

## GLOSSAR Glossary

### A

|                                      |  |
|--------------------------------------|--|
| <b>abend.</b>                        | Abnormal end of task   |
| <b>abend reason code.</b>            | A 4-byte hexadecimal code that uniquely identifies a problem with DB2. A complete list of DB2 abend reason codes and their explanations is contained in <i>DB2 Messages and Codes</i> .  |
| <b>abnormal end of task (abend).</b> | Termination of a task, job, or subsystem because of an error condition that recovery facilities cannot resolve during execution.   |
| <b>access method services.</b>       | The facility that is used to define and reproduce VSAM key-sequenced data sets.  |
| <b>access path.</b>                  | The path that is used to locate data that is specified in SQL statements. An access path can be indexed or sequential.   |
| <b>active log.</b>                   | The portion of the DB2 log to which log records are written as they are generated. The active log always contains the most recent log records, whereas the archive log holds those records that are older and no longer fit on the active log.   |
| <b>address space.</b>                | A range of virtual storage pages that is identified by a number (ASID) and a collection of segment and page tables that map the virtual pages to real pages of the computer's memory   |
| <b>address space connection.</b>     | The result of connecting an allied address space to DB2. Each address space that contains a task that is connected to DB2 has exactly one address space connection, even though more than one task control block (TCB) can be present. See also <i>allied address space</i> and <i>task control block</i> .                                |
| <b>agent.</b>                        | As used in DB2, the structure that associates all processes that are involved in a DB2 unit of work. An <i>allied agent</i> is generally synonymous with an <i>allied thread</i> . <i>System agents</i> are units of work that process independently of the allied agent, such as prefetch processing, deferred writes, and service tasks. |
| <b>alias.</b>                        | An alternative name that can be used in SQL statements to refer to a table or view in the same or a remote DB2 subsystem.  |
| <b>allied address space.</b>         | An area of storage that is external to DB2 and that is connected to DB2. An allied address space is capable of requesting DB2 services.  |
| <b>allied thread.</b>                | A thread that originates at the local DB2 subsystem and that can access data at a remote DB2 subsystem.  |
| <b>ambiguous cursor.</b>             | A database cursor that is not defined with the FOR FETCH ONLY clause or the FOR<UPDATE OF clause, is not defined on a read-only result table, is not the target of a WHERE CURRENT clause on an SQL UPDATE or DELETE statement, and is in a plan or package that contains either PREPARE or EXECUTE IMMEDIATE SQL statements.              |
| <b>APAR.</b>                         | Authorized program analysis report.  |
| <b>APAR fix corrective service.</b>  | A temporary correction of a DB2 defect. The correction is temporary, because it is usually replaced at a later date by a more permanent correction, such as a program temporary fix (PTF).   |
| <b>APF.</b>                          | Authorized program facility.   |
| <b>API.</b>                          | Application programming interface.   |

|   |  |
|---|--|
| <b>APPL.</b>                                      | A VTAM network definition statement that is used to define DB2 to VTAM as an application program that uses SNA LU 6.2 protocols.   |
| <b>application.</b>                               | A program or set of programs that performs a task; for example, a payroll application.   |
| <b>application-directed connection.</b>           | A connection that an application manages using the SQL CONNECT statement.  |
| <b>application plan.</b>                          | The control structure that is produced during the bind process. DB2 uses the application plan to process SQL statements that it encounters during statement execution.   |
| <b>application process.</b>                       | The unit to which resources and locks are allocated. An application process involves the execution of one or more programs.  |
| <b>application programming interface (API).</b>   | A functional interface that is supplied by the operating system or by a separately orderable licensed program that allows an application program that is written in a high-level language to use specific data or functions of the operating system or licensed program. |
| <b>application requester.</b>                     | The component on a remote system that generates DRDA requests for data on behalf of an application. An application requester accesses a DB2 database server using the DRDA application-directed protocol.  |
| <b>application server.</b>                        | The target of a request from a remote application. In the DB2 environment, the application server function is provided by the distributed data facility and is used to access DB2 data from remote applications  |
| <b>archive log.</b>                               | The portion of the DB2 log that contains log records that have been copied from the active log.  |
| <b>ASCII.</b>                                     | An encoding scheme that is used to represent strings in many environments, typically on PCs and workstations. Contrast with <i>EBCDIC</i> and <i>Unicode</i>   |
| <b>attachment facility.</b>                       | An interface between DB2 and TSO, IMS, CICS, or batch address spaces. An attachment facility allows application programs to access DB2.  |
| <b>attribute.</b>                                 | A characteristic of an entity. For example, in database design, the phone number of an employee is one of that employee's attributes.  |
| <b>authorization ID.</b>                          | A string that can be verified for connection to DB2 and to which a set of privileges is allowed. It can represent an individual, an organizational group, or a function, but DB2 does not determine this representation.   |
| <b>authorized program analysis report (APAR).</b> | A report of a problem that is caused by a suspected defect in a current release of an IBM licensed program.  |
| <b>authorized program facility (APF).</b>         | A facility that permits the identification of programs that are authorized to use restricted functions.  |
| <b>auxiliary index.</b>                           | An index on an auxiliary table in which each index entry refers to a LOB.  |
| <b>auxiliary table.</b>                           | A table that stores columns outside the table in which they are defined. Contrast with <i>base table</i> .   |

## B

|   |   |
|---|---|
| <b>backward log recovery.</b>                 | The fourth and final phase of restart processing during which DB2 scans the log in a backward direction to apply UNDO log records for all aborted changes.  |
| <b>base table.</b>                            | (1) A table that is created by the SQL CREATE TABLE statement and that holds persistent data. Contrast with <i>result table</i> and <i>temporary table</i> .<br>(2) A table containing a LOB column definition. The actual LOB column data is not stored with the base table. The base table contains a row identifier for each row and an indicator column for each of its LOB columns. Contrast with <i>auxiliary table</i> .   |
| <b>base table space.</b>                      | A table space that contains base tables.  |
| <b>basic sequential access method (BSAM).</b> | An access method for storing or retrieving data blocks in a continuous sequence, using either a sequential access or a direct access device.  |
| <b>before trigger.</b>                        | A trigger that is defined with the trigger activation time BEFORE.  |
| <b>binary large object (BLOB).</b>            | A sequence of bytes where the size of the value ranges from 0 bytes to 2 GB.1. Such a string does not have an associated CCSID.   |
| <b>binary string.</b>                         | A sequence of bytes that is not associated with a CCSID. For example, the BLOB data type is a binary string.  |
| <b>bind.</b>                                  | The process by which the output from the SQL precompiler is converted to a usable control structure, often called an access plan, application plan, or package. During this process, access paths to the data are selected and some authorization checking is performed. The types of bind are:<br><b>automatic bind.</b> (More correctly, <i>automatic rebind</i> )A process by which SQL statements are bound automatically (without a user issuing a BIND command) when an application process begins execution and the bound application plan or package it requires is not valid.<br><b>dynamic bind.</b> A process by which SQL statements are bound as they are entered.<br><b>incremental bind.</b> A process by which SQL statements are bound during the execution of an application process, because they could not be bound during the bind process, and VALIDATE(RUN) was specified.<br><b>static bind.</b> A process by which SQL statements are bound after they have been precompiled. All static SQL statements are prepared for execution at the same time. |
| <b>BLOB.</b>                                  | Binary large object   |
| <b>BMP.</b>                                   | Batch Message Processing (IMS).   |
| <b>bootstrap data set (BSDS).</b>             | A VSAM data set that contains name and status information for DB2, as well as RBA range specifications, for all active and archive log data sets. It also contains passwords for the DB2 directory and catalog, and lists of conditional restart and checkpoint records.  |
| <b>BSAM.</b>                                  | Basic sequential access method.   |
| <b>BSDS.</b>                                  | Bootstrap data set.   |
| <b>buffer pool.</b>                           | Main storage that is reserved to satisfy the buffering requirements for one or more table spaces or indexes.  |
| <b>built-in function.</b>                     | A function that DB2 supplies. Contrast with <i>user-defined function</i> .  |

## C

- CAF.** Call attachment facility.
- call attachment facility (CAF).** A DB2 attachment facility for application programs that run in TSO or MVS batch. The CAF is an alternative to the DSN command processor and provides greater control over the execution environment.
- cascade delete.** The way in which DB2 enforces referential constraints when it deletes all descendent rows of a deleted parent row.
- cast function.** A function that is used to convert instances of a (source) data type into instances of a different (target) data type. In general, a cast function has the name of the target data type. It has one single argument whose type is the source data type; its return type is the target data type.
- catalog.** In DB2, a collection of tables that contains descriptions of objects such as tables, views, and indexes.
- catalog table.** Any table in the DB2 catalog.
- CCSID.** Coded character set identifier.
- CDB.** Communications database.
- CEC.** Central electronic complex. See *central processor complex*.
- central electronic complex (CEC).**  
See *central processor complex*.
- central processor complex (CPC).**  
A physical collection of hardware (such as an ES/3090) that consists of main storage, one or more central processors, timers, and channels.
- character large object (CLOB).** A sequence of bytes representing single-byte characters or a mixture of single- and double-byte characters where the size of the value can be up to 2 GB. In general, character large object values are used whenever a character string might exceed the limits of the VARCHAR type.
- character set.** A defined set of characters.
- character string.** A sequence of bytes that represent bit data, single-byte characters, or a mixture of single-byte and multibyte characters.
- check constraint.** See *table check constraint*.
- check integrity.** The condition that exists when each row in a table conforms to the table check constraints that are defined on that table. Maintaining check integrity requires DB2 to enforce table check constraints on operations that add or change data.
- check pending.** A state of a table space or partition that prevents its use by some utilities and some SQL statements because of rows that violate referential constraints, table check constraints, or both.
- checkpoint.** A point at which DB2 records internal status information on the DB2 log; the recovery process uses this information if DB2 abnormally terminates.
- CI.** Control interval.
- CICS.** Represents one of the following products:  
**CICS Transaction Server for OS/390:** Customer Information Control System Transaction Server for OS/390

