

24-7 Operation of an application or database 24 hours a day, 7 days a week. Because the database can never be shut down, performing maintenance is a challenge.

Abstract data types In SQL 1999 the ability to define more complex data domains that support inheritance for storing objects.

Accessibility A design goal to make the application usable by as many users as possible, including those with physical challenges. One solution is to support multiple input and output methods.

ACID transactions The acronym for transactions that specifies the four required elements of a safe transaction: atomicity, consistency, isolation, and durability.

Active data objects (ADO) Microsoft's component (COM) approach to connect program code and Web server scripts to a database. Provides SQL statement and row-level access to virtually any database.

Active server pages (ASP) Microsoft's Web pages that enable you to run script programs on the server. Useful for providing access to a server database for Internet users.

Address, physical The location of data stored in memory or on a file system. Used to establish a pointer to a specific piece of data.

ad hoc query An unplanned question asked by users or administrators. Because it is unplanned, it is hard to support with non-relational databases.

Administrative tasks Jobs that need to be performed to keep the application running, such as updating data in lookup tables, backing up the database, and assigning users to groups

Advanced Encryption Standard (AES) A single-key encryption system to replace DES, based on a Belgian encryption system: Rijndael. It supports key lengths of 128, 192 and 256 bits, making it considerably more secure than DES.

Aesthetics An application design goal, where layout, colors, and artwork are used to improve the appearance of the application—not detract from it. By its nature, the value of any design is subjective.

Aggregation The generic name for several SQL functions that operate across the selected rows. Common examples include SUM, COUNT, and AVERAGE.

Aggregation association A relationship where individual items become elements in a new class. For example, an Order contains Items. In UML, the association is indicated with a small open diamond on the association end. *See composition.*

Alias A temporary name for a table or a column. Often used when you need to refer to the same table more than once, as in a self-join.

ALL A SQL SELECT clause often used with subqueries. Used in a WHERE clause to match all of the items in a list. For example, Price > ALL (...) means that the row matches only if Price is greater than the largest value in the list.

ALLOW FILTERING Cassandra query option required for supporting most WHERE clause options including AND, OR, and some range questions. Might reduce query performance so it is a warning that complex queries should be avoided.

ALTER TABLE A SQL data definition command that changes the structure of a table. To improve performance, some systems limit the changes to adding new columns. In these situations to make major changes, you have to create a new table and copy the old data.

Anchor tag The HTML tag that signifies a link. Denoted with <A>.

ANY A SQL SELECT clause often used with subqueries. Used in a WHERE clause to match at least one of the items in a list. For example, Price > ANY (...) means that the row matches as long as Price is greater than at least one item in the list.

Application A complete system that performs a specific collection of tasks. It typically consists of integrated forms and reports and generally contains menus and a Help system.

Application Design Guide A standard set of design principles that should be followed when building applications. The standard makes it easier for users to operate new applications, since techniques they learn in one system will work in another.

Application generator A DBMS tool that assists the developer in creating a complete application package. Common tools include menu and toolbar generators and an integrated context-sensitive Help system

Association Connections between classes or entities. Generally, they represent business rules. For example, an order can be placed by one customer. It is important to identify whether the association is one-to-one, one-to-many, or many-to-many.

Association role In UML the point where an association attaches to a class. It can be named, and generally shows multiplicity, aggregation, or composition.

Association rules A data mining technique that examines a set of transactions to see which items are commonly purchased together.

Atomic The smallest single-valued form of a data element. A table cannot be in first normal form if the cells contain non-atomic data. The definition is subjective depending on the application. For instance, an address such as 123 Main Street is usually considered to be atomic even though it refers to both a house number and a street.

Atomicity The transaction element that specifies that all changes in a transaction must succeed or fail together.

Attribute A feature or characteristic of an entity. An attribute will become a column in a data table. Employee attributes might include name, address, date hired, and phone.

Authentication Providing a verification system to determine who actually wrote a message. Common systems use a dual-key encryption system.

Autonumber A type of data domain where the DBMS automatically assigns a unique identification number for each new row of data. Useful for generating primary keys.

B+tree An indexed data storage method that is efficient for a wide range of data access tasks. Tree searches provide a consistent level of performance that is not affected by the size of the database.

Back end In a client-server system, the back end usually consists of a central database. In general, hardware and data placed at the back

end is designed to be centralized and shared. *See front end.*

Base table A table that contains data about a single basic entity. It generally contains no foreign keys, so data can be entered into this table without reference to other tables. For example, Customer would be a base table; Order would not.

Behavioral security Emphasizes the role of people or employees in security. It is related to logical security but involves interesting problems because it deals with mistakes that people make, such as tricking people to reveal passwords.

BETWEEN A SQL comparison operator that determines whether an item falls between two values. Often useful for dates.

Binary large object (BLOB) A data domain for undefined, large chunks of data. A BLOB (or simple object) type can hold any type of data, but the programmer is often responsible for displaying, manipulating, and searching the data.

Binary search A search technique for sorted data. Start at the middle of the data. If the search value is greater than the middle value, split the following data in half. Keep reducing by half until the value is found.

Bitmap index A compact, high-speed indexing method where the key values and conditions are compressed to a small size that can be stored and searched rapidly.

Boolean algebra Creating and manipulating logic queries connected with AND, OR, and NOT conditions.

Bound control A control on a form that is tied to a column in the database. When data is entered or changed, the changes are automatically saved to the data table.

Boyce-Codd normal form (BCNF) All dependencies must be explicitly shown through keys. There cannot be a hidden dependency between nonkey and key columns.

Browser, Web A software package on a client personal computer used to access and display Web pages from the Internet.

Brute force attack An attempt to break a security system by trying every possible

combination of passwords or encryption keys.

Business intelligence (BI) The general process of analyzing data to find patterns. Tools include automated statistical systems and user-driven exploratory options. Also called data mining.

Business rules The conditions and assumptions that describe how an organization operates. One-to-one and one-to-many rules are particularly important. For instance, a common business rule is that a sale is placed by only one customer.

Call-level interface (CLI) A set of libraries that enable programmers to work in a language outside the DBMS (e.g., C++) and utilize the features of the DBMS. The DBMS provides the communication libraries and handles much of the data exchange itself.

Cascading delete When tables are linked by data, if you delete a row in a higher level table, matching rows in other tables are deleted automatically. For example, if you delete Customer 1173, all orders placed by that customer are also deleted.

Cascading style sheets (CSS) A style sheet that defines how elements are to be displayed on a Web page. Cascading means that a style can be overridden by declaring it inside the HTML file, but to maintain consistency, styles should be defined only in the style sheet.

Cascading triggers Multiple events that arise when a change that fires a trigger on one table causes a change in a second table, that triggers a change in a third table and so on.

CASE A SQL operator supported by some systems. It examines multiple conditions (cases) and takes the appropriate action when it finds a match.

Cassandra Query Language (CQL) An interactive query system for retrieving data from a Cassandra database. It includes commands similar in syntax to SQL but with limited options. In particular the SELECT command does not support JOINS.

