DB2 Universal Database

Glossary

DB2 Universal Database

Glossary

Before using this information and the product it supports, be sure to read the general information under “Appendix B. Notices” on page 75.

Order publications through your IBM representative or the IBM branch office serving your locality or by calling 1-800-879-2755 in the United States or 1-800-IBM-4YOU in Canada.

Contents

Glossary

A

abend. See *abnormal end of task*.

abend reason code. A 4-byte hexadecimal code that uniquely identifies a problem with DB2 UDB for OS/390.

abnormal end of task (abend). In DB2 UDB for OS/390, the termination of a task, job, or subsystem because of an error condition that recovery facilities cannot resolve during execution.

abnormal termination. (1) A system failure or operator action that causes a job to end unsuccessfully. (2) In DB2, exits that are not under program control, such as a `trap` or `segv`.

absolute path. The full path name of an object. Absolute path names begin at the highest level, or ″root″ directory (which is identified by the forward slash (/) or back slash (\) character).

access function. A user-provided function that converts the data type of text stored in a column to a type that can be processed by the Text Extender.

access method services. A facility that is used to define and reproduce VSAM key-sequenced data sets.

access path. (1) The method that is selected by the optimizer for retrieving data from a specific table. For example, an access path can involve the use of an index, a sequential scan, or a combination of the two. (2) The path that is used to locate data that is specified in SQL statements. An access path can be indexed or sequential.

access plan. The set of access paths that are selected by the optimizer to evaluate a particular SQL statement. The access plan specifies the order of operations to resolve the execution plan, the implementation methods (such as JOIN), and the access path for each table referenced in the statement.

accounting string. User-defined accounting information that is sent to DRDA® servers by DB2 Connect. This information can be specified at one of these locations:

- The client workstation using the SQLESACT API or the DB2ACCOUNT environment variable
- The DB2 Connect workstation using the DFT_ACCOUNT_STR database manager configuration parameter.

active log. (1) In DB2 UDB, the primary and secondary log files that are currently needed for recovery and rollback. Contrast with *archive log*. (2) The portion of the DB2 UDB for OS/390 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 records that are older and no longer fit on the active log.

adjacent nodes. Two nodes connected by at least one path that connects no other nodes.

administrative authority. A level of authority that gives a user privileges over a set of objects. For example, DBADM authority gives privileges over all objects in a database, and SYSADM authority gives privileges over all objects in a system.

administrative support table. A table that is used by a DB2 extender to process user requests on image, audio, and video objects. Some administrative support tables identify user tables and columns that are enabled for an extender. Other administrative support tables contain attribute information about objects in enabled columns. Also called a *metadata table*.

ADSM. See *Tivoli Storage Manager*.

Advanced Peer-to-Peer Networking (APPN). An extension to SNA that features distributed network control, dynamic definition of network resources, and automated resource registration and directory lookup.

Advanced Peer-to-Peer Networking (APPN) network. A collection of interconnected network nodes and their client end nodes.

Glossary

Advanced program-to-program communication (APPC). The general facility that characterizes the LU 6.2 architecture and its various implementations in products.

after-image. In DB2 replication, the updated content of a source table element that is recorded in a change data table or in a database log or journal. Contrast with *before-image*.

agent. (1) A separate process or thread that carries out all DB2 requests that are made by a particular client application. (2) In DB2 UDB for OS/390, the structure that associates all processes that are involved in a DB2 UDB for OS/390 unit of work. An *allied agent* is generally synonymous with an *allied thread*. *System agents* are units of work that process independently of the allied agent, such as prefetch processing, deferred writes, and service tasks.

agent site. In the Data Warehouse Center, the location, defined by a single network host name, where an agent application is installed.

aggregate function. Synonym for *column function*.

alert. An action, such as a beep or warning, that is generated when a performance variable exceeds or falls below its warning or alarm threshold.

alias. An alternative name used to identify a table, view, database, or nickname. An alias can be used in SQL statements to refer to a table or view in the same DB2 subsystem or a remote DB2 subsystem.

alias chain. A series of table aliases that refer to each other in a sequential, nonrepeating fashion.

allied address space. In DB2 UDB for OS/390, an area of storage that is external to and connected to DB2 UDB for OS/390. An allied address space is capable of requesting DB2 UDB for OS/390 services.

allied thread. A thread that originates at the local DB2 UDB for OS/390 subsystem and that can access data at a remote DB2 UDB for OS/390 subsystem.

allocated cursor. In DB2 UDB for OS/390, a cursor that is defined for stored procedure result sets by using the SQL statement ALLOCATE CURSOR.

already verified. An LU 6.2 security option that allows DB2 UDB for OS/390 to provide the user's verified authorization ID when allocating a conversation. The user is not validated by the partner subsystem.

ambiguous cursor. (1) A cursor that cannot be determined to be updatable or read-only from its definition or context. (2) In DB2 UDB for OS/390, 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.

APF. See *authorized program facility*.

API. See *application programming interface*.

APPC. See *advanced program-to-program communication*.

APPL. A VTAM® network definition statement that is used to define DB2 UDB for OS/390 to VTAM as an application program that uses SNA LU 6.2 protocols.

application. A program or set of programs that performs a task; for example, a payroll application.

application ID. A string that uniquely identifies an application across networks. An ID is generated at the time that the application connects to the database. This ID is known on both the client and the server and can be used to correlate the two parts of the application.

application plan. The control structure that is produced during the bind process. DB2 UDB for OS/390 uses the application plan to process SQL statements that it encounters during statement execution.

application process. The unit to which resources and locks are allocated. An application process involves the running of one or more programs.

application programming interface (API). (1) A functional interface supplied by the operating system or by a separately orderable licensed program. An API allows an application program that is written in a high-level language to use specific data or functions of the operating system or the licensed programs. (2) In DB2, a function within the interface, for example, the get error message API.

application requester. A facility that accepts a database request from an application process and passes it to an application server.

application server. The local or remote database manager to which the application process is connected.

Apply program. In DB2 replication, a program that is used to refresh or update a target table, depending on the applicable source-to-target rules. Contrast with *Capture program* and *Capture trigger*.

Apply qualifier. In DB2 replication, a character string that identifies subscription definitions that are unique to each instance of the Apply program.