|  |   |
|--|---|
|  | <b>CICS/ESA:</b> Customer Information Control System/Enterprise Systems Architecture  |
| <b>CICS attachment facility.</b>               | <b>CICS/MVS:</b> Customer Information Control System/Multiple Virtual Storage<br>A DB2 subcomponent that uses the MVS subsystem interface (SSI) and cross storage linkage to process requests from CICS to DB2 and to coordinate resource commitment. |
| <b>CIDF.</b>                                   | Control interval definition field.  |
| <b>claim.</b>                                  | A notification to DB2 that an object is being accessed. Claims prevent drains from occurring until the claim is released, which usually occurs at a commit point. Contrast with <i>drain</i> .  |
| <b>claim class.</b>                            | A specific type of object access that can be one of the following:<br>Cursor stability (CS)<br>Repeatable read (RR)<br>Write  |
| <b>claim count.</b>                            | A count of the number of agents that are accessing an object.   |
| <b>class of service.</b>                       | A VTAM term for a list of routes through a network, arranged in an order of preference for their use.   |
| <b>clause.</b>                                 | In SQL, a distinct part of a statement, such as a SELECT clause or a WHERE clause.  |
| <b>client.</b>                                 | See <i>requester</i> .  |
| <b>CLIST.</b>                                  | Command list. A language for performing TSO tasks.  |
| <b>CLOB.</b>                                   | Character large object.   |
| <b>CLPA.</b>                                   | Create link pack area.  |
| <b>clustering index.</b>                       | An index that determines how rows are physically ordered in a table space.  |
| <b>coded character set.</b>                    | A set of unambiguous rules that establish a character set and the one-to-one relationships between the characters of the set and their coded representations.   |
| <b>coded character set identifier (CCSID).</b> | A 16-bit number that uniquely identifies a coded representation of graphic characters. It designates an encoding scheme identifier and one or more pairs consisting of a character set identifier and an associated code page identifier.             |
| <b>column.</b>                                 | The vertical component of a table. A column has a name and a particular data type (for example, character, decimal, or integer).  |
| <b>column function.</b>                        | An operation that derives its result by using values from one or more rows. Contrast with <i>scalar function</i> .  |
| <b>"come from" checking.</b>                   | An LU 6.2 security option that defines a list of authorization IDs that are allowed to connect to DB2 from a partner LU.  |
| <b>command.</b>                                | A DB2 operator command or a DSN subcommand. A command is distinct from an SQL statement.  |
| <b>command recognition character (CRC).</b>    | A character that permits an MVS console operator or an IMS subsystem user to route DB2 commands to specific DB2 subsystems.   |
| <b>commit.</b>                                 | The operation that ends a unit of work by releasing locks so that the database changes that are made by that unit of work can be perceived by other processes.  |
| <b>commit point.</b>                           | A point in time when data is considered consistent.   |

|  |   |
|--|---|
| <b>committed phase.</b>                          | The second phase of the multisite update process that requests all participants to commit the effects of the logical unit of work.  |
| <b>common service area (CSA).</b>                | In MVS, a part of the common area that contains data areas that are addressable by all address spaces.  |
| <b>communications database (CDB).</b>            | A set of tables in the DB2 catalog that are used to establish conversations with remote database management systems.  |
| <b>comparison operator.</b>                      | A token (such as =, >, <) that is used to specify a relationship between two values.  |
| <b>compression dictionary.</b>                   | The dictionary that controls the process of compression and decompression. This dictionary is created from the data in the table space or table space partition.  |
| <b>concurrency.</b>                              | The shared use of resources by more than one application process at the same time.  |
| <b>conditional restart.</b>                      | A DB2 restart that is directed by a user-defined conditional restart control record (CRCR).   |
| <b>connection ID.</b>                            | An identifier that is supplied by the attachment facility and that is associated with a specific address space connection.  |
| <b>consistency token.</b>                        | A timestamp that is used to generate the version identifier for an application. See also <i>version</i> .   |
| <b>constraint.</b>                               | A rule that limits the values that can be inserted, deleted, or updated in a table. See <i>referential constraint</i> , <i>table check constraint</i> , and <i>uniqueness constraint</i> .  |
| <b>control interval (CI).</b>                    | A fixed-length area or direct access storage in which VSAM stores records and creates distributed free space. Also, in a key-sequenced data set or file, the set of records pointed to by an entry in the sequence-set index record. The control interval is the unit of information that VSAM transmits to or from direct access storage. A control interval always includes an integral number of physical records. |
| <b>control interval definition field (CIDF).</b> | In VSAM, a field located in the 4 bytes at the end of each control interval; it describes the free space, if any, in the control interval.  |
| <b>conversation.</b>                             | Communication, which is based on LU 6.2 or Advanced Program-to-Program Communication (APPC), between an application and a remote transaction program over an SNA logical unit-to-logical unit (LU-LU) session that allows communication while processing a transaction.   |
| <b>coordinator.</b>                              | The system component that coordinates the commit or rollback of a unit of work that includes work that is done on one or more other systems.  |
| <b>correlated columns.</b>                       | A relationship between the value of one column and the value of another column.   |
| <b>correlated subquery.</b>                      | A subquery (part of a WHERE or HAVING clause) that is applied to a row or group of rows of a table or view that is named in an outer subselect statement.   |
| <b>correlation ID.</b>                           | An identifier that is associated with a specific thread. In TSO, it is either an authorization ID or the job name.  |
| <b>correlation name.</b>                         | An identifier that designates a table, a view, or individual rows of a table or view within a single SQL statement. It can be defined in any FROM clause or in the first clause of an UPDATE or DELETE statement.   |

|                                      |   |
|--------------------------------------|---|
| <b>cost category.</b>                | A category into which DB2 places cost estimates for SQL statements at the time the statement is bound. A cost estimate can be placed in either of the following cost categories: v A: Indicates that DB2 had enough information to make a cost estimate without using default values. v B: Indicates that some condition exists for which DB2 was forced to use default values for its estimate. The cost category is externalized in the COST_CATEGORY column of the DSN_STATEMNT_TABLE when a statement is explained. |
| <b>CPC.</b>                          | Central processor complex.  |
| <b>CRC.</b>                          | Command recognition character.  |
| <b>CRCR.</b>                         | Conditional restart control record. See also <i>conditional restart</i> .   |
| <b>create link pack area (CLPA).</b> | An option used during IPL to initialize the link pack pageable area.  |
| <b>created temporary table.</b>      | A table that holds temporary data and is defined with the SQL statement CREATE GLOBAL TEMPORARY TABLE. Information about created temporary tables is stored in the DB2 catalog, so this kind of table is persistent and can be shared across application processes. Contrast with <i>declared temporary table</i> . See also <i>temporary table</i> .   |
| <b>cross-memory linkage.</b>         | A method for invoking a program in a different address space. The invocation is synchronous with respect to the caller.   |
| <b>CS.</b>                           | Cursor stability.   |
| <b>CSA.</b>                          | Common service area.  |
| <b>CT.</b>                           | Cursor table.   |
| <b>current status rebuild.</b>       | The second phase of restart processing during which the status of the subsystem is reconstructed from information on the log.   |
| <b>cursor.</b>                       | A named control structure that an application program uses to point to a row of interest within some set of rows, and to retrieve rows from the set, possibly making updates or deletions.  |
| <b>cursor stability (CS).</b>        | The isolation level that provides maximum concurrency without the ability to read uncommitted data. With cursor stability, a unit of work holds locks only on its uncommitted changes and on the current row of each of its cursors.  |
| <b>cursor table (CT).</b>            | The copy of the skeleton cursor table that is used by an executing application process.   |
| <b>cycle.</b>                        | A set of tables that can be ordered so that each table is a descendent of the one before it, and the first table is a descendent of the last table. A self-referencing table is a cycle with a single member.   |

## D

|   |   |
|---|---|
| <b>DASD.</b>                              | Direct access storage device.   |
| <b>database.</b>                          | A collection of tables, or a collection of table spaces and index spaces.   |
| <b>database access thread.</b>            | A thread that accesses data at the local subsystem on behalf of a remote subsystem.   |
| <b>database administrator (DBA).</b>      | An individual who is responsible for designing, developing, operating, safeguarding, maintaining, and using a database.   |
| <b>database descriptor (DBD).</b>         | An internal representation of a DB2 database definition, which reflects the data definition that is in the DB2 catalog. The objects that are defined in a database descriptor are table spaces, tables, indexes, index spaces, and relationships.   |
| <b>database management system (DBMS).</b> | A software system that controls the creation, organization, and modification of a database and the access to the data stored within it.   |
| <b>database request module (DBRM).</b>    | A data set member that is created by the DB2 precompiler and that contains information about SQL statements. DBRMs are used in the bind process.  |
| <b>database server.</b>                   | The target of a request from a local application or an intermediate database server. In the DB2 environment, the database server function is provided by the distributed data facility to access DB2 data from local applications, or from a remote database server that acts as an intermediate database server. |
| <b>DATABASE 2 Interactive (DB2I).</b>     | The DB2 facility that provides for the execution of SQL statements, DB2 (operator) commands, programmer commands, and utility invocation.   |
| <b>data definition name (ddname).</b>     | The name of a data definition (DD) statement that corresponds to a data control block containing the same name.   |
| <b>Data Language/I (DL/I).</b>            | The IMS data manipulation language; a common high-level interface between a user application and IMS.   |
| <b>data sharing.</b>                      | The ability of two or more DB2 subsystems to directly access and change a single set of data.   |
| <b>data sharing group.</b>                | A collection of one or more DB2 subsystems that directly access and change the same data while maintaining data integrity.  |
| <b>data sharing member.</b>               | A DB2 subsystem that is assigned by XCF services to a data sharing group.   |
| <b>data space.</b>                        | A range of up to 2 GB of contiguous virtual storage addresses that a program can directly manipulate. Unlike an address space, a data space can hold only data; it does not contain common areas, system data, or programs.   |
| <b>data type.</b>                         | An attribute of columns, literals, host variables, special registers, and the results of functions and expressions.   |
| <b>date.</b>                              | A three-part value that designates a day, month, and year.  |
| <b>date duration.</b>                     | decimal integer that represents a number of years, months, and days.  |
| <b>datetime value.</b>                    | A value of the data type DATE, TIME, or TIMESTAMP.  |
| <b>DBA.</b>                               | Database administrator.   |
| <b>DBCLOB.</b>                            | Double-byte character large object.   |