Certificate authority A company that ensures the validity of public keys and the applicant's identify for dual-key encryption systems.

Check box A square button that signifies a choice. By the design guide, users can

select multiple options with check boxes, as opposed to option buttons that signify mutually exclusive choices.

Clarity The goal of making an application easier to use through elegant design and organization that matches user tasks so that the purpose and use of the application is clear to the user.

Class A descriptor for a set of objects with similar structure, behavior, and relationships. That is, a class is the model description of the business entity. A business model might have an Employee class, where one specific employee is an object in that class.

Class diagram A graph of classes connected through relationships. It is designed to show the static structure of the model. Similar to the entity-relationship diagram.

Class hierarchy A graph that highlights the inheritance relationships between classes.

Classification analysis A data mining technique that classifies groups of objects, such as customers. It determines which factors are important classifier variables.

Client/server A technique for organizing systems where a few computers hold most of the data, which is retrieved by individuals using personal computer clients.

Cloud computing Using servers that are connected to the Internet. Typically the servers are leased as virtual machines that can be expanded as needed. Users rely on Web-based clients to run applications and retrieve data.

Cluster, data A physical data storage technique to improve performance by storing related data in the same data blocks so that the operating system retrieves the related data in one pass.

Cluster, server A collection of computer servers that share the workload. If one machine fails, the others pick up the load. New servers can be added at any time to improve performance.

Cluster analysis A data mining technique that groups elements of a dataset, often based on how close the items are to each other.

Cold site A facility that can be leased from a disaster backup specialist. A cold site contains

power and telecommunication lines, but no computer. In the event of a disaster, a company calls the computer vendor and begs for the first available machine to be sent to the cold site.

Collaboration diagram A UML diagram to display interactions among objects. It does not show time as a separate dimension. It is used to model processes.

Column A column in a table represents an attribute or measure on the object defined by the table. Columns have a defined data type. A simple example would be a Phone Number in a Customer table. In a relational DBMS, columns hold a single value for each row. In a non-relational DBMS, columns may contain repeating or complex data.

Column family The early term for a table in the Cassandra non-relational DBMS. The term still arises in some documentation and error messages.

Combo box A combination of a list box and a text box that is used to enter new data or to select from a list of items. A combo box saves space compared to a list box since the list is displayed only when selected by the user. Known as a select box on Web forms.

Comma separated values (CSV) A method of storing data for transfer to different computers or applications. Tabular data is stored in rows with the columns separated by commas. The data is stored in a simple text file.

Command button A button on a form that is designed to be clicked. The designer writes the code that is activated when the button is clicked.

Common gateway interface (CGI) With Web servers, CGI is a predefined system for transferring data across the Internet. Current scripting languages hide the details, so you can simply retrieve data as it is needed.

Common language runtime (CLR) Microsoft's base programming language. It is embedded into versions of SQL Server from 2005, so database procedures can be written in CLR languages such as Visual Basic and C# including access to all of their functions. Needed to use RegEx within SQL Server. Note that CLR support is turned off by default.

Composite key A primary key that consists of more than one column. Indicates a many-to-

many relationship between the columns.

Composite primary key In non-relational DBMSs, particularly Cassandra, a composite key uses two or more columns to define the partitioning key so both values are required to retrieve the associated data. It is defined with an extra set of parentheses: PRIMARY KEY (a, b), others).

Composition association A relationship in which an object is composed of a collection of other objects. For example, a bicycle is built from components. In UML, it is indicated with a small filled diamond on the association end.

Compound primary key In non-relational DBMSs, particularly Cassandra, a compound key defines a many-to-many relationship across multiple columns. But unlike the composite key, it uses only the first column to partition the data so all elements of the many side are retrieved by specifying only the value for the first key column. The definition does not include extra parentheses: PRIMARY KEY (a, b, c, ...).

Computer-aided software engineering (CASE) Computer programs that are designed to support the analysis and development of computer systems. They make it easier to create, store, and share diagrams and data definitions. Some versions can analyze existing code and generate new code.

Concatenate A programming operation that appends one string on the end of a second string. For example, LastName & " " & FirstName could yield "Smith, John".

Concatenated key *See composite key.*

Concurrent access Performing two (or more) operations on the same data at the same time. The DBMS must sequence the operations so that some of the changes are not lost.

Confidence In data mining with association rules, a measure of the strength of a rule measured by the percentage of transactions with item A that also contain item B. The probability that B is in the basket given that A is already there.

Consistency, application The goal of making an application easier to use by using the same features, colors, and commands throughout. Modern applications also strive for consistency with a common design guide.

Consistency, transaction The transaction requirement that specifies all data must remain internally consistent when changes are committed and can be validated by application checks.

Constraint In SQL, a constraint is a rule that is enforced on the data. For example, there can be primary-key and foreign-key constraints that limit the data that can be entered into the declared columns. Other business rules can form constraints, such as $\text{Price} > 0$.

Context-sensitive help Help messages that are tailored to the specific task the user is performing.

Context sensitive menu A menu that changes depending on the object selected by the user.

Control break A report consisting of grouped data uses control breaks to separate the groups. The break is defined on the key variable that identifies each member of the group.

Controls The generic term for an item placed on a form. Typical controls consist of text boxes, combo boxes, and labels.

Correlated subquery A subquery that must be reevaluated for each row of the main query. Can be slow on some systems. Can often be avoided by creating a temporary table and using that in the subquery instead.

CREATE DOMAIN A SQL data definition command to create a new data domain that is composed of existing domain types.

CREATE SCHEMA A SQL data definition command to create a new logical grouping of tables. With some systems it is equivalent to creating a new database. This command is not available in Oracle or SQL Server.

CREATE TABLE A SQL data definition command to create a new table. The command is often generated with a program.

CREATE VIEW A SQL command to create a new view or saved query.

Cross join Arises when you do not specify a join condition for two tables. It matches every row in the first table with every row in the second table. Also known as the Cartesian product. It should be avoided.

Crosstab A special SQL query (not offered by all systems) that creates a tabular output

based on two groups of data. Access uses a TRANSFORM command to create a cross tabulation.

Cursor, graphics The current location pointer in a graphical environment.

Cursor, database A row pointer that tracks through a table, making one row of data active at a time.

Cylinder Disk drives are partitioned into cylinders (or sectors) that represent a portion of a track.

Data administration Planning and coordination required to define data consistently throughout the company.

Data administrator (DA) The person in charge of the data resources of a company. The DA is responsible for data integrity, consistency, and integration.

Data bound A control that contains a link to a table and column data source. When a data-bound control is displayed, the data is retrieved from the database and shown in the control. Changes are written to the defined column and table.

Data definition language (DDL) A set of commands that are used to define data, such as CREATE TABLE. Graphical interfaces are often easier to use, but the data definition commands are useful for creating new tables with a program.

Data device Storage space allocated to hold database tables, indexes, and rollback data. *See tablespace.*

Data dictionary Holds the definitions of all of the data tables and describes the type of data that is being stored.