APPN. See *Advanced Peer-to-Peer Networking*

archive log. (1) The set of log files that are closed and are no longer needed for normal processing. These files are retained for use in roll-forward recovery. Contrast with *active log*. (2) The portion of the DB2 UDB for OS/390 log that contains log records that are copied from the active log.

argument. A value passed to or returned from a function or procedure at run time.

asynchronous. Without regular time relationship; unexpected and unpredictable with respect to the processing of program instructions. Contrast with *synchronous*.

asynchronous batched update. A process in which all changes to the source are recorded and applied to existing target data at specified intervals. Contrast with *asynchronous continuous update*.

asynchronous continuous update. A process in which all changes to the source are recorded and applied to existing target data after being committed in the base table. Contrast with *asynchronous batched update*.

attach. In DB2, to remotely access objects at the instance level.

attachment facility. An interface between DB2 UDB for OS/390 and TSO, IMS™, CICS, or batch address spaces. An attachment facility allows application programs to access DB2 UDB for OS/390.

attribute. In SQL database design, a characteristic of an entity. For example, the phone number of an employee is one of that employee's attributes.

authority. See *administrative authority.*

authorization ID. (1) A character string in a statement that designates a set of privileges. It is used by the database manager for authorization checking and as an implicit qualifier for the names of objects such as tables, views, and indexes. (2) A string that can be verified for connection to DB2 UDB for OS/390 and to which a set of privileges is allowed. An authorization ID can represent an individual, an organizational group, or a function, but DB2 UDB for OS/390 does not determine this representation.

authorized program facility (APF). In DB2 UDB for OS/390, a facility that permits the identification of programs that are authorized to use restricted functions.

autocommit. To automatically commit the current unit of work after each SQL statement.

automatic rebind. (1) A feature that automatically rebinds an invalidated package without requiring a **bind** command to be entered manually or a bind file to be present. (2) In DB2 UDB for OS/390, 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. See also *bind*.

auxiliary index. In DB2 UDB for OS/390, an index on an auxiliary table in which each index entry refers to an LOB.

auxiliary table. In DB2 UDB for OS/390, a table that stores columns outside the table in which they are defined. Contrast with *base table*.

B

backup pending. The state of a database or table space that prevents an operation from being performed until the database or table space is backed up.

backward log recovery. The fourth and final phase of restart processing during which DB2 UDB for OS/390 scans the log in a backward direction to apply UNDO log records for all aborted changes.

base aggregate table. In DB2 replication, a type of target table that contains data aggregated from a source table or a point-in-time table at intervals.

base table. (1) A table created with the CREATE TABLE statement. Such a table has both its description and data physically stored in the database. Contrast with *view*. (2) In DB2 UDB for OS/390: (a) A table that is created by the SQL CREATE TABLE statement and that holds persistent data. Contrast with *result table* and *temporary table*. (b) A table that contains an LOB column definition. The actual LOB column data is not stored with the base table. The base table contains a row ID for each row and an indicator column for each of its LOB columns. Contrast with *auxiliary table*.

base table space. In DB2 UDB for OS/390, a table space that contains base tables.

basic conversation. An LU 6.2 conversation between two transaction programs using the APPC basic conversation API. Contrast with *mapped conversation*.

basic predicate. A predicate that compares two values.

basic sequential access method (BSAM). An access method that DB2 UDB for OS/390 uses for storing or retrieving data blocks in a continuous sequence, using either a sequential access or a direct access device.

before-image. In DB2 replication, the content of a source table column prior to a refresh, as recorded in a change data table or in a database log or journal. Contrast with *after-image*.

before trigger. In DB2 UDB for OS/390, a trigger that is defined with the trigger activation time BEFORE.

binary integer. A basic data type that can be further classified as small integer or large integer.

binary large object (BLOB). A sequence of bytes with a size ranging from 0 bytes to 2 gigabytes. This string does not have an associated code page and character set. Image, audio, and video objects are stored in BLOBs. Compare to *character large object (CLOB).*

binary string. In DB2 UDB for OS/390, a sequence of bytes that is not associated with a CCSID. For example, the BLOB data type is a binary string.

bind. (1) In SQL, the process by which the output from the SQL precompiler is converted to a usable structure called an *access plan*. During this process, access paths to the data are selected and some authorization checking is performed. (2) In DB2 UDB for OS/390, the process by which the output from the DBMS precompiler is converted to a usable control structure (which is called a *package* or an *application plan*). During the process, access paths to the data are selected and some authorization checking is performed. See also *automatic rebind*, *dynamic bind*, *incremental bind*, *static bind*.

bindery object name. A 48-byte character string that contains the name of a bindery object on the NetWare file server. The database manager configuration field, objectname, uniquely represents a DB2 server instance, and is stored as an object in the bindery on a NetWare file server.

bind file. A file produced by the precompiler when the **bind** command or API is used with the BINDFILE option. This file includes information about all SQL statements in the application program.

bit data. Data with character type CHAR or VARCHAR that is not associated with a coded character set and therefore is never converted.

BLOB. See *binary large object*.

block. A string of data elements recorded or transmitted as a unit.

blocking. An option that is specified when binding an application. It allows caching of multiple rows of information by the communications subsystem so that each FETCH statement does not require the transmission of one row for each request across the network. Contrast with *data blocking*.

bootstrap data set (BSDS). A VSAM data set that contains name and status information for DB2 UDB for OS/390, as well as RBA range specifications, for all active and archive log data sets. It also contains passwords for the DB2 UDB for OS/390 directory and catalog, and lists of conditional restart and checkpoint records.

broadcast join. A join in which all partitions of a table are sent to all nodes.

browser. A Text Extender function that enables you to display text on a computer monitor.

BSAM. See *basic sequential access method*.

BSDS. See *bootstrap data set*.

buffer pool. In DB2 UDB for OS/390, main storage that is reserved to satisfy the buffering requirements for one or more table spaces or indexes.

built-in function. An SQL function that is provided by DB2 and appears in the SYSIBM schema. Contrast with *user-defined function*.

business metadata. Data that describes information assets in business terms. Business metadata is stored in the information catalog and accessed by users to find and understand the information they need. For example, business metadata for a program would contain a description of what the program does and what tables it uses. Contrast with *technical metadata*.

business name. In the Data Warehouse Center, a name that refers to a step. Each step has a business name and a DB2 table name that is associated with the step. Business names are generally used by warehouse users; DB2 table names are used in SQL statements.

byte reversal. A technique in which numeric data is stored with the least significant byte first.