|   |   |
|---|---|
| <b>DBCS.</b>                            | Double-byte character set.  |
| <b>DBD.</b>                             | Database descriptor.  |
| <b>DBID.</b>                            | Database identifier.  |
| <b>DBMS.</b>                            | Database management system.   |
| <b>DBRM.</b>                            | Database request module.  |
| <b>DB2 catalog.</b>                     | Tables that are maintained by DB2 and contain descriptions of DB2 objects, such as tables, views, and indexes.  |
| <b>DB2 command.</b>                     | An instruction to the DB2 subsystem allowing a user to start or stop DB2, to display information on current users, to start or stop databases, to display information on the status of databases, and so on.  |
| <b>DB2 for VSE &amp; VM.</b>            | The IBM DB2 relational database management system for the VSE and VM operating systems.   |
| <b>DB2I.</b>                            | DATABASE 2 Interactive.   |
| <b>DB2I Kanji Feature.</b>              | The tape that contains the panels and jobs that allow a site to display DB2I panels in Kanji.   |
| <b>DB2 PM.</b>                          | DATABASE 2 Performance Monitor.   |
| <b>DCLGEN.</b>                          | Declarations generator.   |
| <b>DDF.</b>                             | Distributed data facility.  |
| <b>ddname.</b>                          | Data definition name.   |
| <b>deadlock.</b>                        | Unresolvable contention for the use of a resource such as a table or an index.  |
| <b>declarations generator (DCLGEN).</b> | A subcomponent of DB2 that generates SQL table declarations and COBOL, C, or PL/I data structure declarations that conform to the table. The declarations are generated from DB2 system catalog information. DCLGEN is also a DSN subcommand.   |
| <b>declared temporary table.</b>        | A table that holds temporary data and is defined with the SQL statement DECLARE GLOBAL TEMPORARY TABLE. Information about declared temporary tables is not stored in the DB2 catalog, so this kind of table is not persistent and can only be used by the application process that issued the DECLARE statement. Contrast with <i>created temporary table</i> . See also <i>temporary table</i> . |
| <b>default value.</b>                   | A predetermined value, attribute, or option that is assumed when no other is explicitly specified.  |
| <b>deferred write.</b>                  | The process of asynchronously writing changed data pages to disk.   |
| <b>degree of parallelism.</b>           | The number of concurrently executed operations that are initiated to process a query.   |
| <b>delete rule.</b>                     | The rule that tells DB2 what to do to a dependent row when a parent row is deleted. For each relationship, the rule might be CASCADE, RESTRICT, SET NULL, or NO ACTION.   |
| <b>dependent.</b>                       | An object (row, table, or table space) that has at least one parent. The object is also said to be a dependent (row, table, or table space) of its parent. See <i>parent row</i> , <i>parent table</i> , <i>parent table space</i> .  |

|  |   |
|--|---|
| <b>dependent row.</b>  | A row that contains a foreign key that matches the value of a primary key in the parent row.  |
| <b>dependent table.</b>  | A table that is a dependent in at least one referential constraint.   |
| <b>descendent.</b>   | An object that is a dependent of an object or is the dependent of a descendent of an object.  |
| <b>descendent row.</b>   | A row that is dependent on another row, or a row that is a descendent of a dependent row.   |
| <b>descendent table.</b>   | A table that is a dependent of another table, or a table that is a descendent of a dependent table.   |
| <b>DFHSM.</b>  | Data Facility Hierarchical Storage Manager.   |
| <b>DFP.</b>  | Data Facility Product (in MVS).   |
| <b>dimension.</b>  | A data category such as time, products, or markets. The elements of a dimension are referred to as members. Dimensions offer a very concise, intuitive way of organizing and selecting data for retrieval, exploration, and analysis. See also <i>dimension table</i> .   |
| <b>dimension table.</b>  | The representation of a dimension in a star schema. Each row in a dimension table represents all of the attributes for a particular member of the dimension. See also <i>dimension</i> , <i>star schema</i> , and <i>star join</i> .  |
| <b>direct access storage device (DASD).</b>                      | A device in which access time is independent of the location of the data.   |
| <b>directory.</b>  | The DB2 system database that contains internal objects such as database descriptors and skeleton cursor tables.   |
| <b>distinct type.</b>  | A user-defined data type that is internally represented as an existing type (its source type), but is considered to be a separate and incompatible type for semantic purposes.  |
| <b>Distributed Computing Environment MVS/ESA™ (DCE MVS/ESA).</b> | A set of technologies that are provided by the Open Software Foundation to implement distributed computing.   |
| <b>distributed data facility (DDF).</b>                          | A set of DB2 components through which DB2 communicates with another RDBMS.  |
| <b>Distributed Relational Database Architecture (DRDA).</b>      | A connection protocol for distributed relational database processing that is used by IBM's relational database products. DRDA includes protocols for communication between an application and a remote relational database management system, and for communication between relational database management systems. |
| <b>DL/I.</b>   | Data Language/I.  |
| <b>double-byte character large object (DBCLOB).</b>              | A sequence of bytes representing double-byte characters where the size of the values can be up to 2 GB. In general, double-byte character large object values are used whenever a double-byte character string might exceed the limits of the VARGRAPHIC type.  |
| <b>double-byte character set (DBCS).</b>                         | A set of characters, which are used by national languages such as Japanese and Chinese, that have more symbols than can be represented by a single byte. Each character is 2 bytes in length. Contrast with <i>single-byte character set</i> and <i>multibyte character set</i> .                                   |
| <b>drain.</b>  | The act of acquiring a locked resource by quiescing access to that object.  |

|                     |  |
|---------------------|--|
| <b>drain lock.</b>  | A lock on a claim class that prevents a claim from occurring.  |
| <b>DRDA.</b>        | Distributed Relational Database Architecture.  |
| <b>DRDA access.</b> | An open method of accessing distributed data that you can use to can connect to another database server to execute packages that were previously bound at the server location. You use the SQL CONNECT statement or an SQL statement with a three-part name to identify the server. Contrast with <i>private protocol access</i> . |
| <b>DSN.</b>         | (1) The default DB2 subsystem name.<br>(2) The name of the TSO command processor of DB2.<br>(3) The first three characters of DB2 module and macro names.  |
| <b>duration.</b>    | A number that represents an interval of time. See <i>date duration</i> , <i>labeled duration</i> , and <i>time duration</i> .  |
| <b>dynamic SQL.</b> | SQL statements that are prepared and executed within an application program while the program is executing. In dynamic SQL, the SQL source is contained in host language variables rather than being coded into the application program. The SQL statement can change several times during the application program's execution     |

## E

|  |   |
|--|---|
| <b>EA-enabled table space.</b>                 | A table space or index space that is enabled for extended addressability and that contains individual partitions (or pieces, for LOB table spaces) that are greater than 4 GB.  |
| <b>EBCDIC.</b>                                 | Extended binary coded decimal interchange code. An encoding scheme that is used to represent character data in the OS/390, MVS, VM, VSE, and OS/400 <sup>®</sup> environments. Contrast with <i>ASCII</i> and <i>Unicode</i> .  |
| <b>EDM pool.</b>                               | A pool of main storage that is used for database descriptors, application plans, authorization cache, application packages, and dynamic statement caching.  |
| <b>EID.</b>                                    | Event identifier.   |
| <b>embedded SQL.</b>                           | SQL statements that are coded within an application program. See <i>static SQL</i> .  |
| <b>enclave.</b>                                | In Language Environment, an independent collection of routines, one of which is designated as the main routine. An enclave is similar to a program or run unit.   |
| <b>EOM.</b>                                    | End of memory.  |
| <b>EOT.</b>                                    | End of task.  |
| <b>Equi join.</b>                              | A join operation in which the join-condition has the form <i>expression = expression</i> .  |
| <b>error page range.</b>                       | A range of pages that are considered to be physically damaged. DB2 does not allow users to access any pages that fall within this range.  |
| <b>ESDS.</b>                                   | Entry sequenced data set.   |
| <b>ESMT.</b>                                   | External subsystem module table (in IMS).   |
| <b>EUR.</b>                                    | IBM European Standards.   |
| <b>exception table.</b>                        | A table that holds rows that violate referential constraints or table check constraints that the CHECK DATA utility finds.  |
| <b>exclusive lock.</b>                         | A lock that prevents concurrently executing application processes from reading or changing data. Contrast with <i>share lock</i> .  |
| <b>exit routine.</b>                           | A user-written (or IBM-provided default) program that receives control from DB2 to perform specific functions. Exit routines run as extensions of DB2.  |
| <b>expression.</b>                             | An operand or a collection of operators and operands that yields a single value.  |
| <b>extended recovery facility (XRF).</b>       | A facility that minimizes the effect of failures in MVS, VTAM, the host processor, or high-availability applications during sessions between high-availability applications and designated terminals. This facility provides an alternative subsystem to take over sessions from the failing subsystem. |
| <b>external function.</b>                      | A function for which the body is written in a programming language that takes scalar argument values and produces a scalar result for each invocation. Contrast with <i>sourced function</i> , <i>built-in function</i> , and <i>SQL function</i> .   |
| <b>external routine.</b>                       | A user-defined function or stored procedure that is based on code that is written in an external programming language.  |
| <b>External subsystem module table (ESMT).</b> |   |

The name of the external subsystem module table, which specifies which attachment modules must be loaded by IMS.

## F

- fallback.** The process of returning to a previous release of DB2 after attempting or completing migration to a current release.
- field procedure.** A user-written exit routine that is designed to receive a single value and transform (encode or decode) it in any way the user can specify.
- filter factor.** A number between zero and one that estimates the proportion of rows in a table for which a predicate is true.
- fixed-length string.** A character or graphic string whose length is specified and cannot be changed. Contrast with *varying-length string*.
- foreign key.** A column or set of columns in a dependent table of a constraint relationship. The key must have the same number of columns, with the same descriptions, as the primary key of the parent table. Each foreign key value must either match a parent key value in the related parent table or be null.
- forward log recovery.** The third phase of restart processing during which DB2 processes the log in a forward direction to apply all REDO log records.
- free space.** The total amount of unused space in a page; that is, the space that is not used to store records or control information is free space.
- full outer join.** The result of a join operation that includes the matched rows of both tables that are being joined and preserves the unmatched rows of both tables. See also *join*.
- function.** A mapping, embodied as a program (the function body), invocable by means of zero or more input values (arguments), to a single value (the result). See also *column function* and *scalar function*. Functions can be user-defined, built-in, or generated by DB2. (See *built-in function*, *cast function*, *external function*, *sourced function*, *SQL function*, and *user-defined function*.)