Data hierarchy A structured ordering of data where higher levels contain aggregates of the lower levels. Some common hierarchies include dates (year, quarter, month, day), and geography (country, state, city).

Data independence Separates the data from the programs, which often enables the data definition to be changed without altering the program.

Data integrity Keeping accurate data, which means few errors and means that the data reflects the true state of the business. A DBMS

enables you to specify constraints or rules that help maintain integrity, such as prices must always be greater than 0.

Data manipulation language (DML) A set of commands used to alter the data. *See INSERT, DELETE, and UPDATE.*

Data mining Searching databases for unknown patterns and information. Tools include statistical analysis, pattern-matching techniques, and data segmentation analysis, classification analysis, association rules, and cluster analysis.

Data normalization The process of creating a well-behaved set of tables to efficiently store data, minimize redundancy, and ensure data integrity. *See first, second, and third normal form.*

Data replication In a distributed system, placing duplicate copies of data on several servers to reduce overall transmission time and costs.

Data repository A complete listing of all terms used in a database design, including column names and tables. *See data dictionary.*

Data type A type of data that can be held by a column. Each DBMS has predefined system domains (integer, float, string, etc.). Some systems support user-defined domains that are named combinations of other data types.

Data volume The estimated size of the database. Computed for each table by multiplying the estimated number of rows times the average data length of each row.

Data warehouse A specialized database that is optimized for management queries. Data is extracted from online transaction processing systems. The data is cleaned and optimized for searching and analysis. Generally supported by parallel processing and RAID storage.

Database A collection of data stored in a standardized format, designed to be shared by multiple users. A collection of tables for a particular business situation.

Database administration The technical aspects of creating and running the database. The basic tasks are performance monitoring, backup and recovery, and assigning and controlling security.

Database administrator (DBA) A specialist who is trained in the administration of a particular DBMS. DBAs are trained in the details of installing, configuring, and operating the DBMS.

Database cursor A variable created within a programming language that defines a SELECT statement and points to one row of data at a time. Data on that row can be retrieved or edited using the programming language.

Database engine The heart of the DBMS. It is responsible for storing, retrieving, and updating the data.

Database management system (DBMS) Software that defines a database, stores the data, supports a query language, produces reports, and creates data entry screens.

Datasheet A gridlike form that displays rows and columns of data. Generally used as a subform, a datasheet displays data in the least amount of space possible.

Deadlock A situation that exists when two (or more) processes each have a lock on a piece of data that the other one needs.

Default values Values that are displayed and entered automatically. Used to save time at data entry.

Degree, tree The maximum number of children allowed beneath one node in a B-tree. Most systems choose an odd number greater than or equal to three.

DELETE A SQL data manipulation command that deletes rows of data. It is always used with a WHERE clause to specify which rows should be deleted.

Deletion anomaly Problems that arise when you delete data from a table that is not in third normal form. For example, if all customer data is stored with each order, when you delete an order, you could lose all associated customer data.

DeMorgan's law An algebraic law that states: To negate a condition that contains an AND or an OR connector, you negate each of the two clauses and switch the connector. An AND becomes an OR and vice versa.

DENSE_RANK A data mining extension to SQL that rank orders data. If ties exist, they

receive the same dense rank, but the next value receives a dense rank that is one unit higher. Compare to RANK.

Dependence An issue in data normalization. An attribute A depends on another attribute B if the values of A change in response to changes in B. For example, a customer's name depends on the CustomerID (each employee has a specific name). On the other hand, a customer's name does not depend on the OrderID. Customers do not change their names each time they place an order.

Depth, tree The number of levels in a B-tree or the number of nodes between the root and the leaves.

Derived class A class that is created as an extension of another class. The programmer need only define the new attributes and methods. All others are inherited from the higher-level classes. *See inheritance.*

DESC The modifier in the SQL SELECT ... ORDER BY statement that specifies a descending sort (e.g., Z ... A). ASC can be used for ascending, but it is the default, so it is not necessary.

Dimension An attribute in an OLAP cube that is used to group and search the data.

Direct access A data storage method where the physical location is computed from the logical key value. Data can be stored and retrieved with no searches.

Direct manipulation of objects A graphical interface method that is designed to mimic real-world actions. For example, you can copy files by dragging an icon from one location to another.

Disaster plan A contingency plan that is created and followed if a disaster strikes the computer system. Plans include off-site storage of backups, notifying personnel, and establishing operations at a safe site.

DISTINCT An SQL keyword used in the SELECT statement to remove duplicate rows from the output.

Distributed database Multiple independent databases that operate on two or more computers that are connected and share data over a network. The databases are usually in different physical locations. Each database is

controlled by an independent DBMS.

Dockable toolbar A toolbar that users can drag to any location on the application window. It is generally customized with options and buttons to perform specific tasks.

Domain-key normal form (DKNF) The ultimate goal in designing a database. Each table represents one topic, and all of the business rules are expressed in terms of domain constraints and key relationships. That is, all of the business rules are explicitly described by the table rules.

Drag-and-drop A graphical interface technique where actions are defined by holding down a mouse key, dragging an icon, and dropping the icon on a new object.

Drill down The act of moving from a display of summary data to more detail. Commonly used in examining data in a data warehouse or OLAP application. *See Roll up.*

Drive head The mechanism that reads and writes data onto a disk. Modern drives have several drive heads.

DROP TABLE A SQL data definition command that completely removes a table from the database—including the definition. Use it sparingly.

Dual-key encryption An encryption technique that uses two different keys: one private and one public. The public key is published so anyone can retrieve it. To send an encrypted message to someone, you use the person's public key. At that point, only the person's private key will decrypt the message. Encrypting a message first with your private key can also be used to verify that you wrote the message.

Durability The transaction element that specifies that when a transaction is committed, all changes are permanently saved even if there is a hardware or system failure.

Edit The Microsoft DAO command to alter data on the current row.

Electronic data interchange (EDI) Exchanging data over networks with external agents such as suppliers, customers, and banks.

Encapsulation In object-oriented programming, the technique of defining

attributes and methods within a common class. For example, all features and capabilities of an Employee class would be located together. Other code objects can use the properties and methods but only by referencing the Employee object.

Encryption Encoding data with a key value so the data becomes unreadable. Two general types of encryption are used today: single key (e.g., DES) and dual key (e.g., RSA).

Entity An item in the real world that we wish to identify and track.

Entity-relationship diagram (ERD) A graph that shows the associations (relationships) between business entities. Under UML, the class diagram displays similar relationships.

Equi-join A SQL equality join condition. Rows from two tables are joined if the columns match exactly. Equi-join is the most common join condition. Theta joins support inequality conditions.

Error handling Special programming code used to trap errors. The try/catch or On Error Goto syntaxes are common. The goal is to catch errors and handle them automatically without interfering with the user.

Event Something that arises during database or form operations. Events are named and developers can write code that is executed when a specific event is triggered.

