table Designer.
scripting	Pertaining to the automation of user actions or the configuration of a standard state on a computer by means of scripts.
SDK	A set of routines (usually in one or more libraries) designed to allow developers to more easily write programs for a given computer, operating system, or user interface.
search condition	In a WHERE or HAVING clause, predicates that specify the conditions that the source rows must meet to be included in the SQL statement.
search key	The value that is to be searched for in a document or any collection of data.
secondary database	A read-only copy of a primary database.
secondary DPM server	A DPM server that protects one or more primary DPM servers in addition to file and application data.
secondary protection	A type of protection in which data on the protected server is protected by a primary DPM server and the replica on the primary DPM server is protected by a secondary DPM server.
secondary server	In a log shipping configuration, the server instance where the secondary database resides. At regular intervals, the secondary server copies the latest log backup from the primary database and restores the log to the secondary database. The secondary server is a warm standby server.
secret provider	A one-time-password generation and validation component for user authentication.
securable	Entities that can be secured with permissions. The most prominent securables are servers and databases, but discrete permissions can be set at a much finer level.
Secure Sockets Layer	The protocol that improves the security of data communication by using a combination of data encryption, digital certificates, and public key cryptography. SSL enables authentication and increases data integrity and privacy over networks. SSL does not provide authorization or nonrepudiation.
security extension	A component in Reporting Services that authenticates a user or group to a report server.
security ID	In Windows-based systems, a unique value that identifies a user, group, or computer account within an enterprise. Every account is issued a SID when it is created.

Term	Definition
security identifier	In Windows-based systems, a unique value that identifies a user, group, or computer account within an enterprise. Every account is issued a SID when it is created.
segmentation	A data mining technique that analyzes data to discover mutually exclusive collections of records that share similar attributes sets.
self-join	A join in which records from a table are combined with other records from the same table when there are matching values in the joined fields. A self-join can be an inner join or an outer join. In database diagrams, a self-join is called a reflexive relationship.
self-service registration model	A Certificate Lifecycle Manager (CLM) registration model in which a certificate subscriber performs or requests certificate management activities directly using a Web-based interface.
self-tracking entity	An entity built from a Text Template Transformation Toolkit (T4) template that has the ability to record changes to scalar, complex, and navigation properties.
Semantic Model Definition Language	A set of instructions that describe layout and query information for reports created in Report Builder.
semantic object	An object that can be represented by a database object or other real-world object.
semantic validation	The process of confirming that the elements of an XML file are logically valid.
semiadditive measure	A measure that can be summed along one or more, but not all, dimensions in a cube.
sensitive cursor	A cursor that can reflect data modifications made to underlying data by other users while the cursor is open.
sensitive data	Personally identifiable information (PII) that is protected in special ways by law or policy.
sequenced collection	A collection of destination objects of a sequenced relationship object.
sequenced relationship	A relationship in a repository that specifies explicit positions for each destination object within the collection of destination objects.
serial number	A number assigned to a specific inventory item to identify it and differentiate it from similar items with the same item number.
serialization	The process of converting an object's state information into a form that can be stored or transported. During serialization, an object writes its current state to temporary or persistent storage. Later, the object can be recreated by reading, or deserializing, the object's state from storage.
server	A computer that provides shared resources, such as files or printers, to network users.



<b>Term</b>	<b>Definition</b>
server collation	The collation for an instance of SQL Server.
server cursor	A cursor implemented on the server.
server name	A name that uniquely identifies a server computer on a network.
server subscription	A subscription to a merge publication with an assigned priority value used for conflict detection and resolution.
service	A program, routine, or process that performs a specific system function to support other programs.
Service Broker	A technology that helps developers build scalable, secure database applications.
service connection point	An Active Directory node on which system administrators can define Certificate Lifecycle Manager (CLM) management permissions for users and groups.
service principal	A globally-unique name associated with a service entity at an endpoint node within a Kerberos peer-to-peer communication.
service principal name	The name by which a client uniquely identifies an instance of a service. It is usually built from the DNS name of the host. The SPN is used in the process of mutual authentication between the client and the server hosting a particular service.
service program	A program that uses Service Broker functionality. A service program may be a Transact-SQL stored procedure, a SQLCLR stored procedure, or an external program.
session	A period of time when a connection is active and communication can take place. For the purpose of data communication between functional units, session also refers to all the activities that take place during the establishment, maintenance, and release of the connection.
session state	In ASP.NET, a variable store created on the server for the current user; each user maintains a separate Session state on the server. Session state is typically used to store user-specific information between postbacks.
set	A grouping of dimension members or items from a data source that are named and treated as a single unit and can be referenced or reused multiple times.
setup initialization file	A text file, using the Windows .ini file format, that stores configuration information allowing SQL Server to be installed without a user having to be present to respond to prompts from the Setup program.
setup repair	An error reporting process that may run during the setup of a program if a problem occurs.
shadow copy	A static image of a set of data, such as the records displayed as the result of a query.
shapefile	A public domain format for the interchange of spatial data in geographic information systems. Shapefiles have the file name extension ".shp".