## G

|  |  |
|--|--|
| <b>GB.</b>                               | Gigabyte (1 073 741 824 bytes).  |
| <b>GBP.</b>                              | Group buffer pool.   |
| <b>generalized trace facility (GTF).</b> | An MVS service program that records significant system events such as I/O interrupts, SVC interrupts, program interrupts, or external interrupts.                              |
| <b>generic resource name.</b>            | A name that VTAM uses to represent several application programs that provide the same function in order to handle session distribution and balancing in a Sysplex environment. |
| <b>getpage.</b>                          | An operation in which DB2 accesses a data page.  |
| <b>GIMSMP.</b>                           | The load module name for the System Modification Program/Extended, a basic tool for installing, changing, and controlling changes to programming systems.                      |
| <b>graphic string.</b>                   | A sequence of DBCS characters.   |
| <b>gross lock.</b>                       | The <i>shared</i> , <i>update</i> , or <i>exclusive</i> mode locks on a table, partition, or table space.  |
| <b>group buffer pool (GBP).</b>          | A coupling facility cache structure that is used by a data sharing group to cache data and to ensure that the data is consistent for all members.                              |
| <b>GTF.</b>                              | Generalized trace facility.  |

## H

|                            |   |
|----------------------------|---|
| <b>help panel.</b>         | A screen of information presenting tutorial text to assist a user at the terminal.  |
| <b>hiperspace.</b>         | A range of up to 2 GB of contiguous virtual storage addresses that a program can use as a buffer. Like a data space, a hiperspace can hold user data; it does not contain common areas or system data. Unlike an address space or a data space, data in a hiperspace is not directly addressable. To manipulate data in a hiperspace, bring the data into the address space in 4-KB blocks. |
| <b>home address space.</b> | The area of storage that MVS currently recognizes as <i>dispatched</i> .  |
| <b>host language.</b>      | A programming language in which you can embed SQL statements.   |
| <b>host program.</b>       | An application program that is written in a host language and that contains embedded SQL statements.  |
| <b>host structure.</b>     | In an application program, a structure that is referenced by embedded SQL statements.   |
| <b>host variable.</b>      | In an application program, an application variable that is referenced by embedded SQL statements.   |
| <b>HSM.</b>                | Hierarchical storage manager.   |

## I

|                                 |   |
|---------------------------------|---|
| <b>ICF.</b>                     | Integrated catalog facility.  |
| <b>IDCAMS.</b>                  | An IBM program that is used to process access method services commands. It can be invoked as a job or jobstep, from a TSO terminal, or from within a user's application program.  |
| <b>IDCAMS LISTCAT.</b>          | A facility for obtaining information that is contained in the access method services catalog.   |
| <b>identify.</b>                | A request that an attachment service program in an address space that is separate from DB2 issues via the MVS subsystem interface to inform DB2 of its existence and to initiate the process of becoming connected to DB2.  |
| <b>identity column.</b>         | A column that provides a way for DB2 to automatically generate a numeric value for each row. The generated values are unique if cycling is not used. Identity columns are defined with the AS IDENTITY clause. Uniqueness of values can be ensured by defining a single-column unique index using the identity column. A table can have no more than one identity column. |
| <b>IFCID.</b>                   | Instrumentation facility component identifier.  |
| <b>IFI.</b>                     | Instrumentation facility interface.   |
| <b>IFI call.</b>                | An invocation of the instrumentation facility interface (IFI) by means of one of its defined functions.   |
| <b>IFP.</b>                     | IMS Fast Path.  |
| <b>image copy.</b>              | An exact reproduction of all or part of a table space. DB2 provides utility programs to make full image copies (to copy the entire table space) or incremental image copies (to copy only those pages that have been modified since the last image copy).   |
| <b>IMS.</b>                     | Information Management System.  |
| <b>IMS attachment facility.</b> | A DB2 subcomponent that uses MVS subsystem interface (SSI) protocols and cross-memory linkage to process requests from IMS to DB2 and to coordinate resource commitment.  |
| <b>IMS DB.</b>                  | Information Management System Database.   |
| <b>IMS TM.</b>                  | Information Management System Transaction Manager.  |
| <b>in-abort.</b>                | A status of a unit of recovery. If DB2 fails after a unit of recovery begins to be rolled back, but before the process is completed, DB2 continues to back out the changes during restart.  |
| <b>in-commit.</b>               | A status of a unit of recovery. If DB2 fails after beginning its phase 2 commit processing, it "knows," when restarted, that changes made to data are consistent. Such units of recovery are termed <i>in-commit</i> .  |
| <b>independent.</b>             | An object (row, table, or table space) that is neither a parent nor a dependent of another object.  |
| <b>index.</b>                   | A set of pointers that are logically ordered by the values of a key. Indexes can provide faster access to data and can enforce uniqueness on the rows in a table.   |
| <b>index key.</b>               | The set of columns in a table that is used to determine the order of index entries.   |

|   |   |
|---|---|
| <b>index partition.</b>                                       | A VSAM data set that is contained within a partitioning index space.  |
| <b>index space.</b>   | A page set that is used to store the entries of one index.  |
| <b>indicator variable.</b>                                    | A variable that is used to represent the null value in an application program. If the value for the selected column is null, a negative value is placed in the indicator variable.  |
| <b>indoubt.</b>   | A status of a unit of recovery. If DB2 fails after it has finished its phase 1 commit processing and before it has started phase 2, only the commit coordinator knows if an individual unit of recovery is to be committed or rolled back. At emergency restart, if DB2 lacks the information it needs to make this decision, the status of the unit of recovery is <i>indoubt</i> until DB2 obtains this information from the coordinator. More than one unit of recovery can be indoubt at restart. |
| <b>indoubt resolution.</b>                                    | The process of resolving the status of an indoubt logical unit of work to either the committed or the rollback state.   |
| <b>inflight.</b>  | A status of a unit of recovery. If DB2 fails before its unit of recovery completes phase 1 of the commit process, it merely backs out the updates of its unit of recovery at restart. These units of recovery are termed <i>inflight</i> .  |
| <b>inner join.</b>  | The result of a join operation that includes only the matched rows of both tables being joined. See also <i>join</i> .  |
| <b>inoperative package.</b>                                   | A package that cannot be used because one or more user-defined functions or procedures that the package depends on were dropped. Such a package must be explicitly rebound. Contrast with <i>invalid package</i> .  |
| <b>install.</b>   | The process of preparing a DB2 subsystem to operate as an MVS subsystem.  |
| <b>installation verification scenario.</b>                    | A sequence of operations that exercises the main DB2 functions and tests whether DB2 was correctly installed.   |
| <b>instrumentation facility component identifier (IFCID).</b> | A value that names and identifies a trace record of an event that can be traced. As a parameter on the START TRACE and MODIFY TRACE commands, it specifies that the corresponding event is to be traced.  |
| <b>instrumentation facility interface (IFI).</b>              | A programming interface that enables programs to obtain online trace data about DB2, to submit DB2 commands, and to pass data to DB2.   |
| <b>Interactive System Productivity Facility (ISPF).</b>       | An IBM licensed program that provides interactive dialog services.  |
| <b>intermediate database server.</b>                          | The target of a request from a local application or a remote application requester that is forwarded to another database server. In the DB2 environment, the remote request is forwarded transparently to another database server if the object that is referenced by a three-part name does not reference the local location.  |
| <b>internal resource lock manager (IRLM).</b>                 | An MVS subsystem that DB2 uses to control communication and database locking.   |
| <b>invalid package.</b>                                       | A package that depends on an object (other than a user-defined function) that is dropped. Such a package is implicitly rebound on invocation. Contrast with <i>inoperative package</i> .  |
| <b>IRLM.</b>  | Internal resource lock manager.   |
| <b>ISO.</b>   | International Standards Organization.   |

- isolation level.** The degree to which a unit of work is isolated from the updating operations of other units of work. See also *cursor stability*, *read stability*, *repeatable read*, and *uncommitted read*.
- ISPF.** Interactive System Productivity Facility.
- ISPF/PDF.** Interactive System Productivity Facility/Program Development Facility.

## J

**Japanese Industrial Standards Committee (JISC).**

An organization that issues standards for coding character sets.

**Java Archive (JAR).**

A file format that is used for aggregating many files into a single file.

**JCL.**

Job control language.

**JES.**

MVS Job Entry Subsystem.

**JIS.**

Japanese Industrial Standard.

**job control language (JCL).**

A control language that is used to identify a job to an operating system and to describe the job's requirements.

**Job Entry Subsystem (JES).**

An IBM licensed program that receives jobs into the system and processes all output data that is produced by the jobs.

**join.**

A relational operation that allows retrieval of data from two or more tables based on matching column values. See also *equijoin*, *full outer join*, *inner join*, *left outer join*, *outer join*, and *right outer join*.

## K

- KB.** Kilobyte (1024 bytes).
- Kerberos.** A network authentication protocol that is designed to provide strong authentication for client/server applications by using secret-key cryptography.
- Kerberos ticket.** A transparent application mechanism that transmits the identity of an initiating principal to its target. A simple ticket contains the principal's identity, a session key, a timestamp, and other information, which is sealed using the target's secret key.
- key.** A column or an ordered collection of columns identified in the description of a table, index, or referential constraint.
- key-sequenced data set (KSDS).** A VSAM file or data set whose records are loaded in key sequence and controlled by an index.
- KSDS.** Key-sequenced data set.