EXCEPT A SQL operator that examines rows from two SELECT statements. It returns all rows from one statement except those that would be returned by the second statement. Sometimes implemented as a SUBTRACT command. *See UNION.*

EXISTS A SQL keyword used to determine if subqueries return any rows of data.

Expert system (ES) A system with a knowledge base consisting of data and rules that enables a novice to make decisions as effectively as an expert.

Extensible markup language (XML) A tag-based notation system that is used to assign names and structure to data. It was mainly designed for transferring data among diverse systems.

Extraction, transformation, and

transportation (ETT) The three steps in populating a data warehouse from existing files or databases. Extraction means selecting the data you want. Transformation is generally the most difficult step and requires making the data consistent. Transportation implies that the data has to be physically moved over a network to the data warehouse. Oracle refers to the topic as ETL (loading).

Fact table The table or query holding the facts to be presented in an OLAP cube.

Fault tolerance Various methods of building a system so that if something fails, other components pick up the load without losing the entire transaction or application.

Feasibility study A quick examination of the problems, goals, and expected costs of a proposed system. The objective is to determine whether the problem can reasonably be solved with a computer system.

Feedback A design feature where the application provides information to the user as tasks are accomplished or errors arise. Feedback can be provided in many forms (e.g., messages, visual cues, or audible reminders).

FETCH The command used in SQL cursor programming to retrieve the next row of data into memory.

First normal form (1NF) A table is in 1NF when there are no repeating groups within it. Each cell can contain only one value. For example, how many items can be placed in one Order table? The items repeat, so they must be split into a separate table.

Fixed-width storage Storing each row of data in a fixed number of bytes per column.

Fixed-with-overflow storage Storing a portion of the row data in a limited number of bytes, and moving extra data to an overflow location.

Focus In a window environment, a form or control has focus when it is the one that will receive keystrokes. It is usually highlighted.

For Each ... Next In VBA, an iteration command to automatically identify objects in a group and apply some operation to that collection. Particularly useful when dealing with cells in a spreadsheet.

Foreign key A column in one table that is a primary key in a second table. It does not need to be a key in the first table. For example, in an Order table, CustomerID is a foreign key because it is a primary key in the Customer table.

Forms development The process of designing and creating input forms to collect data and store it in the database.

Forms generator A DBMS tool that enables you to set up input forms on the screen.

Fourth normal form (4NF) There cannot be hidden dependencies between key columns. A multi-valued dependency exists when a key determines two separate but independent attributes. Split the table to make the two dependencies explicit.

FROM The SQL SELECT clause that signifies the tables from which the query will retrieve data. Used in conjunction with the JOIN or INNER JOIN statement.

Front end In a client-server or multi-tier system, the forms and applications that are displayed or run on the user's computer. The portion of the application seen and manipulated by the user. *See back end.*

FULL JOIN A join that matches all rows from both tables if they match, plus all rows from the left table that do not match, and all rows from the right table that do not match. Rarely used and rarely available. *See left join and right join.*

Function A procedure designed to perform a specific computation. The difference between a function and a subroutine is that a function returns a specific value (not including the parameters).

Generalization association A relationship among classes that begins with a generic class. More detailed classes are derived from it and inherit the properties and methods of the higher level classes.

Geocode Assigning location coordinates of latitude and longitude to a dataset.

Geographic information system (GIS) Designed to identify and display relationships among business data and locations. A good example of the use of objects in a database environment.

Globally-unique identifier (GUID) A large number that can be created on one computer and be different from all other numbers created. Often used for generated keys in a replicated database.

Gossip A small communication channel in the Cassandra DBMS used to monitor server node status and metadata across database replicas.

GRANT The SQL command to give someone access to specific tables or queries.

Graphics interchange file (GIF) One standard method of storing graphical images. Commonly used for images shared on the Internet.

Group break A report that splits data into groups. The split-point is called a break. Also known as a control break.

GROUP BY A SQL SELECT clause that computes an aggregate value for each item in a group. For example, SELECT Department, SUM(Salary) FROM Employee GROUP BY Department; computes and lists the total employee salaries for each department.

Hashed-key access *See direct access.* Hash refers to the function used to reduce a key value to a numbered location—usually modulo division by a prime number.

HAVING A SQL clause used with the GROUP BY statement. It restricts the output to only those groups that meet the specified condition.

Heads-down data entry Touch typists concentrate on entering data without looking at the screen. Forms for this task should minimize keystrokes and use audio cues.

Help system A method for displaying, sequencing, and searching help documentation. Developers need to write the help files in a specific format and then use a help compiler to generate the final help file.

Hidden dependency A dependency specified by business rules that is not shown in the table structure. It generally indicates that the table needs to be normalized further and is an issue with Boyce-Codd or fourth normal form.

Hierarchical database An older DBMS type that organizes data in hierarchies that can be rapidly searched from top to bottom, e.g.,

Customer – Order – OrderItem.

Horizontal partition Splitting a table into groups based on the rows of data. Rows that are seldom used can be moved to slower, cheaper storage devices.

Hot site A facility that can be leased from a disaster backup specialist. A hot site contains all the power, telecommunication facilities, and computers necessary to run a company. In the event of a disaster, a company collects its backup data, notifies workers, and moves operations to the hot site.

Human factors design An attempt to design computer systems that best accommodate human users.

Hypertext link Hypertext (e.g., Web) documents consist of text and graphics with links that retrieve new pages. Clicking on a link is the primary means of navigation and obtaining more information.

Hypertext markup language (HTML) A display standard that is used to create documents to be shared on the Internet. Several generators will create HTML documents from standard word processor files.

Icon A small graphical representation of some idea or object. Typically used in a graphical user interface to execute commands and manipulate underlying objects.

IN A SQL WHERE clause operator typically used with subqueries. It returns a match if the selected item matches one of the items in the list. For example, WHERE ItemID IN (115, 235, 536) returns a match for any of the items specified. Typically, another SELECT statement is inserted in the parentheses.

Index A sorted list of key values from the original table along with a pointer to the rest of the data in each row. Used to speed up searches and data retrieval.

Indexed sequential access method (ISAM) A data storage method that relies on an index to search and retrieve data faster than a pure sequential search.

Inequality join A SQL join where the comparison is made with an inequality (greater than or less than) instead of an equality operator. Useful for placing data into categories based on ranges of data. The general form is

also known as a theta join.

Inheritance In object-oriented design, the ability to define new classes that are derived from higher-level classes. New classes inherit all prior properties and methods, so the programmer only needs to define new properties and methods.

INNER JOIN A SQL equality join condition. Rows from two tables are joined if the columns match exactly. The most common join condition. Rows that have no match in the other table are not displayed.

Input mask A special string that defines how data can be entered. Used to control the way users enter values into text boxes, such as requiring numeric values for currency items.

InputBox A predefined simplistic Window form that might be used to get one piece of data from the user. But it is better to avoid it and create your own form.

INSERT Two SQL commands that insert data into a table. One version inserts a single row at a time. The other variation copies selected data from one query and appends it as new rows in a different table.