Term	Definition
sharding	A technique for partitioning large data sets, which improves performance and scalability, and enables distributed querying of data across multiple tenants.
shared code	Code that is specifically designated to exist without modification in the server project and the client project.
shared data source item	Data source connection information that is encapsulated in an item. Can be managed as an item in the report server folder namespace.
shared dimension	A dimension created within a database that can be used by any cube in the database.
shared lock	A lock created by nonupdate (read) operations.
shared schedule	Schedule information that can be referenced by multiple items.
shopping basket analysis	A standard data mining algorithm that analyzes a list of transactions to make predictions about which items are most frequently purchased together.
showplan	A report showing the execution plan for an SQL statement.
SID	In Windows-based systems, a unique value that identifies a user, group, or computer account within an enterprise. Every account is issued a SID when it is created.
significance	One of the arguments of the FLOOR function.
Silverlight business application	A template that provides many common features for building a business application with a Silverlight client. It utilizes WCF RIA Services for authentication and registration services.
simple client	A type of cache client that does not have a routing table and thus does not need network connectivity to all cache hosts in the cache cluster. Because data traveling to simple clients from the cluster may need to travel across multiple cache hosts, simple clients may not perform as fast as a routing clients.
Simple Mail Transfer Protocol	A member of the TCP/IP suite of protocols that governs the exchange of electronic mail between message transfer agents.
simple recovery model	A database recovery mode that minimally logs all transactions sufficiently to ensure database consistency after a system crash or after restoring a data backup. The database is recoverable only up to the time of its most recent data backup, and restoring individual pages is unsupported.
single-byte character set	A character encoding in which each character is represented by 1 byte. Single byte character sets are mathematically limited to 256 characters.
single-precision	Of or pertaining to a floating-point number having the least precision among two or more options commonly offered by a programming language, such as single-precision versus double-precision.

Term	Definition
single-user mode	A state in which only one user can access a resource.
sink	A device or part of a device that receives something from another device.
SKU	A unique identifier, usually alphanumeric, for a product. The SKU allows a product to be tracked for inventory purposes. An SKU can be associated with any item that can be purchased. For example, a shirt in style number 3726, size 8 might have a SKU of 3726-8.
sleep	To suspend operation without terminating.
slice	A subset of the data in a cube, specified by limiting one or more dimensions by members of the dimension.
Slicers	A feature that provides one-click filtering controls that make it easy to narrow down the portion of a data set that's being looked at.
sliding window	A window of fixed length L that moves along a timeline according to the stream's events. With every event on the timeline, a new window is created, starting at the event's start time.
slipstream	To integrate updates, patches or service packs into the base installation files of the original software, so that the resulting files will allow a single step installation of the updated software.
slipstream installation	A type of installation that integrates the base installation files for an operating system or program with its service packs, updates or patches, and enables them to be installed in a single step.
smart card	A plastic (credit card-sized or smaller) device with an embedded microprocessor and a small amount of storage that is used, with an access code, to enable certificate-based authentication. Smart cards securely store certificates, public and private keys, passwords, and other types of personal information.
Smart Card Personalization Control	An ActiveX control that performs all Certificate Lifecycle Manager (CLM) smart card application management activities on a client computer.
smart card profile	A Certificate Lifecycle Manager (CLM) profile created when a request is performed using a profile template that only includes smart card-based certificate templates.
smart card reader	A device that is installed in computers to enable the use of smart cards for enhanced security features.
Smart Card Self Service Control	Software installed on a client computer that enables end users and administrators to manage smart cards by providing a connection from the client computer to the smart card.
smart card unblocking	The action of binding a smart card with administrative credentials to reset the the personal identification number (PIN) attempt counter.
SMO	An application programming interface that supports the incorporation of SQL Server administration into any COM or OLE Automation application.

Term	Definition
SMTP	A member of the TCP/IP suite of protocols that governs the exchange of electronic mail between message transfer agents.
snap-in	A type of tool that you can add to a console supported by Microsoft Management Console (MMC). A stand-alone snap-in can be added by itself; an extension snap-in can be added only to extend the function of another snap-in.
snapshot	A static report that contains data captured at a specific point in time.
snapshot isolation level	A transaction isolation level in which each read operation performed by a transaction returns all data as it existed at the start of the transaction. Because a snapshot transaction does not use locks to protect read operations, it will not block other transactions from modifying any data read by the snapshot transaction.
snapshot replication	A replication in which data is distributed exactly as it appears at a specific moment in time and does not monitor for updates to the data.
Snapshot Share	A share available for the storage of snapshot files. Snapshot files contain the schema and data for published tables.
snapshot window	A window that is defined according to the start and end times of the event in the stream, instead of a fixed grid along the timeline.
snowflake schema	An extension of a star schema such that one or more dimensions are defined by multiple tables. In a snowflake schema, only primary dimension tables are joined to the fact table. Additional dimension tables are joined to primary dimension tables.
soft page	A rendered page that can be slightly larger than the size specified using the InteractiveHeight and InteractiveWidth properties of a report (HTML and WinForm control).
soft page-break renderer	A rendering extension that maintains the report layout and formatting so that the resulting file is optimized for screen-based viewing and delivery, such as on a Web page or in the ReportViewer controls.
software development kit	A set of routines (usually in one or more libraries) designed to allow developers to more easily write programs for a given computer, operating system, or user interface.
software profile	A Certificate Lifecycle Manager (CLM) profile created when a request is performed using a profile template that only includes software-based certificate templates.
software transformer	A software module or routine that modifies the events (data) into a format expected by the output device, and emits the data to that device.
solution explorer	A component of Microsoft SQL Server Management Studio that allows you to view and manage items and perform item management tasks in a solution or a project.
solve order	The order of evaluation (from highest to lowest solve order) and calculation (from lowest to highest solve order) for calculated members, custom members, custom rollup formulas, and calculated cells in a single calculation pass of a multidimensional cube.