## L

|                               |   |
|-------------------------------|---|
| <b>labeled duration.</b>      | A number that represents a duration of years, months, days, hours, minutes, seconds, or microseconds.   |
| <b>large object (LOB).</b>    | A sequence of bytes representing bit data, single-byte characters, double-byte characters, or a mixture of single- and double-byte characters. A LOB can be up to 2 GB.1 byte in length. See also <i>BLOB</i> , <i>CLOB</i> , and <i>DBCLOB</i> . |
| <b>latch.</b>                 | A DB2 internal mechanism for controlling concurrent events or the use of system resources.  |
| <b>LCID.</b>                  | Log control interval definition.  |
| <b>LDS.</b>                   | Linear data set.  |
| <b>leaf page.</b>             | A page that contains pairs of keys and RIDs and that points to actual data. Contrast with <i>nonleaf page</i> .   |
| <b>left outer join.</b>       | The result of a join operation that includes the matched rows of both tables that are being joined, and that preserves the unmatched rows of the first table. See also <i>join</i> .  |
| <b>linear data set (LDS).</b> | A VSAM data set that contains data but no control information. A linear data set can be accessed as a byte-addressable string in virtual storage.   |
| <b>linkage editor.</b>        | A computer program for creating load modules from one or more object modules or load modules by resolving cross references among the modules and, if necessary, adjusting addresses.  |
| <b>link-edit.</b>             | The action of creating a loadable computer program using a linkage editor.  |
| <b>L-lock.</b>                | Logical lock.   |
| <b>load module.</b>           | A program unit that is suitable for loading into main storage for execution. The output of a linkage editor.  |
| <b>LOB.</b>                   | Large object.   |
| <b>LOB lock.</b>              | A lock on a LOB value.  |
| <b>LOB table space.</b>       | A table space that contains all the data for a particular LOB column in the related base table.   |
| <b>local subsystem.</b>       | The unique RDBMS to which the user or application program is directly connected (in the case of DB2, by one of the DB2 attachment facilities).  |
| <b>lock.</b>                  | A means of controlling concurrent events or access to data. DB2 locking is performed by the IRLM.   |
| <b>lock duration.</b>         | The interval over which a DB2 lock is held.   |
| <b>lock escalation.</b>       | The promotion of a lock from a row, page, or LOB lock to a table space lock because the number of page locks that are concurrently held on a given resource exceeds a preset limit.   |
| <b>locking.</b>               | The process by which the integrity of data is ensured. Locking prevents concurrent users from accessing inconsistent data.  |
| <b>lock mode.</b>             | A representation for the type of access that concurrently running programs can have to a resource that a DB2 lock is holding.   |

|   |   |
|---|---|
| <b>lock object.</b>                             | The resource that is controlled by a DB2 lock.  |
| <b>lock promotion.</b>                          | The process of changing the size or mode of a DB2 lock to a higher level.   |
| <b>lock size.</b>                               | The amount of data controlled by a DB2 lock on table data; the value can be a row, a page, a LOB, a partition, a table, or a table space.   |
| <b>log.</b>                                     | A collection of records that describe the events that occur during DB2 execution and that indicate their sequence. The information thus recorded is used for recovery in the event of a failure during DB2 execution.   |
| <b>logical index partition.</b>                 | The set of all keys that reference the same data partition.   |
| <b>logical lock (L-lock).</b>                   | The lock type that transactions use to control intra- and inter-DB2 data concurrency between transactions. Contrast with <i>physical lock (P-lock)</i> .  |
| <b>logical recovery pending (LRECP).</b>        | The state in which the data and the index keys that reference the data are inconsistent.  |
| <b>logical unit.</b>                            | An access point through which an application program accesses the SNA network in order to communicate with another application program.   |
| <b>logical unit of work (LUW).</b>              | The processing that a program performs between synchronization points.  |
| <b>logical unit of work identifier (LUWID).</b> | A name that uniquely identifies a thread within a network. This name consists of a fully-qualified LU network name, an LUW instance number, and an LUW sequence number.   |
| <b>log initialization.</b>                      | The first phase of restart processing during which DB2 attempts to locate the current end of the log.   |
| <b>log record sequence number (LRSN).</b>       | A number that DB2 generates and associates with each log record. DB2 also uses the LRSN for page versioning. The LRSNs that a particular DB2 data sharing group generates form a strictly increasing sequence for each DB2 log and a strictly increasing sequence for each page across the DB2 group. |
| <b>log truncation.</b>                          | A process by which an explicit starting RBA is established. This RBA is the point at which the next byte of log data is to be written.  |
| <b>long string.</b>                             | A string whose actual length, or a varying-length string whose maximum length, is greater than 255 bytes or 127 double-byte characters. Any LOB column, LOB host variable, or expression that evaluates to a LOB is considered a long string.   |
| <b>LRECP.</b>                                   | Logical recovery pending.   |
| <b>LRH.</b>                                     | Log record header.  |
| <b>LRSN.</b>                                    | Log record sequence number.   |
| <b>LUW.</b>                                     | Logical unit of work.   |
| <b>LUWID.</b>                                   | #Logical unit of work identifier.   |

## M

|  |   |
|--|---|
| <b>materialize.</b>                    | (1) The process of putting rows from a view or nested table expression into a work file for additional processing by a query.<br>(2) The placement of a LOB value into contiguous storage. Because LOB values can be very large, DB2 avoids materializing LOB data until doing so becomes absolutely necessary. |
| <b>MB.</b>                             | Megabyte (1 048 576 bytes).   |
| <b>migration.</b>                      | The process of converting a DB2 subsystem with a previous release of DB2 to an updated or current release. In this process, you can acquire the functions of the updated or current release without losing the data you created on the previous release.  |
| <b>mixed data string.</b>              | A character string that can contain both single-byte and double-byte characters.  |
| <b>MLPA.</b>                           | Modified link pack area.  |
| <b>MODEENT.</b>                        | A VTAM macro instruction that associates a logon mode name with a set of parameters representing session protocols. A set of MODEENT macro instructions defines a logon mode table.   |
| <b>mode name.</b>                      | A VTAM name for the collection of physical and logical characteristics and attributes of a session.   |
| <b>MPP.</b>                            | Message processing program (in IMS).  |
| <b>MSS.</b>                            | Mass Storage Subsystem.   |
| <b>MTO.</b>                            | Master terminal operator.   |
| <b>multibyte character set (MBCS).</b> | A character set that represents single characters with more than a single byte. Contrast with <i>single-byte character set</i> and <i>double-byte character set</i> . See also <i>Unicode</i> .   |
| <b>multisite update.</b>               | Distributed relational database processing in which data is updated in more than one location within a single unit of work.   |
| <b>must-complete.</b>                  | A state during DB2 processing in which the entire operation must be completed to maintain data integrity.   |
| <b>MVS.</b>                            | Multiple Virtual Storage.   |
| <b>MVS/ESA.</b>                        | Multiple Virtual Storage/Enterprise Systems Architecture.   |

## N

|                                      |  |
|--------------------------------------|--|
| <b>nested table expression.</b>      | A fullselect in a FROM clause (surrounded by parentheses).   |
| <b>network identifier (NID).</b>     | The network ID that is assigned by IMS or CICS, or if the connection type is RRSAF, the OS/390 RRS unit of recovery ID (URID).                   |
| <b>NID.</b>                          | Network ID.  |
| <b>nonleaf page.</b>                 | A page that contains keys and page numbers of other pages in the index (either leaf or nonleaf pages). Nonleaf pages never point to actual data. |
| <b>nonpartitioning index.</b>        | Any index that is not a partitioning index.  |
| <b>NPI</b>                           | see <i>nonpartitioning index</i> .   |
| <b>NRE.</b>                          | Network recovery element.  |
| <b>NUL.</b>                          | In C, a single character that denotes the end of the string.   |
| <b>null.</b>                         | A special value that indicates the absence of information.   |
| <b>NUL-terminated host variable.</b> | A varying-length host variable in which the end of the data is indicated by the presence of a NUL terminator.                                    |
| <b>NUL terminator.</b>               | In C, the value that indicates the end of a string. For character strings, the NUL terminator is X'00'.  |

## O

**OASN (origin application schedule number).**

In IMS, a 4-byte number that is assigned sequentially to each IMS schedule since the last cold start of IMS. The OASN is used as an identifier for a unit of work. In an 8-byte format, the first 4 bytes contain the schedule number and the last 4 bytes contain the number of IMS sync points (*commit points*) during the current schedule. The OASN is part of the NID for an IMS connection.

**OBID.**

Data object identifier.

**originating task.**

In a parallel group, the primary agent that receives data from other execution units (referred to as *parallel tasks*) that are executing portions of the query in parallel.

**OS/390.**

Operating System/390.

**OS/390 OpenEdition Distributed Computing Environment (OS/390 OE DCE).**

A set of technologies that are provided by the Open Software Foundation to implement distributed computing.

**outer join.**

The result of a join operation that includes the matched rows of both tables that are being joined and preserves some or all of the unmatched rows of the tables that are being joined. See also *join*.

## P

|                                    |  |
|------------------------------------|--|
| <b>package.</b>                    | An object containing a set of SQL statements that have been statically bound and that is available for processing. A package is sometimes also called an <i>application package</i> .  |
| <b>package list.</b>               | An ordered list of package names that may be used to extend an application plan.   |
| <b>package name.</b>               | The name of an object that is created by a BIND PACKAGE or REBIND PACKAGE command. The object is a bound version of a database request module (DBRM). The name consists of a location name, a collection ID, a package ID, and a version ID.   |
| <b>page.</b>                       | A unit of storage within a table space (4 KB, 8 KB, 16 KB, or 32 KB) or index space (4 KB). In a table space, a page contains one or more rows of a table. In a LOB table space, a LOB value can span more than one page, but no more than one LOB value is stored on a page.                        |
| <b>page set.</b>                   | Another way to refer to a table space or index space. Each page set consists of a collection of VSAM data sets.  |
| <b>Parallel group.</b>             | A set of consecutive operations that executed in parallel and that have the same number of parallel tasks.   |
| <b>parallel I/O processing.</b>    | A form of I/O processing in which DB2 initiates multiple concurrent requests for a single user query and performs I/O processing concurrently (in <i>parallel</i> ) on multiple data partitions.   |
| <b>Parallel Sysplex.</b>           | A set of MVS systems that communicate and cooperate with each other through certain multisystem hardware components and software services to process customer workloads.   |
| <b>parallel task.</b>              | The execution unit that is dynamically created to process a query in parallel. It is implemented by an MVS service request block.  |
| <b>parent row.</b>                 | A row whose primary key value is the foreign key value of a dependent row.   |
| <b>parent table.</b>               | A table whose primary key is referenced by the foreign key of a dependent table.   |
| <b>parent table space.</b>         | A table space that contains a parent table. A table space containing a dependent of that table is a dependent table space.   |
| <b>participant.</b>                | An entity other than the commit coordinator that takes part in the commit process. The term participant is synonymous with <i>agent</i> in SNA.  |
| <b>partition.</b>                  | A portion of a page set. Each partition corresponds to a single, independently extendable data set. Partitions can be extended to a maximum size of 1, 2, or 4 GB, depending on the number of partitions in the partitioned page set. All partitions of a given page set have the same maximum size. |
| <b>partitioning index.</b>         | The index that represents the partitioning values  |
| <b>partitioned data set (PDS).</b> | A data set in direct access storage that is divided into partitions, which are called members. Each partition can contain a program, part of a program, or data. The term partitioned data set is synonymous with program library.   |
| <b>partitioned page set.</b>       | A partitioned table space or an index space. Header pages, space map pages, data pages, and index pages reference data only within the scope of the partition.   |