Insertion anomaly Problems that arise when you try to insert data into a table that is not in third normal form. For example, if you find yourself repeatedly entering the same data (e.g., a customer's address), the table probably needs to be redefined.

Internet A collection of computers loosely connected to exchange information worldwide. Owners of the computers make files and information available to other users.

INTERSECT A set operation on rows of data from two SELECT statements. Only rows that are in both statements will be retrieved. *See UNION.*

Intranet A network internal to a company that uses Internet technologies to share data.

Isolation The transaction requirement that says the system must give each transaction the perception that it is running in isolation with no concurrent access issues.

Isolation level Used to assign locking properties in transactions. At a minimum, it is used to specify optimistic or pessimistic locks.

Some systems support intermediate levels.

Iteration Causing a section of code to be executed repeatedly, such as the need for a loop to track through each row of data. Typical commands include Do ... Loop, and For ... Next.

Java A programming language developed by Sun Microsystems that is supposed to be able to run unchanged on diverse computers. Originally designed as a control language for embedded systems, Java is targeted for Internet applications. The source of many bad puns in naming software products.

Java 2 enterprise edition (J2EE) A back-end server-based system for building complex applications. It is based on Java but consists of an entire environment.

JDBC A set of methods to connect Java code to databases. Similar in purpose to ADO, but works only in Java. Sometimes referred to as Java Database Connectivity.

JOIN When data is retrieved from more than one table, the tables must be joined by some column or columns of data. *See INNER JOIN and LEFT JOIN.*

Keyspace The top level namespace in the Cassandra non-relational DBMS. It functions like a schema in SQL to separate tables by application. Table names must be unique within a keyspace but different keyspaces can have table names that exist in other keyspaces.

Key-value pair A common method of identifying and transferring data in Web applications that is used for storage in non-relational DBMSs such as Cassandra. A key number uniquely identifies a collection of data contained in the value. For example, (CID, Customer Data) would enable a single number (CID) to retrieve data associated with the specified customer.

Label A simple text item displayed on a form or report. Values cannot be altered by the user.

Latency Time delay in a system. In a Web-based system, the delay created by slow links. Long download times create higher latency which leads to more server conflicts.

Leaves The bottom nodes in a B-tree (opposite from the root).

LEFT JOIN An outer join that includes all of the rows from the “left” table, even if there are no matching rows in the “right” table. The missing values are indicated by Nulls. *See right join and inner join.* Left and right are defined by the order the tables are listed; left is first.

Lifetime The length of time that a programming variable stays available. For example, variables created within subroutines are created when the routine is executed and then destroyed when it exits. Global variables stay alive for all routines within the module.

Lift The potential gain attributed to an association rule compared to purchases without the rule.

LIKE The SQL pattern-matching operator used to compare string values. The standard uses percent (%) to match any number of characters, and underscore (_) to match a single character. Some systems (e.g., Access) use an asterisk (*) and a question mark (?) instead.

LIMIT The CQL keyword used in the Cassandra SELECT command to set the maximum number of rows retrieved. A default value is always used even if the LIMIT command is not specified, so beware of queries that do not retrieve the expected number of rows.

List One type of Cassandra data collection (along with map and set). It enables a single column to hold multiple values, such as multiple e-mail addresses for one person. A list stores items as a sequence (1, 2, 3, ...)

List box A control on a form that displays a list of choices. The list is always displayed and takes up a fixed amount of space on the screen.

List of values (LOV) An important technique for selecting data on Oracle forms. Used instead of combo/select boxes, the Oracle form maintains a small, buffered list of data that can be selected for a text box. Particularly useful in distributed databases because only portions of the list are sent to the user.

Local area network (LAN) A collection of personal computers within a small geographic area. All components of the network are owned or controlled by one company.

Local variable A variable defined within a subroutine. It can be accessed only within that subroutine and not from other procedures.

Logic Logic statements that define a program's purpose and structure. Can be written in pseudocode independently of the program's syntax. Logic structures include loops, conditions, subroutines, and input/output commands.

Logical security Determining which users should have access to which data. It deals with preventing three data problems: (1) unauthorized disclosure, (2) unauthorized modification, and (3) unauthorized withholding.

Loop Each loop must have a beginning, an end condition, and some way to increment a variable. *See iteration.*

Map One type of Cassandra data collection (along with list and set). It enables a single column to hold multiple values stored as key-value pairs. For instance, a phone map might hold items such as cell=1111, home=2222, work=3333.

Market basket analysis *See association rules.*

Master-detail A common one-to-many relationship often found on business forms, where the main form (e.g., Order) displays data for the master component, and a subform (e.g., Order Items) displays detail data. Sometimes called a parent-child relationship.

Measure The numeric data displayed in an OLAP cube.

Menu A set of application commands grouped together—usually on a toolbar. It provides an easy reference for commonly used commands and highlights the structure of the application.

Metadata Data about data, or the description of the data tables and columns. Usually held in the data dictionary. For example, table definitions and column domains are metadata.

Method A function or operation that a class can perform. For example, a Customer class would generally have an AddNew method that is called whenever a new customer object is added to the database.

Microsoft Assistance Markup Language (MAML) A help system authoring language similar to XML introduced by Microsoft for use with the Windows Vista operating system. Discontinued.

Middle tier In a multi-tier application, a set of programs that lies between the front and back ends. It is generally used to define and process business rules.

Modal form A form that takes priority on the screen and forces the user to deal with it before continuing. It should be avoided because it interrupts the user.

Module (or package) A collection of subroutines, generally related to a common purpose.

MsgBox A predefined method in Windows for displaying a brief message on the screen and presenting a few limited choices to the user. Because it is modal and interrupts the user, it should be used sparingly.

Multidimensional expressions (MDX) A language for querying OLAP data that was initiated by Microsoft, but in the process of becoming a standard. It is designed to handle relatively complex computations and OLAP cubes.

Multiplicity The UML term for signifying the quantities involved in an association. It is displayed on an association line with a minimum value, an ellipsis (...), and a maximum value or asterisk (*) for many. For example, a customer can place from zero to many orders, so the multiplicity is (0...*).

N-ary association An association among three or more classes. It is drawn as a diamond on a UML class diagram. The term comes from extending English terms: unary means one, binary means two, ternary means four; so N-ary means many.

Naming conventions Program teams should name their variables and controls according to a consistent format. One common approach is to use a three letter prefix to identify the type of variable, followed by a descriptive name.

Nested conditions Conditional statements that are placed inside other conditional statements. For example, If (x > 0) Then ... If (y < 4) Then ... Some nested conditions can be replaced with a Case statement.

Nested query *See subquery.*

Network attached storage (NAS) Disk drives or other storage media that are connected to servers via a network connection instead

of building them into the computer itself. The separation provides some autonomy, making it easier to provide backup and to replace the storage devices. Typically the drives are RAID configured.