Term	Definition
sort order	A way to arrange data based on value or data type. You can sort data alphabetically, numerically, or by date. Sort orders use an ascending (1 to 9, A to Z) or descending (9 to 1, Z to A) order.
source	A disk, file, document, or other collection of information from which data is taken or moved.
source	The SSIS data flow component that makes data from different external data sources available to the other components in the data flow.
source	A synchronization provider that enumerates any changes and sends them to the destination provider.
source adapter	A data flow component that extracts data from a data store.
source code control	A set of features that include a mechanism for checking source code in and out of a central repository. It also implies a version control system that can manage files through the development lifecycle, keeping track of which changes were made, who made them, when they were made, and why.
source control	A set of features that include a mechanism for checking source code in and out of a central repository. It also implies a version control system that can manage files through the development lifecycle, keeping track of which changes were made, who made them, when they were made, and why.
source cube	The cube on which a linked cube is based.
source database	A database on the Publisher from which data and database objects are marked for replication as part of a publication that is propagated to Subscribers. For a database view, the database on which the view is created.
source object	The single object to which all objects in a particular collection are connected by way of relationships that are all of the same relationship type.
source partition	An Analysis Services partition that is merged into another and is deleted automatically at the end of the merger process.
source provider	A synchronization provider that enumerates any changes and sends them to the destination provider.
sparkline	A miniature chart that can be inserted into text or embedded within a cell on a worksheet to illustrate highs, lows, and trends in your data.
sparse column	A column that reduces the storage requirement for null values at the cost of more overhead to retrieve nonnull values.
sparse file	A file that is handled in a way that requires much less disk space than would otherwise be needed. Sparse support allows an application to create very large files without committing disk space for those regions of the file that contain only zeros. For example, you can use sparse support to work with a 42-GB file in which you need to write data only to the first 64 KB (the rest of the file is zeroed).
sparsity	The relative percentage of a multidimensional structure's cells that do not contain data.

Term	Definition
spatial data	Data that is represented by 2D or 3D images. Spatial data can be further subdivided into geometric data (data that can use calculations involving Euclidian geometry) and geographic data (data that identifies geographic locations and boundaries on the earth).
SPN	The name by which a client uniquely identifies an instance of a service. It is usually built from the DNS name of the host. The SPN is used in the process of mutual authentication between the client and the server hosting a particular service.
SQL	A database query and programming language widely used for accessing, querying, updating, and managing data in relational database systems.
SQL database	A database based on Structured Query Language (SQL).
SQL expression	Any combination of operators, constants, literal values, functions, and names of tables and fields that evaluates to a single value.
SQL Native Client	A stand-alone data access API that is used for both OLE DB and ODBC.
SQL Server	A family of Microsoft relational database management and analysis systems for e-commerce, line-of-business, and data warehousing solutions.
SQL Server 2005 Express Edition	An edition of a Microsoft relational database design and management system for e-commerce, line-of-business, and data warehousing solutions.
SQL Server 2005 Mobile Edition	SQL Server product name (edition)
SQL Server Agent	A Microsoft Windows service that executes scheduled administrative tasks, which are called jobs, and stores the information in SQL Server.
SQL Server Analysis Services	A feature of Microsoft SQL Server that supports online analytical processing (OLAP) and data mining for business intelligence applications. Analysis Services organizes data from a data warehouse into cubes with precalculated aggregation data to provide rapid answers to complex analytical queries.
SQL Server Browser	The Windows service that listens for incoming requests for Microsoft SQL Server resources and provides information about SQL Server instances installed on the computer.
SQL Server Compact	A Microsoft relational database management and analysis system for e-commerce, line-of-business, and data warehousing solutions.
SQL Server component	A SQL Server program module developed to perform a specific set of tasks - e.g., data transformation, data analysis, reporting.
SQL Server Configuration Manager	A tool to manage the services associated with SQL Server, to configure the network protocols used by SQL Server, and to manage the network connectivity configuration from SQL Server client computers.

Term	Definition
SQL Server Connection Director	A connectivity technology where applications based on different data access technologies (.NET or native Win32) can share the same connection information. Connection information can be centrally managed for such client applications.
SQL Server Data Mining Content Viewer	A viewer that displays the content that is contained in the content schema rowset of the mining model.
SQL Server Data Mining Content Viewer Controls	A set of server-side controls that allow a user to browse complex mining models from any computer that has Microsoft Internet Explorer installed.
SQL Server Data Quality Services	A knowledge-based data-quality system that enables users to perform knowledge discovery and management, data cleansing, data matching, integration with reference data services, and integrated profiling.
SQL Server data-tier application project	A Visual Studio project used by database developers to create and develop a DAC. DAC projects get full support from Visual Studio and VSTS source code control, versioning, and development project management.
SQL Server End-User Recovery	A tool for SQL Server that enables backup administrators to authorize end users to recover backups of SQL Server databases from DPM, without further action from the backup administrator.
SQL Server EUR	A tool for SQL Server that enables backup administrators to authorize end users to recover backups of SQL Server databases from DPM, without further action from the backup administrator.
SQL Server Execute Package Utility	A graphical user interface that is used to run a Integration Services package.
SQL Server Express	An edition of a Microsoft relational database design and management system for e-commerce, line-of-business, and data warehousing solutions.
SQL Server instance	A copy of SQL Server running on a computer.
SQL Server instance auto-protection	A type of protection that enables DPM to automatically identify and protect databases that are added to instances of SQL Server that are configured for auto-protection.
SQL Server login	An account stored in SQL Server that allows users to connect to SQL Server.

Term	Definition
SQL Server Management Objects	An application programming interface that supports the incorporation of SQL Server administration into any COM or OLE Automation application.
SQL Server Management Studio	A suite of management tools included with Microsoft SQL Server for configuring, managing, and administering all components within Microsoft SQL Server.
SQL Server Master Data Services	A master data management application to consistently define and manage the critical data entities of an organization.
SQL Server Power Pivot for Excel	A SQL Server add-in for Excel.
SQL Server Profiler	A graphical user interface for monitoring an instance of the SQL Server database engine or an instance of Analysis Services.
SQL Server Reporting Services	A server-based report generation environment for enterprise, Web-enabled reporting functionality so you can create reports that draw content from a variety of data sources, publish reports in various formats, and centrally manage security and subscriptions.
SQL Server Service Broker	A technology that helps developers build scalable, secure database applications.
SQL Server Store for Office Applications	A SQL Server feature that enables storage of data from an Office application on a SQL server.
SQL Server Trace	A set Transact-SQL system stored procedures to create traces on an instance of the SQL Server Database Engine.
SQL Server Utility	A way to organize and monitor SQL Server resource health. It enables administrators to have a holistic view of their environment.
SQL Server Utility dashboard	A dashboard that provides an at-a-glance summary of resource health for managed SQL Server instances and data-tier applications. Can also be referred to as the SQL Server Utility detail view or the list view with details.
SQL Server Utility Explorer	A hierarchical tree displaying the objects in the SQL Server Utility.
SQL Server Utility viewpoints	A feature of SQL Server Utility that provides administrators a holistic view of resource health through an instance of SQL Server that serves as a utility control point (UCP).