C

cache. A buffer that contains frequently accessed instructions and data; it is used to reduce access time.

Cache Manager. In Net.Data®, the program that manages a cache for one workstation. The Cache Manager can manage multiple caches.

cache structure. A coupling facility structure that stores data that can be available to all members of a Parallel Sysplex®. A DB2 UDB for OS/390 data sharing group uses cache structures as group buffer pools.

caching. The process of storing frequently used results from a request to the Web server locally for quick retrieval, until it is time to refresh the information.

CAF. See *call attachment facility*.

call attachment facility (CAF). A DB2 UDB for OS/390 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.

call level interface (CLI). A callable API for database access, which is an alternative to an embedded SQL API. In contrast to embedded SQL, the CLI does not require precompiling or binding by the user, but instead provides a standard set of functions to process SQL statements and related services at run time.

Capture program. In DB2 replication, a program that reads database log or journal records to capture data about changes made to DB2 source tables. Contrast with *Apply program* and *Capture trigger*.

Capture trigger. In DB2 replication, a mechanism that captures delete, update, and insert operations performed on non-IBM source tables. Contrast with *Capture program* and *Apply program.*

cardinality. The number of rows in a database table.

Glossary

cascade. In the Data Warehouse Center, to run a sequence of events. When a step cascades to another step, the steps run sequentially or concurrently. A step can also cascade to a program, which runs after the step finishes running.

cascade delete. The way in which DB2 UDB for OS/390 enforces referential constraints when it deletes all descendent rows of a deleted parent row.

cascade rejection. In DB2 replication, the process of rejecting a replication transaction because it is associated with a transaction that had a conflict detected and was itself rejected.

CASE expression. In DB2 UDB for OS/390, an expression that allows another expression to be selected based on the evaluation of one or more conditions.

cast function. A function used to convert instances of a data type (origin) into instances of a different data type (target). In general, cast functions have the name of the target data type. They have a single argument whose type is the origin data type; their return type is the target data type.

catalog. A set of tables and views maintained by the database manager. These tables and views contain information about the database, such as descriptions of tables, views, and indexes.

catalog node. The node at which the catalog tables reside. The catalog node can be a different node for each database.

catalog table. Any table in the DB2 UDB for OS/390 catalog.

catalog view. A view of a system table created by the Text Extender for administration purposes. A catalog view contains information about the tables and columns that are enabled for use by the Text Extender.

CCD table. See *consistent-change-data table*.

CCSID. See *coded character set identifier*.

CDB. See *communications database*.

CDRA. See *Character Data Representation Architecture*.

CD table. See *change data table*.

CEC. Central electronic complex. 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.

CFRM policy. In DB2 UDB for OS/390, a declaration by an MVS administrator regarding the allocation rules for a coupling facility structure.

change aggregate table. In DB2 replication, a type of target table that contains data aggregations based on changes recorded for a source table.

change data (CD) table. A replication control table at the source server that contains changed data for a replication source table.

Character Data Representation Architecture (CDRA). An architecture used to achieve consistent representation, processing, and interchange of string data.

character large object (CLOB). A sequence of characters (single-byte, multibyte, or both) up to 2 gigabytes. A CLOB can be used to store large text objects. Also called character large object string. Compare to *binary large object (BLOB)*.

character string. A sequence of bytes or characters.

character string delimiter. The characters used to enclose character strings in delimited ASCII files that are imported or exported. See *delimiter*.

CHECK clause. In SQL, an extension to the SQL CREATE TABLE and SQL ALTER TABLE statements that specifies a table check constraint.

check condition. A restricted form of search condition used in check constraints.

check constraint. A constraint that specifies a check condition that is not false for each row of the table on which the constraint is defined.

check integrity. In DB2 UDB for OS/390, 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 UDB for OS/390 to enforce table check constraints on operations that add or change data.

check pending. A state into which a table can be put where only limited activity is allowed on the table and constraints are not checked when the table is updated.

checkpoint. A point at which DB2 UDB for OS/390 records internal status information on the log; the recovery process uses this information if the subsystem abnormally terminates.

CI. See *control interval*.

CICS. An IBM® licensed program that provides online transaction-processing services and management for critical business appliations. In DB2 UDB for OS/390 information, this term represents the following products:
CICS Transaction Server for OS/390®: Customer Information Control Center Transaction Server for OS/390
CICS/ESA: Customer Information Control System/Enterprise Systems Architecture
CICS/MVS: Customer Information Control System/Multiple Virtual Storage

CICS attachment facility. A DB2 UDB for OS/390 subcomponent that uses the MVS subsystem interface (SSI) and cross storage linkage to process requests from CICS to DB2 UDB for OS/390 and to coordinate resource commitment.

CIDF. See *control interval definition field*.

circular log. A database log in which records are overwritten if they are no longer needed by an active database. Consequently, if a failure occurs, lost data cannot be restored during forward recovery. Contrast with *recoverable log*.

claim. In DB2 UDB for OS/390, a notification to the DBMS that an object is being accessed. Claims prevent drains from occurring until the claim is released, which usually occurs at a commit point. See also *drain*.

claim class. In DB2 UDB for OS/390, a specific type of object access that can be one of the following types: cursor stability (CS), repeatable read (RR), write.

claim count. In DB2 UDB for OS/390, a count of the number of agents that are accessing an object.

class of service. In DB2 UDB for OS/390, a VTAM term for a list of routes through a network, arranged in an order of preference for their use.

clause. In DB2 UDB for OS/390 SQL, a distinct part of a statement, such as a SELECT clause or a WHERE clause.

cleanse. The process of manipulating the data extracted from operational systems so as to make it usable by the data warehouse.

CLI. See *call level interface*.

client. (1) Any program (or workstation that it is running on) that communicates with and accesses a database server. (2) See *requester*.

cliette. A long-running process in Net.Data Live Connection that serves requests from the Web server. The Connection Manager schedules cliette processes to serve these requests.

CLIST. Command list. A language that DB2 UDB for OS/390 uses to perform TSO tasks.