Network database An older DBMS type that expanded the hierarchical database by supporting multiple connections between entities. A network database is characterized by the requirement that all connections had to be supported by an index.

Node An entry in a B-tree that contains a key value and links to other nodes. The top-most node is the root, the bottom nodes are leaves.

Non-relational database Several newer database designs were developed to optimize performance for specific tasks--primarily for Web-based applications. The data storage does not follow data normalization rules and imposes limitations on queries to store and retrieve huge amounts of data.

Normalization *See data normalization.*

NoSQL Read as either Not SQL or Not Only SQL. Often used to refer to non-relational databases where data storage does not follow normalization rules. More specifically, the tools rarely support JOIN queries. By restricting operations allowed, the systems are optimized for specific data storage and retrieval tasks for giant databases.

NOT The SQL negation operator. Used in the WHERE clause to reverse the truth value of a statement. *See DeMorgan's law.*

NotInList An event corresponding to a combo box in Access. It is triggered when a user enters a value that does not yet exist in the selected list. Often used to add new data to a table, such as new customers.

Null A missing (or currently unassigned) value.

Object An instance or particular example of a class. For example, in an Employee class, one individual employee would be an object. In a relational environment, a class is stored as a table, while an individual row in the table contains data for one object.

Object browser A tool provided within Microsoft software that displays properties and methods for available objects.

Object-oriented database management system (OODBMS) A database system that holds objects, including properties and methods. It supports links between objects, including inheritance.

Object-oriented programming A programming methodology where code is encapsulated within the definition of various objects (or classes). Systems are built from (hopefully) reusable objects. You control the system by manipulating object properties and calling object methods.

Offset, file A method of identifying the location of data within an operating system file. The offset is the number of bytes from the start of the file. It is often used as a pointer to the physical data items.

Online analytical processing (OLAP) The use of a database for data analysis. The focus is on retrieval of the data. The primary goals are to provide acceptable response times, maintain security, and make it easy for users to find the data they need.

Online transaction processing (OLTP) The use of a database for transaction processing. It consists of many insert and update operations and supports hundreds of concurrent accesses. High-speed storage of data, reliability, and data integrity are primary goals. Examples include airline reservations, online banking, and retail sales.

Open database connectivity (ODBC) A standard created by Microsoft to enable software to access a variety of databases. Each DBMS vendor provides an ODBC driver. Application code can generally be written once. To change the DBMS, you simply install and set up the proper ODBC driver.

Optimistic lock A transaction lock that does not block other processes. If the data is changed between read and write steps, the system generates an error that must be handled by code.

Option button A round button that is used to indicate a choice. By the design guide, option buttons signify mutually exclusive choices, as opposed to check boxes.

ORDER BY The clause in the SQL SELECT statement that lists the columns to sort the output. The modifiers ASC and DESC are

used to specify ascending and descending sort orders.

Outer join A generic term that represents a left join or a right join. It returns rows from a table, even if there is no matching row in the other table.

Pack A maintenance operation that must be periodically performed on a database to remove fragments of deleted data.

Package A UML mechanism to group logical elements together. It is useful for isolating sections of a design. Packages can provide an overview of the entire system without having to see all the details.

Page footer A report element that appears at the bottom of every page. Often used for page numbers.

Page header A report element that appears at the top of every page. Often used for column headings and subtitles.

Parameter A variable that is passed to a subroutine or function and used in its computations.

Partition, SQL A SQL 2003 standard used to implement OLAP computations. Similar to a GROUP BY clause because you specify columns whose values are used to define the partitions. But partitions offer additional options and enable you to display detail and aggregate data at the same time. Also known as a data window.

Pass-by-reference A subroutine parameter that can be altered within the subroutine. If it is altered, the new value is returned to the calling program. That is, the subroutine can alter variables in other parts of the code. Usually a dangerous approach. *See pass-by-value.*

Pass-by-value A subroutine parameter that cannot be altered within the subroutine. Only its value is passed. If the subroutine changes the value, the original value in the calling program is not changed.

Pass-through query SQL queries that are ignored by Access on the front end. They enable you to write complex SQL queries that are specific to the server database.

Peer-to-peer Refers to distributed databases and servers. No single server or DBMS is

in charge of coordinating processing. All devices are considered equal and queries can be addressed to any server (usually the closest one).

Persistent objects In object-oriented programming, the ability to store objects (in a file or database) so that they can be retrieved at a later date.

Persistent stored modules (PSM) In SQL-99 a proposed method for storing methods associated with objects. The module code would be stored and retrieved automatically by the DBMS.

Pessimistic lock A complete isolation level that blocks other processes from reading a locked piece of data until the transaction is complete. Program code will receive an error message if the data element is locked.

Physical security The branch of security that involves physically protecting the equipment and people. It includes disaster planning, physical access to equipment, and risk analysis and prevention.

PivotTable Microsoft's tool for enabling managers to examine OLAP data dynamically—typically inside of an Excel spreadsheet. The data is extracted from the database or OLAP cube, and managers can click buttons to examine summaries or details.

Pointer A logical or physical address of a piece of data.

Polymorphism In a class hierarchy each new class inherits methods from the prior classes. Through polymorphism you can override those definitions and assign a new method (with the same name) to the new class.

Primary key A column or set of columns that identify a particular row in a table. In non-relational DBMSs, the primary key defines how the data will be stored and retrieved.

Private key In a dual-key encryption system, the key that is never revealed to anyone else. A message encrypted with a public key can be decrypted only with the matching private key.

Procedural language A traditional programming language that is based on following procedures and is typically executed one statement at a time. Compared to SQL, which operates on sets of data with one

command.

Procedure A subroutine or function that is designed to perform one specific task. It is generally wise to keep procedures small.

Property An attribute or feature of an entity that we wish to track. The term is often applied in an object-oriented context. *See attribute.*

Prototype An initial outline of an application that is built quickly to demonstrate and test various features of the application. Often used to help users visualize and improve forms and reports.

Pseudo column A computed column in a table that can only display data, not accept updates. Used to predefine row-wise calculations that are commonly used, such as $\text{Value}=\text{price}*\text{quantity}$.

Public key In a dual-key encryption system, the key that is given to the public. A message encrypted with a public key can only be decrypted with the matching private key.

Query by example (QBE) A fill-in-the-form approach to designing queries. You select tables and columns from a list and fill in blanks for conditions and sorting. It is relatively easy to use, requires minimal typing skills, generally comes with a Help system, and is useful for beginners.

RANK A data mining extension to SQL that rank orders data. If ties exist, they receive the same rank, and the next entry receives rank that is equal to its position in the overall list. Compare to DENSE RANK.

Rapid application development (RAD) A systems design methodology that attempts to reduce development time through efficiency and overlapping stages.

Redundant array of independent drives (RAID) A collection of several inexpensive disk drives that are treated as a single drive by the operating system. Data is written in stripes across the drives, supporting parallel operations to substantially improve performance and backup at the same time.