|                                  |   |
|----------------------------------|---|
| <b>partitioned table space.</b>  | A table space that is subdivided into parts (based on index key range), each of which can be processed independently by utilities.  |
| <b>partner logical unit.</b>     | An access point in the SNA network that is connected to the local DB2 subsystem by way of a VTAM conversation.  |
| <b>PCT.</b>                      | Program control table (in CICS).  |
| <b>PDS.</b>                      | Partitioned data set.   |
| <b>PI</b>                        | see <i>partitioning index</i> .   |
| <b>piece.</b>                    | A data set of a nonpartitioned page set.  |
| <b>physical consistency.</b>     | The state of a page that is not in a partially changed state.   |
| <b>plan.</b>                     | See <i>application plan</i> .   |
| <b>plan allocation.</b>          | The process of allocating DB2 resources to a plan in preparation for execution.   |
| <b>plan name.</b>                | The name of an application plan.  |
| <b>plan segmentation.</b>        | The dividing of each plan into sections. When a section is needed, it is independently brought into the EDM pool.   |
| <b>PLT.</b>                      | Program list table (in CICS).   |
| <b>point of consistency.</b>     | A time when all recoverable data that an application accesses is consistent with other data. The term point of consistency is synonymous with <i>sync point</i> or <i>commit point</i> .  |
| <b>postponed abort UR.</b>       | A unit of recovery that was inflight or in-abort, was interrupted by system failure or cancellation, and did not complete backout during restart.   |
| <b>PPT.</b>                      | (1) Processing program table (in CICS).<br>(2) Program properties table (in MVS).   |
| <b>precompilation.</b>           | A processing of application programs containing SQL statements that takes place before compilation. SQL statements are replaced with statements that are recognized by the host language compiler. Output from this precompilation includes source code that can be submitted to the compiler and the database request module (DBRM) that is input to the bind process. |
| <b>predicate.</b>                | An element of a search condition that expresses or implies a comparison operation. <b>prefix.</b> A code at the beginning of a message or record.   |
| <b>primary authorization ID.</b> | The authorization ID used to identify the application process to DB2.   |
| <b>primary index.</b>            | An index that enforces the uniqueness of a primary key.   |
| <b>primary key.</b>              | In a relational database, a unique, nonnull key that is part of the definition of a table. A table cannot be defined as a parent unless it has a unique key or primary key.   |
| <b>principal.</b>                | An entity that can communicate securely with another entity. In Kerberos, principals are represented as entries in the Kerberos registry database and include users, servers, computers, and others.  |
| <b>principal name.</b>           | The name by which a principal is known to the DCE security services.  |
| <b>private connection.</b>       | A communications connection that is specific to DB2.  |
| <b>private protocol access.</b>  | A method of accessing distributed data by which you can direct a query to another DB2 system. Contrast with <i>DRDA access</i> .  |

- private protocol connection.** A DB2 private connection of the application process. See also *private connection*.
- privilege.** The capability of performing a specific function, sometimes on a specific object. The term includes:  
**explicit privileges**, which have names and are held as the result of SQL GRANT and REVOKE statements. For example, the SELECT privilege.  
**implicit privileges**, which accompany the ownership of an object, such as the privilege to drop a synonym one owns, or the holding of an authority, such as the privilege of SYSADM authority to terminate any utility job.
- privilege set.** For the installation SYSADM ID, the set of all possible privileges. For any other authorization ID, the set of all privileges that are recorded for that ID in the DB2 catalog.
- process.** In DB2, the unit to which DB2 allocates resources and locks. Sometimes called an *application process*, a process involves the execution of one or more programs. The execution of an SQL statement is always associated with some process. The means of initiating and terminating a process are dependent on the environment.
- program.** A single compilable collection of executable statements in a programming language.
- program temporary fix (PTF).** A solution or bypass of a problem that is diagnosed as a result of a defect in a current unaltered release of a licensed program. An authorized program analysis report (APAR) fix is corrective service for an existing problem. A PTF is preventive service for problems that might be encountered by other users of the product. A PTF is *temporary*, because a permanent fix is usually not incorporated into the product until its next release.
- protected conversation.** A VTAM conversation that supports two-phase commit flows.
- PTF.** Program temporary fix.

## Q

- QMF.** Query Management Facility.
- QSAM.** Queued sequential access method.
- query block.** The part of a query that is represented by one of the FROM clauses. Each FROM clause can have multiple query blocks, depending on DB2's internal processing of the query.
- query CP parallelism.** Parallel execution of a single query, which is accomplished by using multiple tasks. See also *Sysplex query parallelism*.
- query I/O parallelism.** Parallel access of data, which is accomplished by triggering multiple I/O requests within a single query.
- queued sequential access method (QSAM).** An extended version of the basic sequential access method (BSAM). When this method is used, a queue of data blocks is formed. Input data blocks await processing, and output data blocks await transfer to auxiliary storage or to an output device.

## R

|                                      |   |
|--------------------------------------|---|
| <b>RACF.</b>                         | Resource Access Control Facility, which is a component of the SecureWay Security Server for OS/390.   |
| <b>RAMAC.</b>                        | IBM family of enterprise disk storage system products.  |
| <b>RBA.</b>                          | Relative byte address.  |
| <b>RCT.</b>                          | Resource control table (in CICS attachment facility).   |
| <b>RDB.</b>                          | Relational database.  |
| <b>RDBMS.</b>                        | Relational database management system.  |
| <b>RDBNAM.</b>                       | Relational database name.   |
| <b>RDF.</b>                          | Record definition field.  |
| <b>read stability (RS).</b>          | An isolation level that is similar to repeatable read but does not completely isolate an application process from all other concurrently executing application processes. Under level RS, an application that issues the same query more than once might read additional rows that were inserted and committed by a concurrently executing application process. |
| <b>rebind.</b>                       | The creation of a new application plan for an application program that has been bound previously. If, for example, you have added an index for a table that your application accesses, you must rebind the application in order to take advantage of that index.  |
| <b>record.</b>                       | The storage representation of a row or other data.  |
| <b>record identifier (RID).</b>      | A unique identifier that DB2 uses internally to identify a row of data in a table stored as a record. Compare with <i>row ID</i> .  |
| <b>record identifier (RID) pool.</b> | An area of main storage above the 16-MB line that is reserved for sorting record identifiers during list prefetch processing.   |
| <b>recovery.</b>                     | The process of rebuilding databases after a system failure.   |
| <b>recovery log.</b>                 | A collection of records that describes the events that occur during DB2 execution and indicates their sequence. The recorded information is used for recovery in the event of a failure during DB2 execution.   |
| <b>recovery pending (RECP).</b>      | A condition that prevents SQL access to a table space that needs to be recovered.   |
| <b>recovery token.</b>               | An identifier for an element that is used in recovery (for example, <i>NID</i> or <i>URID</i> ).  |
| <b>RECP.</b>                         | Recovery pending.   |
| <b>redo.</b>                         | A state of a unit of recovery that indicates that changes are to be reapplied to the DASD media to ensure data integrity.   |
| <b>referential constraint.</b>       | The requirement that nonnull values of a designated foreign key are valid only if they equal values of the primary key of a designated table.   |
| <b>referential integrity.</b>        | The state of a database in which all values of all foreign keys are valid. Maintaining referential integrity requires the enforcement of referential  |

|   |  |
|---|--|
|   | constraints on all operations that change the data in a table upon which the referential constraints are defined.  |
| <b>referential structure.</b>                         | A set of tables and relationships that includes at least one table and, for every table in the set, all the relationships in which that table participates and all the tables to which it is related.  |
| <b>registry.</b>                                      | See <i>registry database</i> .   |
| <b>registry database.</b>                             | A database of security information about principals, groups, organizations, accounts, and security policies.   |
| <b>relational database (RDB).</b>                     | A database that can be perceived as a set of tables and manipulated in accordance with the relational model of data.   |
| <b>relational database management system (RDBMS).</b> | A collection of hardware and software that organizes and provides access to a relational database.   |
| <b>relational database name (RDBNAM).</b>             | A unique identifier for an RDBMS within a network. In DB2, this must be the value in the LOCATION column of table SYSIBM.LOCATIONS in the CDB. DB2 publications refer to the name of another RDBMS as a LOCATION value or a location name.   |
| <b>relationship.</b>                                  | A defined connection between the rows of a table or the rows of two tables. A relationship is the internal representation of a referential constraint.   |
| <b>relative byte address (RBA).</b>                   | The offset of a data record or control interval from the beginning of the storage space that is allocated to the data set or file to which it belongs.   |
| <b>remigration.</b>                                   | The process of returning to a current release of DB2 following a fallback to a previous release. This procedure constitutes another migration process.   |
| <b>remote attach request.</b>                         | A request by a remote location to attach to the local DB2 subsystem. Specifically, the request that is sent is an SNA Function Management Header 5.  |
| <b>remote subsystem.</b>                              | Any RDBMS, except the <i>local subsystem</i> , with which the user or application can communicate. The subsystem need not be remote in any physical sense, and might even operate on the same processor under the same MVS system.   |
| <b>reoptimization.</b>                                | The DB2 process of reconsidering the access path of an SQL statement at run time; during reoptimization, DB2 uses the values of host variables, parameter markers, or special registers.   |
| <b>REORG pending (REORP).</b>                         | A condition that restricts SQL access and most utility access to an object that must be reorganized.   |
| <b>REORP.</b>   | REORG pending.   |
| <b>repeatable read (RR).</b>                          | The isolation level that provides maximum protection from other executing application programs. When an application program executes with repeatable read protection, rows referenced by the program cannot be changed by other programs until the program reaches a commit point. |
| <b>request commit.</b>                                | The vote that is submitted to the prepare phase if the participant has modified data and is prepared to commit or roll back.   |
| <b>requester.</b>                                     | The source of a request to access data at a remote server. In the DB2 environment, the requester function is provided by the distributed data facility.  |
| <b>resource allocation.</b>                           | The part of plan allocation that deals specifically with the database resources.   |
| <b>resource control table (RCT).</b>                  | A construct of the CICS attachment facility, created by site-provided macro parameters, that defines authorization and access attributes for transactions or transaction groups.   |