CLOB. See *character large object*.

CLP. See *Command Line Processor*.

CLPA. See *create link pack area*.

clustered index. An index whose sequence of key values closely corresponds to the sequence of rows stored in a table. The degree to which this correspondence exists is measured by statistics that are used by the optimizer.

Glossary

coded character set. A set of unambiguous rules that establishes a character set and the one-to-one relationships between the characters of the set and their coded representations.

coded character set identifier (CCSID). A number that includes an encoding scheme identifier, character set identifiers, code page identifiers, and other information that uniquely identifies the coded graphic character representation.

code page. A set of assignments of characters to code points.

code point. In CDRA, a unique bit pattern that represents a character in a code page.

code set. Encoding values for a character set that provides the interface between the system and its input and output devices. ISO uses code set as the term equivalent to the IBM-defined term code page.

cold start. (1) The process of starting a system or program using an initial program load procedure. Contrast with *warm start*. (2) A process by which DB2 UDB for OS/390 restarts without processing any log records.

collating sequence. The sequence in which the characters are ordered for the purpose of sorting, merging, comparing, and processing indexed data sequentially.

collection. In DB2 UDB for OS/390, a group of packages that have the same qualifier.

collocated join. The result of two tables being joined in which the following conditions are met:

- The tables reside in a single-partition nodegroup in the same database partition; or they are in the same partitioned nodegroup and have the same number of partitioning columns, the columns are partition-compatible, and both tables use the same partitioning function.
- All pairs of the corresponding partitioning key columns participate in the equijoin predicates.

column distribution value. Statistics describing the most frequent values of some column or the quantile values. These values are used in the optimizer to help determine the best access plan.

column function. (1) An operation used in queries that applies to the values from several rows. Column functions include SUM, AVG, MIN, MAX, COUNT, STDDEV, and VARIANCE. Synonym for *aggregate function*. (2) In DB2 UDB for OS/390, an SQL operation that derives its result from a collection of values across one or more rows. Contrast with *scalar function*.

"come from" checking. An LU 6.2 security option that defines a list of authorization IDs that are allowed to connect to DB2 UDB for OS/390 from a partner LU.

command. A DB2 UDB for OS/390 operator command or a DSN subcommand. A command is distinct from an SQL statement.

Command Line Processor (CLP). A character-based interface for entering SQL statements and database manager commands.

command prefix. In DB2 UDB for OS/390, a one- to eight-character command identifier. The command prefix distinguishes the command as belonging to an application or subsystem rather than to OS/390.

command recognition character (CRC). A character that permits an MVS console operator or an IMS subsystem user to route DB2 commands to specific DB2 UDB for OS/390 subsystems.

command scope. In DB2 UDB for OS/390, the scope of command operation in a data sharing group. If a command has *member scope*, the command displays information only from the one member or affects only non-shared resources that are owned locally by that member. If a command has *group scope*, the command displays information from all members, affects non-shared resources that are owned locally by all members, displays information on sharable resources, or affects sharable resources.

commit. The operation that ends a unit of work by releasing locks so that the database changes made by that unit of work can be perceived by other processes. This operation makes the data changes permanent.

commitment control. The establishment of a boundary within the process under which Net.Data is running, where operations on resources are part of a unit of work.

commit point. A point in time when data is considered to be consistent. Synonym for *point of consistency.*

committed phase. In DB2 UDB for OS/390, the second phase of the multi-site update process that requests all participants to commit the effects of the logical unit of work.

common-index table. A DB2 table whose text columns share a common text index. See also *multi-index table*.

Common Programming Interface Communications (CPI-C). An API for applications that require program-to-program communication, using SNA LU 6.2 to create a set of interprogram services.

common service area (CSA). In OS/390, a part of the common area that contains data areas that can be addressed by all address spaces.

common table expression. An expression that defines a result table with a name (qualified SQL identifier) that can be specified as a table name in any FROM clause in the fullselect that follows the WITH clause.

communications database (CDB). A set of tables in the DB2 UDB for OS/390 catalog that are used to establish conversations with remote database management systems.

comparison operator. An infix operator used in comparison expressions. Comparison operators are ¬< (not less than), <= (less than or equal to), ¬= (not equal to), = (equal to), >= (greater than or equal to), > (greater than), and ¬> (not greater than).

complete. A table attribute that indicates that the table contains a row for every primary key value of interest. As a result, a complete source table can be used to perform a refresh of a target table.

complete CCD table. A CCD table that contains all the rows that satisfy the source view and predicates from the source table or view. Contrast with *noncomplete CCD table*.

composite key. An ordered set of key columns of the same table.

compound SQL statement. A block of SQL statements that are executed in a single call to the application server.

compression dictionary. In DB2 UDB for OS/390, 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.

concurrency. The shared use of resources by multiple interactive users or application processes at the same time.

condensed. A table attribute indicating that the table contains current data rather than a history of changes to the data. A condensed table includes no more than one row for each primary key value in the table. As a result, a condensed table can be used to supply current information for a refresh.

condensed CCD table. In DB2 replication, a CCD table that contains only the most current value for a row. This type of table is useful for staging changes to remote locations and for summarizing hot-spot updates. Contrast with *noncondensed CCD table*.

conditional restart. A DB2 UDB for OS/390 restart that is directed by a user-defined conditional restart control record (CRCR).

conflict detection. In update-anywhere replication configurations:

- The process of detecting constraint errors.
- The process of detecting if the same row was updated in the source and target tables during the same replication cycle. When a conflict is detected, the transaction that caused the conflict is rejected. See also *enhanced conflict detection*, *standard conflict detection*, and *row-replica conflict detection*.

connect. In DB2, to access objects at the database level.

connection. (1) An association between an application process and an application server. (2) In data communications, an association established between functional units for conveying information. (3) In SNA, the existence of a communication path between two partner LUs that allows information to be exchanged (for example, two DB2 UDB for OS/390 subsystems that are connected and communicating by way of a conversation).

connection handle. Within the CLI, the data object that contains information associated with a connection. This information includes general status information, transaction status, and diagnostic information.

Glossary

connection ID. In DB2 UDB for OS/390, an identifier that is supplied by the attachment facility and that is associated with a specific address space connection.