Referential integrity A data integrity constraint where data can be entered into a foreign key column only if the data value already exists in the base table. For example, clerks should not be able to enter an Order for

CustomerID 1173 if CustomerID 1173 is not in the Customer table.

Reflexive association A relationship from one class back to itself. Most commonly seen in business in an Employee class, where some employees are managers over other employees.

Reflexive join A situation that exists when a table is joined to itself through a second column. For example, the table `Employee(EmployeeID, ..., ManagerID)` could have a join from ManagerID to EmployeeID.

Regular expression (RegEx) A pattern matching tool for searching complex strings. Search online references for details. It has many options to match individual characters, sets of characters, repetitions, and sequences.

Relational database The most popular type of DBMS. All data is stored in tables (sometimes called relations). Tables are logically connected by the data they hold (e.g., through the key values). Relational databases should be designed through data normalization.

Relationship An association between two or more entities. *See association.*

Repeating groups Groups of data that repeat, such as items being ordered by a customer, multiple phone numbers for a client, and tasks assigned to a worker.

Replicate Make a deliberate copy of a database so it can be distributed to a new location and the contents later synchronized with the master copy.

Replication manager In a distributed database that relies on replication, the manager is an automated system that transfers changes to the various copies of the database. It has to handle conflicts if two people changed the same data before it was replicated.

Report footer A report element that appears at the end of the report. Often used for summary statistics or graphs.

Report header A report element that appears only at the start of the report. Often used for title pages and overviews.

Report services A server-based tool that processes reports and enables users to browse for reports and generate them on demand.

Report writer A DBMS tool that enables you

to set up reports on the screen to specify how items will be displayed or calculated. Most of these tasks are performed by dragging data onto the screen.

Resource file A special text file associated with a form or report. It is commonly used to hold language translations. The form references a tag in the resource file, and the system pulls the matching value from the appropriately translated resource file for the desired language.

Resume A VBA error-handling operator that tells the processor to return to a new location and continue evaluating the code. Resume by itself returns to the line that caused the error. Resume Next returns to the line immediately following the error. Resume <label> sends the processor to a new location.

REVOKE The SQL command used to remove permissions that were granted to certain users.

RIGHT JOIN An outer join that includes all the rows from the “right” table, even if there are no matching rows in the “left” table. The missing values are indicated by Nulls. *See right join and inner join.* Left and right are defined by the order the tables are listed; left is first.

Rivest-Shamir-Adelman (RSA) encryption A dual-key encryption system that was patented in the United States by three mathematicians. *See dual-key encryption.*

Role Security group where permissions are assigned by the task or role that is performed instead of to individual people. Roles make it easier to change employee permissions.

Roll back A database system transaction feature. If an error occurs in a sequence of changes, the preceding changes can be rolled back to restore the database to a safe state with correct data.

Roll forward If an error occurs in processing transactions, the database can be restarted and loaded from a known checkpoint. Then partially completed transactions can be rolled forward to record the interrupted changes.

Roll up The act of aggregating detail data into a display of category summaries. Commonly used in examining data in a data warehouse or OLAP application. *See drill down.*

Root The top-most or entry node in a B-tree.

Row-by-row calculations The way that SQL performs in-line calculations. For example, the statement `SELECT Price*Quantity AS Extended` goes through each row and multiplies the row’s value for Price times the matching value for Quantity.

Scalability The ability to handle increased demands without needing to make major changes to an application. Generally applied to server hardware, it is often supported through server clusters (or farms).

Schema A collection of tables that are grouped together for a common purpose. Generally used as a naming method.

Scope Refers to where a variable is accessible. Variables defined within a subroutine can typically be accessed only by code within that subroutine. Variables defined within a module are globally accessible by any code within that module.

Scroll bars A common graphical interface feature used to move material horizontally or vertically.

Second normal form (2NF) A table is in 2NF if every nonkey column depends on the entire key (not just part of it). This issue arises only if there is a concatenated key (with multiple columns).

Secure sockets layer (SSL) An Internet standard that uses RSA encryption with public keys to establish a secure, encrypted session between a browser and a Web server.

SELECT The primary data retrieval command for SQL. The main components are `SELECT ... FROM ... INNER JOIN ... WHERE`.

Self-join A table joined to itself. *See reflexive join.*

Serialization A transaction requirement that specifies that each transaction is treated completely separately and run as if there were no other transactions.

Set One type of Cassandra data collection (along with list and map). It enables a single column to hold multiple values stored as key-value pairs. Order is unimportant but usually alphabetical. Columns are defined with angle brackets: `email set<text>`

Shell site *See cold site.*

Single-row form An input form that displays data from one row of a table at a time. The most common input form, since the designer has full control over the layout of the form.

Snapshot In a constantly changing database, you can take a snapshot that provides a copy of the data at one point in time.

Snowflake design A database design used for OLAP. A fact table is connected to dimension tables. The dimension tables can be joined to other dimension tables. It is less restrictive than the star design.

Solid-state drive (SSD) A data storage device that uses electronic flash (non-volatile) memory instead of moving disks. Modern versions are usually smaller than rotating disks, draw less power, and are tens or hundreds of times faster at data storage and retrieval.

Sparse table A table that contains missing data for many columns. Often an indicator that non-relational systems might be more efficient at storing and retrieving data because they use space only for actual data.

SQL A standardized database language, used for data retrieval (queries), data definition, and data manipulation.

SQL 1999 (SQL3) Approved in 1999, SQL 1999 was largely designed to add object-oriented features to the SQL language.

SQL 2003 An almost completed new version of SQL. Its primary contribution is a formal definition of integrating XML and multimedia into relational databases.

SQL 2008 The latest release of the SQL standard which is primarily an updated version of SQL 2006. Added support for XQuery and procedural code (persistent stored modules).

SQL injection attack An attack on the application and database usually performed by an outsider. The person enters special text into a text box to try and alter the SQL query. The special text often includes the comment character (--) and quotation marks. Your application code should test all user input and remove those characters.

Star design A database design used for OLAP. A fact table is connected to dimension

tables that provide categories for analysis. More restrictive than the snowflake design, all dimension tables are directly connected to the fact table.

Startup form A form that is used to direct users to different parts of the application. Often used as the first form to appear. Options on the form should match the tasks of the users. Also called a switchboard form.

Storage area network (SAN) A high-speed network that moves disk drives out of the computer and into a separate location. Separating the disk drives from the processor makes it easier to share the data, to back it up, and to scale up by adding new processors.

Style sheet A special file that describes the desired layout, fonts, and styles for a set of Web pages. It is a powerful method to establish and change styles on many Web pages through making minor changes to one file.

Subform form A form that is displayed inside another (main) form. The data in the subform is generally linked to the row currently being displayed on the main form.

Subquery Using a second query to retrieve additional data within the main query. For example, to retrieve all sales where price was greater than the average, the WHERE clause could use a subquery to compute the average price.

Subroutine A separate section of code designed to perform one specific purpose. A subroutine can use parameters to exchange data with the calling routine.