<b>Term</b>	<b>Definition</b>
SQL statement	An SQL or Transact-SQL command, such as SELECT or DELETE, that performs some action on data.
SQL Trace	A set Transact-SQL system stored procedures to create traces on an instance of the SQL Server Database Engine.
SQL writer	A VSS compliant writer provided by the SQL Server that handles the VSS interaction with SQL Server.
SQL Writer Service	A service that permits Windows backup programs to copy SQL Server data files through the Volume Shadow Copy Service framework, while SQL Server is running.
SQL-92	The version of the SQL standard published in 1992.
SSAS	A feature of Microsoft SQL Server that supports online analytical processing (OLAP) and data mining for business intelligence applications. Analysis Services organizes data from a data warehouse into cubes with precalculated aggregation data to provide rapid answers to complex analytical queries.
SSL	The protocol that improves the security of data communication by using a combination of data encryption, digital certificates, and public key cryptography. SSL enables authentication and increases data integrity and privacy over networks. SSL does not provide authorization or nonrepudiation.
SSRCT	A tool that enables DPM administrators to authorize end users to perform self-service recovery of data by creating and managing DPM roles (grouping of users, objects, and permissions).
SSRS	A server-based report generation environment for enterprise, Web-enabled reporting functionality so you can create reports that draw content from a variety of data sources, publish reports in various formats, and centrally manage security and subscriptions.
SSRT	A tool that is used by end users to recover backups from DPM, without any action required from the DPM administrator.
staged data	Data imported into staging tables during the staging process in SQL Server Master Data Services.
staging	The process used in SQL Server Master Data Services to import data into staging tables and then process the staged data as a batch prior to importing it into the master database.
staging process	The process used in SQL Server Master Data Services to import data into staging tables and then process the staged data as a batch prior to importing it into the master database.
staging queue	The batch table in SQL Server Master Data Services where staged records are queued as batches to be processed into the Master Data Services database.
staging table	A table in SQL Server Master Data Services that is populated with business data during the staging process.
standalone server	A computer that runs Windows Server but does not participate in a domain. A standalone server has only its own database of end users, and it processes logon requests by itself. It does not share account information with any other computer and cannot provide access to domain accounts.

Term	Definition
standby file	In a restore operation, a file used during the undo phase to hold a "copy-on-write" pre-image of pages that are to be modified. The standby file allows reverting the undo pass to bring back the uncommitted transactions.
standby server	A server instance containing a copy of a database that can be brought online if the source copy of the database becomes unavailable. Log shipping can be used to maintain a "warm" standby server, called a secondary server, whose copy of the database is automatically updated from log backups at regular intervals. Before failover to a warm standby server, its copy of the database must be brought fully up to date manually. Database mirroring can be used to maintain a "hot" standby server, called a mirror server, whose copy of the database is continuously brought up to date. Failover to the database on a mirror server is essentially instantaneous.
standing query	An instantiation of a query template that runs within the StreamInsight server performing continuous computation over the incoming events.
star join	A join between a fact table (typically a large fact table) and at least two dimension tables.
star query	A query that joins a fact table and a number of dimension tables.
star schema	A relational database structure in which data is maintained in a single fact table at the center of the schema with additional dimension data stored in dimension tables. Each dimension table is directly related to and usually joined to the fact table by a key column.
start angle	The angle of rotation, between 0 and 360, at which the scale will begin. The zero (0) position is located at the bottom of the gauge, and the start angle rotates clockwise. For example, a start angle of 90 degrees starts the scale at the 9 o'clock position.
start cap	The start of a line.
statement	A compiled T-SQL query.
static cursor	A cursor that shows the result set exactly as it was at the time the cursor was opened.
static row filter	A filter available for all types of replication that allows you to restrict the data replicated to a Subscriber based on a WHERE clause.
stemmer	In Full-Text Search, for a given language, a stemmer generates inflectional forms of a particular word based on the rules of that language. Stemmers are language specific.
step into	To execute the current statement and enter Break mode, stepping into the next procedure whenever a call for another procedure is reached.
stewardship portal	A feature of MDS that provides centralized control over master data, including members and hierarchies and enables data model administrators to ensure data quality by developing, reviewing, and managing data models and enforcing them consistently across domains.
stock keeping unit	A unique identifier, usually alphanumeric, for a product. The SKU allows a product to be tracked for inventory purposes. An SKU can be associated with any item that can be purchased. For example, a shirt in style number 3726, size 8 might have a SKU of 3726-8.