Connection Manager. An executable file, `dtwcm`, in Net.Data that is needed to support Live Connection.

consistency token. In DB2 UDB for OS/390, a timestamp that is used to generate the version identifier for an application.

consistent-change-data (CCD) table. In DB2 replication, a type of target table that is used for auditing or staging data or both. See also *complete CCD table*, *condensed CCD table*, *external CCD table*, *internal CCD table*, *noncomplete CCD table*, and *noncondensed CCD table*.

constant. A language element that specifies an unchanging value. Constants are classified as string constants or numeric constants. Contrast with *variable*.

constraint. A rule that limits the values that can be inserted, deleted, or updated in a table. See *check constraint*, *referential constraint*, and *unique constraint*.

container. See *table space container*.

contention. In the database manager, a situation in which a transaction attempts to lock a row or table that is already locked.

Control Center. A graphical interface that shows database objects (such as databases and tables) and their relationship to each other. From the Control Center, you can perform the tasks provided by the DBA Utility, Visual Explain, and Performance Monitor tools. Contrast with *DataJoiner Replication Administration (DJRA) tool*.

control interval (CI). In VSAM, a fixed-length area of 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.

control interval definition field (CIDF). 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.

control metadata. In the Data Warehouse Center, information about changes to the warehouse, such as the date and time that a table is updated by the processing of a step.

control point. (1) In APPN, a component of a node that manages resources of that node and optionally provides services to other nodes in the network. Examples are a system services control point (SSCP) in a type 5 node, a physical unit control point (PUCP) in a type 4 node, a network node control point (NNCP) in a type 2.1 (T2.1) network node, and an end node control point (ENCP) in a T2.1 end node. An SSCP and an NNCP can provide services to other nodes. (2) A component of a T2.1 node that manages the resources of that node. If the T2.1 node is an APPN node, the control point is capable of engaging in control point-to-control point sessions with other APPN nodes. If the T2.1 node is a network node, the control point also provides services to adjacent end nodes in the T2.1 network. See also *physical unit*.

control privilege. The authority to completely control an object. This includes the authority to access, drop, or alter an object, and the authority to extend or revoke privileges on the object to other users.

control server. In DB2 replication, the database location of the applicable subscription definitions and Apply program control tables.

control table. In DB2 replication, a table in which replication source and subscription definitions or other replication control information is stored.

conversation. In APPC, a connection between two transaction programs over a logical unit-logical unit (LU-to-LU) session that allows them to communicate with each other while processing a transaction.

conversational transaction. In APPC, two or more programs communicating using the services of logical units (LUs).

conversation security. In APPC, a process that allows validation of a user ID or group ID and password before establishing a connection.

conversation security profile. The set of user IDs or group IDs and passwords that are used by APPC for conversation security.

Coordinated Universal Time (UTC). Synonym for Greenwich Mean Time.

coordinating agent. The agent that is started when a request is received by the database manager from an application. It remains associated with the application during the life of the application. This agent coordinates subagents that work for the application. See also *subagent*.

coordinator. In DB2 UDB for OS/390, 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.

coordinator node. The node to which the application originally connected and on which the coordinating agent resides.

coordinator subsection. The subsection of an application that starts other subsections (if any) and returns results to the application.

correlated columns. In SQL, a relationship between the value of one column and the value of another column.

correlated reference. A reference to a column of a table that is outside a subquery.

correlated subquery. A subquery that contains a correlated reference to a column of a table that is outside the subquery.

correlation ID. In DB2 UDB for OS/390, an identifier that is associated with a specific thread. In TSO, it is either an authorization ID or the job name.

correlation name. An identifier designating 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.

cost category. A category into which DB2 UDB for OS/390 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:

- A: Indicates that DB2 UDB for OS/390 had enough information to make a cost estimate without using default values.
- B: Indicates that some condition exists for which DB2 UDB for OS/390 was forced to use default values for its estimate.

The cost category is externalized in the COST_CATEGORY column of DSN_STATEMNT_TABLE when a statement is explained.

country code. When accessing the database, the country code of the application is used to determine the date and time presentation (display and print) formats. It is also used with the code page to determine the default collating sequence for the database.

coupling facility. In an OS/390 environment, a special PR/SM™ LPAR logical partition that runs the coupling facility control program and provides high-speed caching, list processing, and locking functions in a Parallel Sysplex.

CP. See *control point*.

CPC. See *central processor complex*.

CPI-C. See *Common Programming Interface Communications*.

CPI-C side information profile. In SNA, the profile that specifies the conversation characteristics to use when allocating a conversation with a remote transaction program. The profile is used by local transaction programs that communicate through CPI Communications. It specifies the partner LU name (the name of the connection profile that contains the remote LU name), the mode name, and the remote transaction program name.

CP name. Control point name. A network-qualified name of a control point that consists of a network ID qualifier that identifies the network to which the control point node belongs.

crash recovery. The process of recovering from an immediate failure.

CRC. See *command recognition character*.

CRCR. In DB2 UDB for OS/390, conditional restart control record. See *conditional restart*.

create link pack area (CLPA). An option used during IPL to initialize the link pack pageable area.

cross-memory linkage. In an OS/390 environment, a method for invoking a program in a different address space. The invocation is synchronous with respect to the caller.

cross-system coupling facility (XCF). A component of OS/390 that provides functions to support cooperation between authorized programs running within a Parallel Sysplex.

cross-system extended services (XES). A set of OS/390 services that enable multiple instances of an application or subsystem, running on different systems in a Parallel Sysplex environment, to implement high-performance, high-availability data sharing by using a coupling facility.

CS. See *cursor stability*.