Subtable In SQL 1999 a subtable inherits all of the columns from a base table. It provides inheritance similar to that of the abstract data types; however, all data is stored in separate columns.

Super-aggregate In a query, totals of totals. An important concept in OLAP queries provide by the ROLLUP option.

Support A data mining measure in association rules, measured by the percent of transactions that contain both items.

Surrogate keys Internally generated keys used to identify objects. They are often better keys than keys created by external sources because they are easier to guarantee

uniqueness. See AutoNumber.

Switchboard form See *startup form*.

Synonym A short name for the full database path. Advantages of the short name are that it is easy to remember and the user never needs to know where the data is located. In addition, if everyone uses the synonym, database administrators can easily move the server databases to different locations just by altering the synonym's properties.

Syntax The specific format of commands that can be created in a program. Programs consist of logical steps, but each command must be given in the proper syntax for the compiler to understand it. Compilers generally check for syntax and prompt you with messages.

Tab order The sequence of controls followed on a form when the user presses the tab or return keys.

Table A collection of data for one class or entity. It consists of columns for each attribute and a row of data for each specific entity or object.

Tablespace In Oracle a tablespace is disk space that is allocated to hold tables, indexes, and other system data. You must first know the approximate size of the database.

Tabular form An input form that displays data in columns and rows. It is used when there are few columns of data or when the user needs to see multiple rows at the same time.

Template A special form (or report) that defines the standard elements to be applied to all forms. It also specifies styles and common code. It is used to improve consistency of forms and reports within an application.

Text box A common form control that is used to display and enter data.

Third normal form (3NF) A table is in third normal form (3NF) if each nonkey column depends on the whole key and nothing but the key.

Three-tier client/server A client/server system with a middle layer to hold code that defines business rules and consolidates access to various transaction servers.

Toggle button A three-dimensional variation of the check box. It is used to signify a choice

of options.

Toolbar A small object in applications that can hold buttons and text menus. Users can execute commands with one or two mouse clicks. Used to hold frequently used commands, and commands that are used across the entire application, such as printing.

Tooltip A short message that is displayed when the user moves the mouse cursor over an item on the screen. Extremely useful for identifying the purpose of icons.

TOP A SQL SELECT clause provided by Access that restricts the displayed output to a specified number of rows. You can set the number of rows directly or use a percentage of the total number.

Transaction In a database application a transaction is a set of changes that must all be made together. Transactions must be identified to the DBMS and then committed or rolled back (if there is an error). For example, a transfer of money from one bank account to another requires two changes to the database—both must succeed or fail together.

Transaction log A sequential file that records transactions as they are being created. If something happens during the transaction update, the DBMS uses the transaction log to complete or roll back the incomplete transactions.

Transaction processing Collecting data for the purpose of recording transactions. Common examples include sales, human resource management, and financial accounting.

Trigger An event that causes a procedure to be executed. For example, clicking a button can be a trigger, as can a change in a data value.

Tunable consistency A concept emphasized in the Cassandra distributed DBMS. A developer can choose the level of consistency required for read and write operations, such as a single node, agreement across all nodes, or consistency across a quorum. Higher consistency (ALL) is closer to SQL but can result in delays.

Tuning Setting indexes, rewriting queries, and setting storage and other parameters to improve the database application performance.

Two-phase commit A mechanism for

handling concurrency and deadlock problems in a distributed database. In the first phase the coordinating DBMS sends updates to the other databases and asks them to prepare the transaction. Once they have agreed, the coordinator sends a message to commit the changes.

Unicode A standard method of storing and displaying a variety of character sets. Almost all current world character sets have been defined, as well as several ancient languages. It uses 2 bytes to represent each character, enabling it to handle over 65,000 characters or ideograms.

Unified Modeling Language (UML) A standardized modeling language for designing and documenting computer and business systems.

UNION An SQL clause to combine rows from two SELECT statements. Both queries must have the same number of columns with the same domains. Most systems also support INTERSECT and EXCEPT (or SUBTRACT) operators.

Universally unique identifier (uuid) A type of large number defined by an ISO standard to enable a computer to generate values that will be unique across any computer. Several types are defined (at least 4). Some use the unique MAC network card address and a timestamp. Others use pure random numbers, which might overlap with a tiny probability. Often used as generated key values in distributed (non-relational) databases.

UPDATE A SQL data manipulation command that changes the values in specified columns. A WHERE clause specifies which rows will be affected.

Use case In UML, a diagram that shows how a specific group of people will use the system.

User interface The look and feel of the application as it is seen by the user. Graphical interfaces are commonly employed in which users can manipulate icons and data on the screen to perform their tasks.

Validation tables Simple tables of one or two columns that contain standardized data for entry into other tables. For example, a list of departments would be stored in a validation table. To enter a department name into an

Employee table, the user would be given a choice of the rows in the validation table.

VARCHAR A common method for storing character data. It stands for variable characters. Each column of data uses the exact amount of bytes needed to store the specific data.

Variable A location in memory used to hold temporary values. Variables have a scope and a lifetime depending on where they are created and how they are defined. They also have a specific data type, although the Variant data type in VBA can hold any common type of data.

Variable length storage A system of storing data columns where only a pointer is stored in the actual row. The actual data is stored in a pooled area.

Vertical partition Splitting a table into groups based on the columns of data. Large columns or columns that are seldom used (e.g., pictures) can be moved to slower, cheaper storage devices.

View A saved query. You can build new queries that retrieve data from the view. A view is saved as an SQL statement—not as the actual data.

Virtual machine (VM) A method of running multiple computer operating systems on one physical computer. Each VM has its own disk space and operating system and appears to be a separate computer. The underlying computer needs to support multiprocessing and have sufficient memory. VMs are relatively easy to backup and transfer to other physical computers.

Visual Basic (VB) A stand-alone programming language sold by Microsoft and used to develop applications for the Windows environment. The professional version supports database connections. The program can be compiled into a stand-alone executable file.

Visual Basic for Applications (VBA) The programming language that underlies almost all of Microsoft's tools, including Access.

Volume table of contents (VTOC) A design tool that can be used to outline the overall structure of an application. It generally shows a sequence of interrelated menus.

WHERE The SQL clause that restricts the

rows that will be used in the query. It can also refer to data in subqueries.

Wide area network (WAN) A network that is spread across a larger geographic area. Parts of the network are outside the control of a single firm. Long-distance connections often use public carriers.

With ... End With In VBA, a shortcut for examining or altering several properties for a single object. Once the object is specified in the With statement, you simply refer to the properties inside the “loop.”

WITH GRANT OPTION A security permission option that transfers the ability to assign permissions to the specified role or user.

World Wide Web (WWW) A first attempt to set up an international database of information. Web browsers display graphical pages of information. Hypertext connections enable you to get related information by clicking highlighted words or icons. A standard method for displaying text and images on client computers.

XML *See extensible markup language.*

XML schema A definitional file that describes the tags, data types, and attributes allowed in an associated XML data file.

XQuery A standardized language to retrieve individual elements or groups of data from an XML file or text group.