Term	Definition
stolen page	A page in Buffer Cache taken for other server requests
stoplist	A specific collection of so-called stopwords, which tend to appear frequently in documents, but are believed to carry no usable information.
stopword	A word that tends to appear frequently in documents and carries no usable information.
storage engine	A component of SQL Server that is responsible for managing the raw physical data in your database. For example, reading and writing the data to disk is a task handled by the storage engine.
storage location	The position at which a particular item can be found: either an addressed location or a uniquely identified location on a disk, tape, or similar medium.
stored procedure	A precompiled collection of SQL statements and optional control-of-flow statements stored under a name and processed as a unit. They are stored in an SQL database and can be run with one call from an application.
stored procedure resolver	A program that is invoked to handle row change-based conflicts that are encountered in an article to which the resolver was registered.
stream	An abstraction of a sequence of bytes, such as a file, an I/O device, an inter-process communication pipe, a TCP/IP socket, or a spooled print job. The relationship between streams and storages in a compound file is similar to that of files and folders.
stream consumer	The structure or device that consumes the output of a query. Examples are an output adapter or another running query.
StreamInsight Event Flow Debugger	A stand-alone tool in the Microsoft StreamInsight platform that provides event-flow debugging and analysis.
StreamInsight platform	The platform, consisting of the StreamInsight server, Event Flow Debugging tool, Visual Studio IDE, and other components, for the development of complex event processing applications.
StreamInsight server	The core engine and adapter framework components of Microsoft StreamInsight. The StreamInsight server can be used to process and analyze the event streams associated with a complex event processing application.
string	A group of characters or character bytes handled as a single entity. Computer programs use strings to store and transmit data and commands. Most programming languages consider strings (such as 2674:gstmn) as distinct from numeric values (such as 470924).
strip line	Horizontal or vertical ranges that set the background pattern of the chart in regular or custom intervals. You can use strip lines to improve readability for looking up individual values on the chart, highlight dates that occur at regular intervals, or highlight a specific key range.
stripe	Horizontal or vertical ranges that set the background pattern of the chart in regular or custom intervals. You can use strip lines to improve readability for looking up individual values on the chart, highlight dates that occur at regular intervals, or highlight a specific key range.

Term	Definition
striped media set	A media set that uses multiple devices, among which each backup is distributed.
strong consistency	A scenario where high availability is enabled and there is more than one copy of a cached object in the cache cluster. All copies of that object remain identical.
Structured Query Language	A database query and programming language widely used for accessing, querying, updating, and managing data in relational database systems.
subquery	A SELECT statement that contains one or more subqueries.
subreport	A report contained within another report.
subscribe	To request data from a Publisher.
subscriber	In Notification Services, the person or process to which notifications are delivered.
subscriber	In replication, a database instance that receives replicated data.
subscriber database	In replication, a database instance that receives replicated data.
subscribing server	A server running an instance of Analysis Services that stores a linked cube.
subscription	A request for a copy of a publication to be delivered to a subscriber.
subscription database	A database at the Subscriber that receives data and database objects published by a Publisher.
subscription event rule	A rule that processes information for event-driven subscriptions.
subscription scheduled rule	One or more Transact-SQL statements that process information for scheduled subscriptions.
subset	A selection of tables and the relationship lines between them that is part of a larger database diagram.
subtract	To perform the basic mathematical operation of deducting something from something else.
Support Count	A dynamic option that displays the number of rows in which the determinant column value determines the dependent column.
Support Percentage	A dynamic option that displays the percentage of rows in which the determinant column determines the dependent column.
surface area	The number of ways in which a piece of software can be attacked.

<b>Term</b>	<b>Definition</b>
suspect tape	A tape that has conflicting identification information, such as the barcode or the on-media identifier.
SVF	A function that returns a single value, such as a string, integer, or bit value.
SVG	An XML-based language for device-independent description of two-dimensional graphics. SVG images maintain their appearance when printed or when viewed with different screen sizes and resolutions. SVG is a recommendation of the World Wide Web Consortium (W3C).
sweep angle	The number of degrees, between 0 and 360 that the scale will sweep in a circle. A sweep angle of 360 degrees produces a scale that is a complete circle.
switch in table	The staging table the user wants to use to switch in their data. The staging table needs to be created before switching partitions with the Manage Partitions Wizard.
switch out table	The staging table the user wants to use for the partition to switch out of the current source table.
symmetric key	The cryptographic key used to both encrypt and decrypt protected content during publishing and consumption.
Sync Manager	A tool used to ensure that a file or directory on a client computer contains the same data as a matching file or directory on a server.
sync provider	A software component that allows a replica to synchronize its data with other replicas.
synchronizati on application	A software component, such as a personal information manager or music database, that hosts a synchronization session and invokes synchronization providers to synchronize disparate data stores.
synchronizati on community	A set of replicas that keep their data synchronized with one another.
synchronizati on manager	A tool used to ensure that a file or directory on a client computer contains the same data as a matching file or directory on a server.
Synchronizati on Manager	A tool used to ensure that a file or directory on a client computer contains the same data as a matching file or directory on a server.
synchronizati on orchestrator	An orchestrator that initiates and controls synchronization sessions.
synchronizati on provider	A software component that allows a replica to synchronize its data with other replicas.
synchronizati on session	A unidirectional synchronization in which the source provider enumerates its changes and sends them to the destination provider, which applies them to its store.

Term	Definition
syndication format	A format used for publishing data on blogs and web sites.
syntactic validation	The process of confirming that an XML file conforms to its schema.
System Configuration Checker	A system preparation tool that helps to avoid setup failures by validating the target machine before a software application is installed.
system databases	A set of five databases present in all instances of SQL Server that are used to store system information.
system functions	A set of built-in functions that perform operations on and return the information about values, objects, and settings in SQL Server.
system locale	A Regional and Language Options setting that specifies the default code pages and associated bitmap font files for a specific computer that affects all of that computer's users. The default code pages and fonts enable a non-Unicode application written for one operating system language version to run correctly on another operating system language version.
system role assignment	Role assignment that applies to the site as a whole.
system role definition	Role definition that conveys site-wide authority.
system stored procedure	A type of stored procedure that supports all of the administrative tasks required to run a SQL Server system.
system stored procedures	A set of SQL Server-supplied stored procedures that can be used for actions such as retrieving information from the system catalog or performing administration tasks.
system table	A table that stores the data defining the configuration of a server and all its tables.
system tables	Built-in tables that form the system catalog for SQL Server.
system variable	A variable provided by DTS.
tab page	A part of a tab control that consists of the tab UI element and the display area, which acts as a container for data or other controls, such as text boxes, combo boxes, and command buttons.
table	A database object that stores data in records (rows) and fields (columns). The data is usually about a particular category of things, such as employees or orders.
table data region	A report item on a report layout that displays data in a columnar format.