CSA. See *common service area.*

CT. See *cursor table*.

current data. In DB2 UDB for OS/390, data within a host structure that is current with (identical to) the data within the base table.

current function path. An ordered list of schema names used in the resolution of unqualified references to functions and data types. In dynamic SQL, the current function path is found in the CURRENT FUNCTION PATH special register. In static SQL, it is defined in the FUNCPATH option for PREP and BIND commands.

current status rebuild. In DB2 UDB for OS/390, the second phase of restart processing during which the status of the subsystem is reconstructed from information on the log.

current working directory. The default directory of a process from which all relative path names are resolved.

cursor. A named control structure used by an application program to point to a specific row within some ordered set of rows. The cursor is used to retrieve rows from a set.

cursor stability (CS). An isolation level that locks any row accessed by a transaction of an application while the cursor is positioned on the row. The lock remains in effect until the next row is fetched or the transaction is terminated. If any data is changed in a row, the lock is held until the change is committed to the database.

cursor table (CT). In DB2 UDB for OS/390, the copy of the skeleton cursor table that is used by an executing application process.

cycle. In DB2 UDB for OS/390, 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

DARI. Database Application Remote Interface. Obsolete term for *stored procedure*.

data area. A memory area used by a program to hold information.

database access thread. In DB2 UDB for OS/390, a thread that accesses data at the local subsystem on behalf of a remote subsystem.

database administrator (DBA). A person who is responsible for the design, development, operation, safeguarding, maintenance, and use of a database.

Database Application Remote Interface (DARI). Obsolete term for *stored procedure*.

database catalog. In the Data Warehouse Center, a collection of tables that contains descriptions of database objects such as tables, views, and indexes.

database client. A workstation used to access a database that is on a database server.

database connection services (DCS) directory. A directory that contains entries for remote databases and the corresponding application requester used to access them.

database descriptor (DBD). An internal representation of a DB2 UDB for OS/390 database definition, which reflects the data definition that is in the DB2 UDB for OS/390 catalog. The objects that are defined in a database descriptor are table spaces, tables, indexes, index spaces, and relationships.

database directory. A directory that contains database access information for all databases to which a client can connect.

database engine. The part of the database manager providing the base functions and configuration files needed to use the database.

database log. A set of primary and secondary log files consisting of log records that record all changes to a database. The database log is used to roll back changes for units of work that are not committed and to recover a database to a consistent state.

database-managed space (DMS) table space. A table space whose space is managed by the database. Contrast with *system-managed space (SMS) table space*.

database management system (DBMS). Synonym for *database manager*.

database manager. A computer program that manages data by providing the services of centralized control, data independence, and complex physical structures for efficient access, integrity, recovery, concurrency control, privacy, and security.

database manager instance. A logical database manager environment similar to an image of the actual database manager environment. You can have several instances of the database manager product on the same workstation. You can use these instances to separate the development environment from the production environment, tune the database manager to a particular environment, and protect sensitive information from a particular group of people.

database node. See *database partition.*

database object. Anything that can be created or manipulated with SQL—for example, tables, views, indexes, packages, triggers, or table spaces.

database partition. A part of the database that consists of its own user data, indexes, configuration files, and transaction logs. Sometimes called a *node* or *database node*.

database request module (DBRM). A data set member that is created by the DB2 UDB for OS/390 precompiler and that contains information about SQL statements. DBRMs are used in the bind process.

database server. A functional unit that provides database services for databases.

database system monitor. A collection of programming APIs that monitor performance and status information about the database manager, databases, and applications using the database manager and DB2 Connect.

data blocking. The process of specifying how many minutes worth of change data will be replicated during a subscription cycle. Contrast with *blocking*.

data currency. In DB2 UDB for OS/390, the state in which data that is retrieved into a host variable in your program is a copy of data in the base table.

data definition language (DDL). A language for describing data and its relationships in a database. Synonym for *data description language*.

data definition name (ddname). In DB2 UDB for OS/390, the name of a data definition (DD) statement that corresponds to a data control block that contains the same name.

data description language. Synonym for *data definition language*.

DataJoiner. A separately available product that provides client applications integrated access to distributed data and provides a single database image of a heterogeneous environment. With DataJoiner, a client application can join data (using a single SQL statement) that is distributed across multiple database management systems or update a single remote data source as if the data were local.

Glossary

DataJoiner Replication Administration (DJRA) tool. A database administration tool that you can use to perform various replication administration tasks. Unlike the Control Center, the DJRA tool can be used to administer replication for non-IBM databases. Contrast with *Control Center*.

DATALINK. A DB2 data type that enables logical references from the database to a file stored outside the database.

data link control (DLC). In SNA, the protocol layer that consists of the link stations that schedule data transfer over a link between two nodes and perform error control for the link.

data manipulation language (DML). A subset of SQL statements used to manipulate data.

datamart. A subset of a data warehouse that contains data tailored for the specific needs of a department or team. A datamart can be a subset of a warehouse for your entire organization, such as data contained in OLAP tools.

data partition. In an OS/390 environment, a VSAM data set that is contained within a partitioned table space.

data sharing. The ability of two or more DB2 UDB for OS/390 subsystems to directly access and change a single set of data.

data sharing group. A collection of one or more DB2 UDB for OS/390 subsystems that directly access and change the same data while maintaining data integrity.

data sharing member. A DB2 UDB for OS/390 subsystem that is assigned by XCF services to a data sharing group.

data space. In DB2 UDB for OS/390, a range of up to 2 gigabytes 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.

data type. In SQL, an attribute of columns, literals, host variables, special registers, and the results of functions and expressions.

Data Warehouse Center. A graphical interface, and the software behind it, that enables you to work with the components of the warehouse. You can use the Data Warehouse Center to define and manage the warehouse data and the processes that create the data in the warehouse.

Data Warehouse Center administrative interface. The user interface to the administration functions of the Data Warehouse Center. The interface can be on the Data Warehouse Center server or on different machines for multiple administrators.

Data Warehouse Center program. A program, supplied with the Data Warehouse Center, that can be started from the Data Warehouse Center and that is automatically defined, for example, DB2 Load programs and transformers.

Data Warehouse Center property. An attibute that applies across sessions of the Data Warehouse Center, such as the warehouse control database that contains the technical metadata. See also *property*.

date. A three-part value that designates a day, month, and year.

date duration. A DECIMAL(8,0) value that represents a number of years, months, and days.

datetime value. A value of the data type DATE, TIME, or TIMESTAMP.

DBA. See *database administrator*.

DBA Utility. A tool that lets DB2 users configure databases and database manager instances, manage the directories necessary for accessing local and remote databases, back up and recover databases or table spaces, and manage media on a system using a graphical interface. The tasks provided by this tool can be accessed from the Control Center.

DBCLOB. See *double-byte character large object*.

DBCS. See *double-byte character set*.

DBD. See *database descriptor*.