|  |   |
|--|---|
| <b>resource definition online.</b>         | A CICS feature that you use to define CICS resources online without assembling tables.  |
| <b>resource limit facility (RLF).</b>      | A portion of DB2 code that prevents dynamic manipulative SQL statements from exceeding specified time limits. The resource limit facility is sometimes called the governor.   |
| <b>resource limit specification table.</b> | A site-defined table that specifies the limits to be enforced by the resource limit facility.   |
| <b>restart pending (RESTP).</b>            | A restrictive state of a page set or partition that indicates that restart (backout) work needs to be performed on the object. All access to the page set or partition is denied except for access by the: v RECOVER POSTPONED command v Automatic online backout (which DB2 invokes after restart if the system parameter LBACKOUT=AUTO) |
| <b>result table.</b>                       | The set of rows that are specified by a SELECT statement.   |
| <b>RID.</b>                                | Record identifier.  |
| <b>RID pool.</b>                           | Record identifier pool.   |
| <b>right outer join.</b>                   | The result of a join operation that includes the matched rows of both tables that are being joined and preserves the unmatched rows of the second join operand. See also <i>join</i> .  |
| <b>RLF.</b>                                | Resource limit facility.  |
| <b>RMID.</b>                               | Resource manager identifier.  |
| <b>RO.</b>                                 | Read-only access.   |
| <b>rollback.</b>                           | The process of restoring data changed by SQL statements to the state at its last commit point. All locks are freed. Contrast with <i>commit</i> .   |
| <b>root page.</b>                          | The page of an index page set that follows the first index space map page. A root page is the highest level (or the beginning point) of the index.  |
| <b>routine.</b>                            | A term that refers to either a user-defined function or a stored procedure.   |
| <b>row.</b>                                | The horizontal component of a table. A row consists of a sequence of values, one for each column of the table.  |
| <b>ROWID.</b>                              | Row identifier.   |
| <b>row identifier (ROWID).</b>             | A value that uniquely identifies a row. This value is stored with the row and never changes.  |
| <b>row lock.</b>                           | A lock on a single row of data.   |
| <b>RRE.</b>                                | Residual recovery entry (in IMS).   |
| <b>RRSAF.</b>                              | Recoverable Resource Manager Services attachment facility. RRSAF is a DB2 subcomponent that uses OS/390 Transaction Management and Recoverable Resource Manager Services to coordinate resource commitment between DB2 and all other resource managers that also use OS/390 RRS in an OS/390 system.                                      |
| <b>RS.</b>                                 | Read stability.   |
| <b>RTT.</b>                                | Resource translation table.   |

## S

|                                     |   |
|-------------------------------------|---|
| <b>savepoint.</b>                   | A named entity that represents the state of data and schemas at a particular point in time within a unit of work. SQL statements exist to set a savepoint, release a savepoint, and restore data and schemas to the state that the savepoint represents. The restoration of data and schemas to a savepoint is usually referred to as <i>rolling back to a savepoint</i> .  |
| <b>scalar function.</b>             | An SQL operation that produces a single value from another value and is expressed as a function name, followed by a list of arguments that are enclosed in parentheses. Contrast with <i>column function</i> .  |
| <b>schema.</b>                      | A logical grouping for user-defined functions, distinct types, triggers, and stored procedures. When an object of one of these types is created, it is assigned to one schema, which is determined by the name of the object. For example, the following statement creates a distinct type T in schema C: <code>CREATEDISTINCT TYPE C.T ...</code>  |
| <b>SDWA.</b>                        | System diagnostic work area.  |
| <b>search condition.</b>            | A criterion for selecting rows from a table. A search condition consists of one or more predicates.   |
| <b>secondary authorization ID.</b>  | An authorization ID that has been associated with a primary authorization ID by an authorization exit routine.  |
| <b>section.</b>                     | The segment of a plan or package that contains the executable structures for a single SQL statement. For most SQL statements, one section in the plan exists for each SQL statement in the source program. However, for cursor-related statements, the DECLARE, OPEN, FETCH, and CLOSE statements reference the same section because, they each refer to the SELECT statement that is named in the DECLARE CURSOR statement. SQL statements such as COMMIT, ROLLBACK, and some SET statements do not use a section. |
| <b>segmented table space.</b>       | A table space that is divided into equal-sized groups of pages called segments. Segments are assigned to tables so that rows of different tables are never stored in the same segment.  |
| <b>self-referencing constraint.</b> | A referential constraint that defines a relationship in which a table is a dependent of itself.   |
| <b>self-referencing table.</b>      | A table with a self-referencing constraint.   |
| <b>sequential data set.</b>         | A non-DB2 data set whose records are organized on the basis of their successive physical positions, such as on magnetic tape. Several of the DB2 database utilities require sequential data sets.   |
| <b>sequential prefetch.</b>         | A mechanism that triggers consecutive asynchronous I/O operations. Pages are fetched before they are required, and several pages are read with a single I/O operation.  |
| <b>server.</b>                      | The target of a request from a remote requester. In the DB2 environment, the server function is provided by the distributed data facility, which is used to access DB2 data from remote applications.   |
| <b>service class.</b>               | An eight-character identifier that is used by MVS Workload Manager to associate customer performance goals with a particular DDF thread or stored procedure. A service class is also used to classify work on parallelism assistants.   |
| <b>session.</b>                     | A link between two nodes in a VTAM network.   |

|  |  |
|--|--|
| <b>session protocols.</b>                | The available set of SNA communication requests and responses.   |
| <b>share lock.</b>                       | A lock that prevents concurrently executing application processes from changing data, but not from reading data. Contrast with <i>exclusive lock</i> .   |
| <b>short string.</b>                     | A string whose actual length, or a varying-length string whose maximum length, is 255 bytes (or 127 double-byte characters) or less. Regardless of length, a LOB string is not a short string.   |
| <b>sign-on.</b>                          | A request that is made on behalf of an individual CICS or IMS application process by an attachment facility to enable DB2 to verify that it is authorized to use DB2 resources.  |
| <b>simple page set.</b>                  | A nonpartitioned page set. A simple page set initially consists of a single data set (page set piece). If and when that data set is extended to 2 GB, another data set is created, and so on up to a total of 32 data sets. DB2 considers the data sets to be a single contiguous linear address space containing a maximum of 64 GB. Data is stored in the next available location within this address space without regard to any partitioning scheme. |
| <b>simple table space.</b>               | A table space that is neither partitioned nor segmented.   |
| <b>single-byte character set (SBCS).</b> | A set of characters in which each character is represented by a single byte. Contrast with <i>double-byte character set</i> or <i>multibyte character set</i> .  |
| <b>SMF.</b>                              | System management facility.  |
| <b>SMP/E.</b>                            | System Modification Program/Extended.  |
| <b>SMS.</b>                              | Storage Management Subsystem.  |
| <b>SNA.</b>                              | Systems Network Architecture.  |
| <b>SNA network. T</b>                    | he part of a network that conforms to the formats and protocols of Systems Network Architecture (SNA).   |
| <b>sourced function.</b>                 | A function that is implemented by another built-in or user-defined function that is already known to the database manager. This function can be a scalar function or a column (aggregating) function; it returns a single value from a set of values (for example, MAX or AVG). Contrast with <i>built-in function</i> , <i>external function</i> , and <i>SQL function</i> .  |
| <b>special register.</b>                 | A storage area that DB2 defines for an application process to use for storing information that can be referenced in SQL statements. Examples of special registers are USER and CURRENT DATE.   |
| <b>SPUFI.</b>                            | SQL Processor Using File Input.  |
| <b>SQL.</b>                              | Structured Query Language.   |
| <b>SQL authorization ID (SQL ID).</b>    | The authorization ID that is used for checking dynamic SQL statements in some situations.  |
| <b>SQLCA.</b>                            | SQL communication area.  |
| <b>SQL communication area (SQLCA).</b>   | A structure that is used to provide an application program with information about the execution of its SQL statements.   |
| <b>SQLDA.</b>                            | SQL descriptor area.   |
| <b>SQL descriptor area (SQLDA).</b>      | A structure that describes input variables, output variables, or the columns of a result table.  |

|  |   |
|--|---|
| <b>SQL/DS.</b>                                 | Structured Query Language/Data System. This product is now obsolete and has been replaced by DB2 for VSE & VM.  |
| <b>SQL function.</b>                           | A user-defined function in which the CREATE FUNCTION statement contains the source code. The source code is a single SQL expression that evaluates to a single value. The SQL user-defined function can return only one parameter.  |
| <b>SQL processing conversation.</b>            | Any conversation that requires access of DB2 data, either through an application or by dynamic query requests.  |
| <b>SQL Processor Using File Input (SPUFI).</b> | SQL Processor Using File Input. A facility of the TSO attachment subcomponent that enables the DB2I user to execute SQL statements without embedding them in an application program.  |
| <b>SQL routine.</b>                            | A user-defined function or stored procedure that is based on code that is written in SQL.   |
| <b>SSI.</b>                                    | Subsystem interface (in MVS).   |
| <b>SSM.</b>                                    | Subsystem member.   |
| <b>stand-alone.</b>                            | An attribute of a program that means it is capable of executing separately from DB2, without using DB2 services.  |
| <b>star join.</b>                              | A method of joining a dimension column of a fact table to the key column of the corresponding dimension table. See also <i>join</i> , <i>dimension</i> , and <i>star schema</i> .   |
| <b>star schema.</b>                            | The combination of a fact table (which contains most of the data) and a number of dimension tables. See also <i>star join</i> , <i>dimension</i> , and <i>dimension table</i> .   |
| <b>statement string.</b>                       | For a dynamic SQL statement, the character string form of the statement.  |
| <b>static SQL.</b>                             | SQL statements, embedded within a program, that are prepared during the program preparation process (before the program is executed). After being prepared, the SQL statement does not change (although values of host variables that are specified by the statement might change). |
| <b>storage group.</b>                          | A named set of disks on which DB2 data can be stored.   |
| <b>stored procedure.</b>                       | A user-written application program that can be invoked through the use of the SQL CALL statement.   |
| <b>string.</b>                                 | See <i>character string</i> or <i>graphic string</i> .  |
| <b>Structured Query Language (SQL).</b>        | A standardized language for defining and manipulating data in a relational database.  |
| <b>subcomponent.</b>                           | A group of closely related DB2 modules that work together to provide a general function.  |
| <b>subpage.</b>                                | The unit into which a physical index page can be divided.   |
| <b>subquery.</b>                               | A SELECT statement within the WHERE or HAVING clause of another SQL statement; a nested SQL statement.  |
| <b>subselect.</b>                              | That form of a query that does not include ORDER BY clause, UPDATE clause, or UNION operators.  |
| <b>subsystem.</b>                              | A distinct instance of a relational database management system (RDBMS).   |
| <b>sync point.</b>                             | See <i>commit point</i> .   |
| <b>synonym.</b>                                | In SQL, an alternative name for a table or view. Synonyms can be used only to refer to objects at the subsystem in which the synonym is defined.  |