<b>Term</b>	<b>Definition</b>
Table Designer	A visual design surface that is used to create and edit TSQL tables and table related objects.
table lock	A lock on a table including all data and indexes.
table reference	A name, expression or string that resolves to a table.
tablespace	A unit of database storage that is roughly equivalent to a file group in SQL Server. Tablespaces allow storage and management of database objects within individual groups.
table-valued function	A user-defined function that returns a table.
Tablix	A data region that can render data in table, matrix, and list format. It is intended to convey the unique functionality of the data region object and the users' ability to combine data formats.
Tablix data region	A data region that can render data in table, matrix, and list format. It is intended to convey the unique functionality of the data region object and the users' ability to combine data formats.
tabular data stream	The SQL Server internal client/server data transfer protocol. TDS allows client and server products to communicate regardless of operating-system platform, server release, or network transport.
tabular query	A standard operation such as search, sort, filter or transform on data in a table.
tail-log backup	A log backup taken from a possibly damaged database to capture the log that has not yet been backed up. A tail-log backup is taken after a failure in order to prevent work loss.
tape backup	A SQL Server backup operation that writes to any tape device supported by the operating system.
target	The database on which an operation acts.
target partition	An Analysis Services partition into which another is merged, and which contains the data of both partitions after the merger.
target queue	In Service Broker, the queue associated with the service to which messages are sent.
target server	A server that receives jobs from a master server.
target type	The type of target, which has certain characteristics and behavior.
task object	A Data Transformation Services (DTS) object that defines pieces of work to be performed as part of the data transformation process. For example, a task can execute an SQL statement or move and transform heterogeneous data from an OLE DB source to an OLE DB destination using the DTS Data Pump.
TDS	The SQL Server internal client/server data transfer protocol. TDS allows client and server products to communicate regardless of operating-system platform, server release, or network transport.

Term	Definition
temporary smart card	A non-permanent smart card issued to a user for replacement of a lost smart card or to a user that requires access for a limited time.
temporary stored procedure	A procedure placed in the temporary database, tempdb, and erased at the end of the session.
temporary table	A table placed in the temporary database, tempdb, and erased at the end of the session.
tenant	A client organization that is served from a single instance of an application by a web service. A company can install one instance of software on a set of servers and offer Software as a Service to multiple tenants.
theater view	A view where the preview is centered in a Power Pivot Gallery SharePoint document library and lets you rotate through the available worksheets. Smaller thumbnails of each worksheet appear lower on the page, on either side.
theta join	A join based on a comparison of scalar values.
thousand separator	A symbol that separates thousands from hundreds within a number that has four or more places to the left of the decimal separator.
thread	A type of object within a process that runs program instructions. Using multiple threads allows concurrent operations within a process and enables one process to run different parts of its program on different processors simultaneously. A thread has its own set of registers, its own kernel stack, a thread environment block, and a user stack in the address space of its process.
throttle	A Microsoft SQL Server tool designed to limit the performance of an instance of the database engine any time more than eight operations are active at the same time.
tick	A regular, rapidly recurring signal emitted by a clocking circuit.
tick count	A monotonically increasing number that is used to uniquely identify a change to an item in a replica.
tile server	A map image caching engine that caches and serves pregenerated, fixed-size map image tiles.
time	A SQL Server system data type that stores a time value from 0:00 through 23:59:59.999999.
time interval	A period of time in which a given event is valid. The valid time interval includes the valid start time, and all moments of time up to, but not including the valid end time.
tokenization	In text mining or Full-Text Search, the process of identifying meaningful units within strings, either at word boundaries, morphemes, or stems, so that related tokens can be grouped. For example, although "San Francisco" is two words, it could be treated as a single token.
tombstone	A marker that is used to represent and track a deleted item and prevent its accidental reintroduction into the synchronization community.
tool	A utility or feature that aids in accomplishing a task or set of tasks.



Term	Definition
topology	The set of participants involved in synchronization and the way in which they are connected to each other.
trace	A collection of events and data returned by the Database Engine.
trace file	A file containing records of activities of a specified object, such as an application, operating system, or network. A trace file can include calls made to APIs, the activities of APIs, the activities of communication links and internal flows, and other information.
tracer token	A performance monitoring tool available for transactional replication. A token (a small amount of data) is sent through the replication system to measure the amount of time it takes for transactions to reach the Distributor and Subscribers.
trail byte	The byte value that is the second half of a double-byte character.
train	To populate a model with data to derive patterns that can be used in prediction or knowledge discovery.
training data set	A set of known and predictable data used to train a data mining model.
trait	An attribute that describes an entity.
trait phrasing	A way of expressing a relationship in which a minor entity describes a major entity.
transaction isolation level	The property of a transaction that controls the degree to which data is isolated for use by one process, and is guarded against interference from other processes.
transaction log	A file that records transactional changes occurring in a database, providing a basis for updating a master file and establishing an audit trail.
transaction log backup	A backup of transaction logs that includes all log records not backed up in previous log backups. Log backups are required under the full and bulk-logged recovery models and are unavailable under the simple recovery model.
transaction retention period	In transactional replication, the amount of time transactions are stored in the distribution database.
transaction rollback	Rollback of a user-specified transaction to the last savepoint inside a transaction or to the beginning of a transaction.
transactional data	Data related to sales, deliveries, invoices, trouble tickets, claims, and other monetary and non-monetary interactions.
transactional replication	A type of replication that typically starts with a snapshot of the publication database objects and data.
Transact-SQL	The language containing the commands used to administer instances of SQL Server, create and manage all objects in an instance of SQL Server, and to insert, retrieve, modify and delete all data in SQL Server tables. Transact-SQL is an extension of the language defined in the SQL standards published by the International Standards Organization (ISO) and the American National Standards Institute (ANSI).