DBID. Database identifier.

DBMS. Database management system. See *database manager*.

DBMS instance connection. A logical connection between an application and an agent process or thread owned by a DB2 instance.

DBRM. See *database request module*.

DB2 CLI. DB2 Call Level Interface. An alternative SQL interface for the DB2 family of products that takes full advantage of DB2 capability.

DB2 command. An instruction to the DB2 UDB for OS/390 subsystem allowing a user to start or stop DB2 UDB for OS/390, to display information on current users, to start or stop databases, to display information on the status of databases, and so on.

DB2 Connect. A product that provides the function necessary (DRDA application requester support) for client applications to read and update data stored in DRDA application servers.

DB2 extender. A program that you can use to store and retrieve data types beyond the traditional numeric and character data, such as image, audio, and video data, and complex documents.

DB2I. In DB2 UDB for OS/390, DATABASE 2 Interactive.

DB2I Kanji Feature. In DB2 UDB for OS/390, the tape that contains the panels and jobs that allow a site to display DB2I panels in Kanji.

DB2 PM. In DB2 UDB for OS/390, DATABASE 2 Performance Monitor.

DB2 SDK. See *DB2 Application Development Client*.

DB2 Application Development Client (DB2 SDK). A collection of tools that help developers create database applications.

DB2 thread. The DB2 UDB for OS/390 structure that describes an application's connection, traces its progress, processes resource functions, and delimits its accessibility to DB2 UDB for OS/390 resources and services.

DB2UEXIT. An optional, user-written executable program that the database manager invokes to move or retrieve archive log files.

DCE. See *Distributed Computing Environment*.

DCE ticket. In an OS/390 environment, 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.

DCLGEN. See *declarations generator*.

DDF. See *distributed data facility*.

DDL. See *data definition language*.

ddname. See *data definition name*.

deadlock. A condition under which a transaction cannot proceed because it is dependent on exclusive resources that are locked by some other transaction, which in turn is dependent on exclusive resources in use by the original transaction.

deadlock detector. A process within the database manager that monitors the states of the locks to determine if a deadlock condition exists. When a deadlock condition is detected, the detector stops one of the transactions involved in the deadlock. This transaction is rolled back and the other transactions proceed.

declarations generator (DCLGEN). A subcomponent of DB2 UDB for OS/390 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 UDB for OS/390 system catalog information. DCLGEN is also a DSN subcommand.

deferred embedded SQL. In DB2 UDB for OS/390, SQL statements that are neither fully static nor fully dynamic. Like static statements, they are embedded within an application, but like dynamic statements, they are prepared during the execution of the application.

definition metadata. In the Data Warehouse Center, information about the format of the data warehouse (the schema), the sources of the data, and the transformations applied in loading the data.

degree of parallelism. In DB2 UDB for OS/390, the number of concurrently executed operations that are initiated to process a query.

delete-connected. In SQL, a table that is a dependent of table P or a dependent of a table to which delete operations from table P cascade.

delete rule. A rule associated with a referential constraint that either restricts the deletion of a parent row or specifies the effect of such a deletion on the dependent rows.

delete trigger. In DB2 UDB for OS/390, a trigger that is defined with the triggering SQL operation DELETE.

delimited identifier. A sequence of characters enclosed within double quotation marks. The sequence must consist of a letter followed by zero or more characters, each of which is a letter, digit, or the underscore character.

delimiter. A character or flag that groups or separates items of data.

delimiter token. A string constant, a delimited identifier, an operator symbol, or any of the special characters shown in syntax diagrams.

dependent. In SQL, an object (row, table, or table space) that has at least one parent. See *parent row*, *parent table*, *parent table space*.

dependent logical unit (DLU). A logical unit that requires assistance from a system services control point (SSCP) to instantiate an LU-to-LU session.

dependent row. A row that contains a foreign key that matches the value of a parent key in the parent row. The foreign key value represents a reference from the dependent row to the parent row.

dependent table. A table that is a dependent in at least one referential constraint.

descendent. An object that is a dependent of an object or is the dependent of a descendent of an object.

descendent row. A row that is dependent on another row or a row that is a descendent of a dependent row.

descendent table. A table that is a dependent of another table or a descendent of a dependent table.

deterministic function. See *not-variant function*.

device name. A name reserved by the system, or a device driver that refers to a specific device.

DFHSM. In an OS/390 environment, Data Facility Hierarchical Storage Manager.

DFP. In an OS/390 environment, Data Facility Product.

dictionary. A collection of language-related linguistic information that the Text Extender uses during text analysis, indexing, retrieval, and highlighting of documents in a particular language.

differential refresh. In DB2 replication, a process in which only changed data is copied to the target table, replacing existing data. Contrast with *full refresh*.

dimension. In the OLAP Starter Kit, a data category, such as time, accounts, products, or markets. Dimensions represent the highest consolidation level in a multidimensional database outline.

directed join. A relational operation in which all of the rows in one or both of the joined tables are rehashed and directed to new database partitions based on the join predicate. If all of the partitioning key columns in a table participate in the equijoin predicates, the other table is rehashed; otherwise (if there is at least one equijoin predicate), both tables are rehashed.

directory. The DB2 UDB for OS/390 system database that contains internal objects such as database descriptors and skeleton cursor tables.

directory services. A portion of the APPN protocols that maintains information about the location of resources in an APPN network.

disable. To restore a database, a text table, or a text column to its condition before it was enabled for the Text Extender by removing the items created during the enabling process.

distinct type. 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.

Distributed Computing Environment (DCE). A set of services and tools that support the creation, use, and maintenance of distributed applications in a heterogeneous computing environment. DCE is independent of the operating system and network; it provides interoperability and portability across heterogeneous platforms.

distributed data facility (DDF). A set of DB2 UDB for OS/390 components through which DB2 UDB for OS/390 communicates with another RDBMS.

distributed directory database. The complete listing of all the resources in the network as maintained in the individual directories scattered throughout an APPN network. Each node has a piece of the complete directory, but it is not necessary for any one node to have the entire list. Entries are created, modified, and deleted through system definition, operator action, automatic registration, and ongoing network search procedures. Synonym for *distributed network directory*.

distributed network directory. See *distributed directory database*.

distributed relational database. A database whose tables are stored on different but interconnected computing systems.