- Sysplex.** See *Parallel Sysplex*.
- Sysplex query parallelism.** Parallel execution of a single query that is accomplished by using multiple tasks on more than one DB2 subsystem. See also *query CP parallelism*.
- system administrator.** The person at a computer installation who designs, controls, and manages the use of the computer system.
- system agent.** A work request that DB2 creates internally such as prefetch processing, deferred writes, and service tasks.
- system conversation.** The conversation that two DB2 subsystems must establish to process system messages before any distributed processing can begin.
- system diagnostic work area (SDWA).**  
The data that is recorded in a SYS1.LOGREC entry that describes a program or hardware error.
- system-directed connection.** A connection that an RDBMS manages by processing SQL statements with three-part names.
- System Modification Program/Extended (SMP/E).**  
A tool for making software changes in programming systems (such as DB2) and for controlling those changes.
- Systems Network Architecture (SNA).**  
The description of the logical structure, formats, protocols, and operational sequences for transmitting information through and controlling the configuration and operation of networks.
- SYS1.DUMPxx data set.** A data set that contains a system dump.
- SYS1.LOGREC.** A service aid that contains important information about program and hardware errors.

## T

|                                   |  |
|-----------------------------------|--|
| <b>table.</b>                     | A named data object consisting of a specific number of columns and some number of unordered rows. See also <i>base table</i> or <i>temporary table</i> .   |
| <b>table check constraint.</b>    | A user-defined constraint that specifies the values that specific columns of a base table can contain.   |
| <b>table function.</b>            | A function that receives a set of arguments and returns a table to the SQL statement that references the function. A table function can be referenced only in the FROM clause of a subselect.  |
| <b>table space.</b>               | A page set that is used to store the records in one or more tables.  |
| <b>table space set.</b>           | A set of table spaces and partitions that should be recovered together for one of these reasons: <ul style="list-style-type: none"> <li>v Each of them contains a table that is a parent or descendent of a table in one of the others.</li> <li>v The set contains a base table and associated auxiliary tables. A table space set can contain both types of relationships.</li> </ul>  |
| <b>task control block (TCB).</b>  | A control block that is used to communicate information about tasks within an address space that are connected to DB2. An address space can support many task connections (as many as one per task), but only one address space connection. See also <i>address space connection</i> .   |
| <b>TB.</b>                        | Terabyte (1 099 511 627 776 bytes).  |
| <b>TCB.</b>                       | Task control block (in MVS).   |
| <b>temporary table.</b>           | A table that holds temporary data; for example, temporary tables are useful for holding or sorting intermediate results from queries that contain a large number of rows. The two kinds of temporary table, which are created by different SQL statements, are the created temporary table and the declared temporary table. Contrast with <i>result table</i> . See also <i>created temporary table</i> and <i>declared temporary table</i> . |
| <b>thread.</b>                    | The DB2 structure that describes an application's connection, traces its progress, processes resource functions, and delimits its accessibility to DB2 resources and services. Most DB2 functions execute under a thread structure. See also <i>allied thread</i> and <i>database access thread</i> .  |
| <b>three-part name.</b>           | The full name of a table, view, or alias. It consists of a location name, authorization ID, and an object name, separated by a period.   |
| <b>time.</b>                      | A three-part value that designates a time of day in hours, minutes, and seconds.   |
| <b>time duration.</b>             | A decimal integer that represents a number of hours, minutes, and seconds.   |
| <b>timeout.</b>                   | Abnormal termination of either the DB2 subsystem or of an application because of the unavailability of resources. Installation specifications are set to determine both the amount of time DB2 is to wait for IRLM services after starting, and the amount of time IRLM is to wait if a resource that an application requests is unavailable. If either of these time specifications is exceeded, a timeout is declared.                       |
| <b>Time-Sharing Option (TSO).</b> | An option in MVS that provides interactive time sharing from remote terminals.   |
| <b>timestamp.</b>                 | A seven-part value that consists of a date and time. The timestamp is expressed in years, months, days, hours, minutes, seconds, and microseconds.   |

|                                 |  |
|---------------------------------|--|
| <b>TMP.</b>                     | Terminal Monitor Program.  |
| <b>to-do.</b>                   | A state of a unit of recovery that indicates that the unit of recovery's changes to recoverable DB2 resources are indoubt and must either be applied to the DASD media or backed out, as determined by the commit coordinator. |
| <b>trace.</b>                   | A DB2 facility that provides the ability to monitor and collect DB2 monitoring, auditing, performance, accounting, statistics, and serviceability (global) data.   |
| <b>TSO.</b>                     | Time-Sharing Option.   |
| <b>TSO attachment facility.</b> | A DB2 facility consisting of the DSN command processor and DB2I. Applications that are not written for the CICS or IMS environments can run under the TSO attachment facility.   |
| <b>type 1 indexes.</b>          | Indexes that were created by a release of DB2 before DB2 Version 4 or that are specified as type 1 indexes in Version 4. Contrast with <i>type 2 indexes</i> . As of Version 7, type 1 indexes are no longer supported.        |
| <b>type 2 indexes.</b>          | Indexes that are created on a release of DB2 after Version 6 or that are specified as type 2 indexes in Version 4 or later.  |

## U

- UDF.** User-defined function.
- UDT.** User-defined data type. In DB2 for OS/390 and z/OS, the term *distinct type* is used instead of user-defined data type. See *distinct type*.
- uncommitted read (UR).** The isolation level that allows an application to read uncommitted data.
- undo.** A state of a unit of recovery that indicates that the changes the unit of recovery made to recoverable DB2 resources must be backed out.
- Unicode.** A standard that parallels the ISO-10646 standard. Several implementations of the Unicode standard exist, all of which have the ability to represent a large percentage of the characters contained in the many scripts that are used throughout the world.
- union.** An SQL operation that combines the results of two select statements. Unions are often used to merge lists of values that are obtained from several tables.
- unique constraint.** An SQL rule that no two values in a primary key, or in the key of a unique index, can be the same.
- unique index.** An index which ensures that no identical key values are stored in a table.
- unlock.** The act of releasing an object or system resource that was previously locked and returning it to general availability within DB2.
- UR.** Uncommitted read.
- URE.** Unit of recovery element.
- URID (unit of recovery ID).** The LOGRBA of the first log record for a unit of recovery. The URID also appears in all subsequent log records for that unit of recovery.
- user-defined data type (UDT).** See *distinct type*.
- user-defined function (UDF).** A function that is defined to DB2 by using the CREATE FUNCTION statement and that can be referenced thereafter in SQL statements. A user-defined function can be an *external function*, a *sourced function*, or an *SQL function*. Contrast with *built-in function*. **UT.** Utility-only access.

## V

|   |   |
|---|---|
| <b>value.</b>   | The smallest unit of data that is manipulated in SQL.   |
| <b>varying-length string.</b>                           | A character or graphic string whose length varies within set limits. Contrast with <i>fixed-length string</i> .   |
| <b>version.</b>   | A member of a set of similar programs, DBRMs, packages, or LOBs.  |
| <b>A version of a program</b>                           | is the source code that is produced by precompiling the program. The program version is identified by the program name and a timestamp (consistency token).   |
| <b>A version of a DBRM</b>                              | is the DBRM that is produced by precompiling a program. The DBRM version is identified by the same program name and timestamp as a corresponding program version.   |
| <b>A version of a package</b>                           | is the result of binding a DBRM within a particular database system. The package version is identified by the same program name and consistency token as the DBRM.  |
| <b>A version of a LOB</b>                               | is a copy of a LOB value at a point in time. The version number for a LOB is stored in the auxiliary index entry for the LOB.   |
| <b>view.</b>  | An alternative representation of data from one or more tables. A view can include all or some of the columns that are contained in tables on which it is defined.   |
| <b>Virtual Storage Access Method (VSAM).</b>            | An access method for direct or sequential processing of fixed- and varying-length records on direct access devices. The records in a VSAM data set or file can be organized in logical sequence by a key field (key sequence), in the physical sequence in which they are written on the data set or file (entry-sequence), or by relative-record number. |
| <b>Virtual Telecommunications Access Method (VTAM).</b> | An IBM licensed program that controls communication and the flow of data in an SNA network.   |
| <b>VSAM.</b>  | Virtual storage access method.  |
| <b>VTAM.</b>  | Virtual Telecommunication Access Method (in MVS).   |

## W

- WLM application environment.** An MVS Workload Manager attribute that is associated with one or more stored procedures. The WLM application environment determines the address space in which a given DB2 stored procedure runs.
- write to operator (WTO).** An optional user-coded service that allows a message to be written to the system console operator informing the operator of errors and unusual system conditions that may need to be corrected.
- WTO.** Write to operator.
- WTOR.** Write to operator (WTO) with reply.

## X

- XMS** Cross memory Service
- XRF.** Extended recovery facility.

## Z

- z/OS.** An operating system for the eServer product line that supports 64-bit real storage.