Term	Definition
transformation	The SSIS data flow component that modifies, summarizes, and cleans data.
transformation input	Data that is contained in a column, which is used during a join or lookup process, to modify or aggregate data in the table to which it is joined.
transformation output	Data that is returned as a result of a transformation procedure.
Trend	A general tendency or inclination, typically determined by the examination of a particular attribute over time.
trusted connection	A Windows network connection that can be opened only by users who have been authenticated by the network.
TSQL pane	One of the tabs that hosts the editor control to allow TSQL code editing.
tumbling window	A hopping window whose hop size is equal to the window size.
tuple	An ordered collection of members that uniquely identifies a cell, based on a combination of attribute members from every attribute hierarchy in the cube.
two-phase commit	A protocol that ensures that transactions that apply to more than one server are completed on all servers or none at all. Two-phase commit is coordinated by the transaction manager and supported by resource managers.
type checking	The process performed by a compiler or interpreter to make sure that when a variable is used, it is treated as having the same data type as it was declared to have.
typed adapter	An adapter that emits only a single event type.
typed event	An event for which the structure of the event payload provided by the source or consumed by the sink is known, and the input or output adapter is designed around this specific event structure.
UCP	A network node that provides the central reasoning point for the SQL Server Utility. It uses Utility Explorer in SQL Server Management Studio (SSMS) to organize and monitor SQL Server resource health.
UDT	A user-written extension to the scalar type system in SQL Server.
unbalanced hierarchy	A hierarchy in which one or more levels do not contain members in one or more branches of the hierarchy.
unbound stream	An event stream that contains the definition of the event model or payload type, but does not define the data source.
uncommittable	Pertaining to a transaction that remains open and cannot be completed. Uncommittable transactions could be considered a subclass of partially failed transactions, where the transaction has encountered an error that prevents its completion, but it is still holding its locks and has to be rolled back by the user.

Term	Definition
uncompress	To restore the contents of a compressed file to its original form.
undeliverable	Not able to be delivered to an intended recipient. If an e-mail message is undeliverable, it is returned to the sender with information added by the mail server explaining the problem; for example, the e-mail address may be incorrect, or the recipient's mailbox may be full.
underlying table	A table referenced by a view, cursor, or stored procedure.
undo	The phase during database recovery that reverses (rolls back) changes made by any transactions that were uncommitted when the redo phase of recovery completed.
undo file	A file that saves the content of the pages in a database after they've been modified by uncommitted, rolled back transactions and before recovery restores them to their previous state. The undo file prevents the changes performed by uncommitted transactions from being lost.
undo phase	The phase during database recovery that reverses (rolls back) changes made by any transactions that were uncommitted when the redo phase of recovery completed.
unenforced relationship	A link between tables that references the primary key in one table to a foreign key in another table, and which does not check the referential integrity during INSERT and UPDATE transactions.
uninitialize	To change the state of an enumerator or data source object so that it cannot be used to access data.
unique index	An index in which no two rows are permitted to have the same index value, thus prohibiting duplicate index or key values.
uniqueifier	A 4-byte column that the SQL Server Database Engine automatically adds to a row to make each index key unique.
Universal Time Coordinate	The standard time common to every place in the world, coordinated by the International Bureau of Weights and Measures. Coordinated Universal Time is used for the synchronization of computers on the Internet.
unknown member	A member of a dimension for which no key is found during processing of a cube that contains the dimension.
unknown tape	Tape that has not been identified by the DPM server.
unmanaged code	Code that is executed directly by the operating system, outside the .NET Framework common language runtime. Unmanaged code must provide its own memory management, type checking, and security support, unlike managed code, which receives these services from the common language runtime.
unmanaged instance	An instance of SQL Server not monitored by a utility control point.
unpivot	To expand values from multiple columns in a single record into multiple records with the same values in a single column.

Term	Definition
unsafe code	Code that is executed directly by the operating system, outside the .NET Framework common language runtime. Unmanaged code must provide its own memory management, type checking, and security support, unlike managed code, which receives these services from the common language runtime.
untyped adapter	An adapter that accepts or emits multiple event types in which the payload structure or the type of fields in the payload are not known in advance. Examples are events from a CSV or text file, a SQL table, or a socket.
update lock	A lock placed on resources (such as row, page, table) that can be updated.
update statistics	A process that recalculates information about the distribution of key values in specified indexes.
updategram	A template that makes it possible to modify a database in Microsoft SQL Server from an existing XML document.
user account	In Active Directory, an object that consists of all the information that defines a domain user, which includes user name, password, and groups in which the user account has membership. User accounts can be stored in either Active Directory or on your local computer.
user database	A database created by a SQL Server user and used to store application data.
user instance	An instance of SQL Server Express that is generated by the parent instance on behalf of a user.
user-defined aggregate function	An aggregate function created against a SQL Server assembly whose implementation is defined in an assembly created in the .NET Frameworks common language runtime.
user-defined type	A user-written extension to the scalar type system in SQL Server.
utility control point	A network node that provides the central reasoning point for the SQL Server Utility. It uses Utility Explorer in SQL Server Management Studio (SSMS) to organize and monitor SQL Server resource health.
Utility Reader	A privilege that allows the user account to connect to the SQL Server Utility, see all viewpoints in the Utility Explorer in SSMS and see settings on the Utility Administration node in Utility Explorer in SSMS.
vacuummer	A tool for data removal.
validity period	The amount of time a defined credential is deemed to be trusted.
value expression	An expression in Multidimensional Expressions (MDX) that returns a value. Value expressions can operate on sets, tuples, members, levels, numbers, or strings.
version	A property that is used to differentiate objects stored in the cache using the same key. Windows Server AppFabric stores the version information using the DataCacheItemVersion class. Every time an object is added or updated in the cache, the version value changes. Versioning is used to maintain data consistency. Optimistic concurrency is achieved by using versioning as opposed to locks.

<b>Term</b>	<b>Definition</b>
version	Metadata that identifies a change made to an item in a replica. It consists of the replica key and the replica tick count for the item.
vertical filtering	Filtering columns from a table. When used as part of replication, the table article created contains only selected columns from the publishing table.
vertical partitioning	The process of splitting a single table into multiple tables based on selected columns. Each of the multiple tables has the same number of rows but fewer columns.
vertical split	A vertical orientation of the CIDER shell.
very large database	A database that has become large enough to be a management challenge, requiring extra attention to people, processes, and processes.
victim	The longest running transaction that has not generated row versions when tempdb runs out of space and the Database Engine forces the version stores to shrink. A message 3967 is generated in the error log for each victim transaction. If a transaction is marked as a victim, it can no longer read the row versions in the version store.
view generation	A repository engine feature that is used to create relational views based on classes, interfaces, and relationships in an information model.
virtual log file	Non-physical files that are derived from one physical log file by the SQL Server Database Engine.
visual total	A displayed, aggregated cell value for a dimension member that is consistent with the displayed cell values for its displayed children.
visualizer	A way to visually represent data in debug mode.
VLDB	A database that has become large enough to be a management challenge, requiring extra attention to people, processes, and processes.
VSS writer	A component within an application that interfaces with the VSS platform infrastructure during backups to ensure that application data is ready for shadow copy creation.
wall-time	The total time taken by a computer to complete a task which is the sum of CPU time, I/O time, and the communication channel delay.
warm latency	The time to create a new workflow instance when the workflow type has already been compiled.
warm standby	A method of redundancy in which the secondary (i.e., backup) system runs in the background of the primary system. Data is mirrored to the secondary server at regular intervals, which means that there are times when both servers do not contain the exact same data.
warm standby server	A standby server that contains a copy of a database that is asynchronously updated, and that can be brought online fairly quickly.
watermark	A threshold used to manage the memory consumption on each cache host. The high watermark and low watermark specify when objects are evicted out of memory.

Term	Definition
Web application	A software program that uses Hypertext Transfer Protocol (HTTP) for its core communication protocol and that delivers Web-based information to the user in the Hypertext Markup Language (HTML) language.
Web pool	A grouping of one or more URLs served by a worker process.
Web Pool Agent	An isolated process under which the Certificate Lifecycle Manager (CLM) web portal runs.
Web project	A collection of files that specifies elements of a Web application.
Web site	A group of related web pages that is hosted by an HTTP server on the World Wide Web or an intranet. The pages in a website typically cover one or more topics and are interconnected through hyperlinks.
Web synchronization	In merge replication, a feature that lets you replicate data by using the HTTPS protocol.
website	A group of related web pages that is hosted by an HTTP server on the World Wide Web or an intranet. The pages in a website typically cover one or more topics and are interconnected through hyperlinks.
weighted close formula	A formula that calculates the average of the high, low, and close prices, while giving extra weight to the close price.
wide character	A 2-byte multilingual character code.
window	A subset of events within a stream that fall within some period of time; that is, a window contains event data along a timeline.
Windows Management Instrumentation	The Microsoft extension to the Distributed Management Task Force (DMTF) Web-based Enterprise Management (WBEM) initiative.
Windows NT Integrated Security	A security mode that leverages the Windows NT authentication process.
witness server	In database mirroring, the server instance that monitors the status of the principal and mirror servers and that, by default, can initiate automatic failover if the principal server fails. A database mirroring session can have only one witness server (or "witness"), which is optional.
WMI	The Microsoft extension to the Distributed Management Task Force (DMTF) Web-based Enterprise Management (WBEM) initiative.
WMI Query Language	A subset of ANSI SQL with semantic changes adapted to Windows Management Instrumentation (WMI).
workbook	In a spreadsheet program, a file containing a number of related worksheets.

Term	Definition
workload governor	A Microsoft SQL Server tool designed to limit the performance of an instance of the database engine any time more than eight operations are active at the same time.
workload group	In Resource Governor, a container for session requests that are similar according to the classification rules that are applied to each request. A workload group allows the aggregate monitoring of resource consumption and a uniform policy that is applied to all the requests in a group.
workstation	A microcomputer or terminal connected to a network.
WQL	A subset of ANSI SQL with semantic changes adapted to Windows Management Instrumentation (WMI).
write back	To update a cube cell value, member, or member property value.
write-ahead log	A transaction logging method in which the log is always written prior to the data.
x-axis	The horizontal reference line on a grid, chart, or graph that has horizontal and vertical dimensions.
XML for Analysis	A specification that describes an open standard that supports data access to data sources that reside on the World Wide Web.
XMLA	A specification that describes an open standard that supports data access to data sources that reside on the World Wide Web.
XQuery	Functional query language that is broadly applicable to a variety of XML data types derived from Quilt, XPath, and XQL. Both Ipedo and Software AG implement their own versions of the W3C's proposed specification for the XQuery language. Also called: XML Query, XQL.
XSL	An XML vocabulary that is used to transform XML data to another form, such as HTML, by means of a style sheet that defines presentation rules.
XSL Transformation	A declarative, XML-based language that is used to present or transform XML data.
XSLT	A declarative, XML-based language that is used to present or transform XML data.

---

## Community Additions