Distributed Relational Database Architecture (DRDA). The architecture that defines formats and protocols for providing transparent access to remote data. DRDA defines two types of functions, the application requester function and the application server function.

distributed request. In a federated database system, an SQL query directed to two or more data sources.

distributed unit of work (DUOW). A unit of work that allows SQL statements to be submitted to multiple relational database management systems, but no more than one system per SQL statement.

DJRA tool. A database administration tool that you can use to perform various replication administration tasks. Unlike the Control Center, the DJRA tool can also be used to administer replication for non-IBM databases. Contrast with *Control Center*.

DLC. See *data link control*.

DLU. See *dependent logical unit*.

DML. See *data manipulation language*.

DMS table space. See *database-managed space table space*.

DNS. See *domain name system*.

Document Access Definition (DAD). A definition that is used to enable an XML Extender column of an XML collection, which is XML formatted.

document model. The definition of the structure of a document in terms of the sections that it contains. The Text Extender uses a document model when indexing.

domain name. The name by which TCP/IP applications refer to a TCP/IP host within a TCP/IP network. A domain name consists of a sequence of names separated by dots.

domain name server (DNS). A TCP/IP network server that manages a distributed directory that is used to map TCP/IP host names to IP addresses.

domain name system. The distributed database system used by TCP/IP to map human-readable machine names into IP addresses.

Domino™ *Go* Web server. The Web server offered by Lotus® Corp. and IBM, that offers both regular and secure connections. ICAPI and GWAPI are the interfaces provided with this server.

double-byte character large object (DBCLOB). A sequence of double-byte characters, where the size can be up to 2 gigabytes. A data type that can be used to store large double-byte text objects. Also called double-byte character large object string. Such a string always has an associated code page.

double-byte character set (DBCS). A set of characters in which each character is represented by two bytes.

double-precision floating point number. In SQL, a 64-bit approximate representation of a real number.

drain. In DB2 UDB for OS/390, the act of acquiring a locked resource by quiescing access to that object.

drain lock. In DB2 UDB for OS/390, a lock on a claim class that prevents a claim from occurring.

DRDA. See *Distributed Relational Database Architecture*.

DRDA access. In DB2 UDB for OS/390, a method of accessing distributed data by which you can connect to another location, using an SQL statement, to execute packages that were previously bound at that location. The SQL CONNECT or three-part name statement is used to identify application servers, and SQL statements are executed using packages that were previously bound at those servers. Contrast with *private protocol access*.

DSN. (1) The default subsystem name for DB2 UDB for OS/390. (2) The name of the TSO command processor of DB2 UDB for OS/390. (3) The first three characters of the names of DB2 UDB for OS/390 modules and macros.

DUOW. See *distributed unit of work*.

duration. In SQL, a number that represents an interval of time. See *date duration*, *labeled duration*, and *time duration*.

dynamic bind. A process by which SQL statements are bound as they are entered. See also *bind*.

dynamic SQL. SQL statements that are prepared and run within a running program. In dynamic SQL, the SQL source is contained in host language variables rather than being coded into the program. The SQL statement might change several times while the program is running.

E

EA-enabled table space. In DB2 UDB for OS/390, 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.

EBCDIC. Extended binary-coded decimal interchange code. A coded character set of 256 8-bit characters.

EDM pool. In DB2 UDB for OS/390, a pool of main storage that is used for database descriptors, application plans, authorization cache, application packages, and dynamic statement caching.

EID. Event identifier.

embedded SQL. SQL statements coded within an application program. See *static SQL*.

EN. See *end node*.

enable. To prepare a database, a text table, or a text column for use by the Text Extender.

enclave. In Language Environment (which is used by DB2 UDB for OS/390), an independent collection of routines, one of which is designated as the main routine. An enclave is similar to a program or run unit.

encoding scheme. A set of rules to represent character data.

end node (EN). In APPN, a node that supports sessions between its local control point and the control point in an adjacent network node.

enhanced conflict detection. Conflict detection that guarantees data integrity among all replicas and the source table. The Apply program locks all replicas or user tables in the subscription set against further transactions. It begins detection after all changes made prior to locking have been captured. See also *conflict detection*, *standard conflict detection*, and *row-replica conflict detection*.

environment handle. A handle that identifies the global context for database access. All data that is pertinent to all objects in the environment is associated with this handle.

environment profile. A script that is provided with the Text Extender that contains settings for environment variables.

EOM. End of memory.

EOT. End of task.

equijoin. A join in which the predicate contains an equals operator, for example, T1.C1 = T2.C2.

error page range. A range of pages that are considered to be physically damaged. DB2 UDB for OS/390 does not allow users to access any pages that fall within this range.

escape character. The symbol that is used to enclose an SQL delimited identifier. The escape character is the double quotation mark, except in COBOL applications, where the user assigns the symbol, which is either a double quotation mark or an apostrophe.

ESDS. In an OS/390 environment, entry sequenced data set.

ESMT. In the OS/390 environment, the external subsystem module table of IMS.

EUC. See *Extended UNIX® Code*.

event monitor. A database object for monitoring and collecting data on database activities over a period of time.

event timing. In DB2 replication, the most precise method of controlling when to start a subscription cycle. It requires that you specify an event and the time when you want the event processed. Contrast with *interval timing* and *on-demand timing*.

exception table. In DB2 UDB for OS/390, a table that holds rows that violate referential constraints or table check constraints that the CHECK DATA utility finds.

exclusive lock. A lock that prevents concurrently executing application processes from accessing database data.

executable statement. An SQL statement that can be embedded in an application program, dynamically prepared and executed, or issued interactively.

exit routine. A program that receives control from another program (such as DB2 UDB for OS/390) to perform specific functions.

explain. To capture detailed information about the access plan that was chosen by the SQL compiler to resolve an SQL statement. The information describes the decision criteria used to choose the access plan.

explainable statement. An SQL statement for which the explain operation can be performed. Explainable statements are SELECT, UPDATE, INSERT, DELETE, and VALUES.

explained statement. An SQL statement for which an explain operation was performed.

explained statistics. Statistics for a database object that was referenced in an SQL statement at the time that the statement was explained.

explain snapshot. A capture of the current internal representation of an SQL query and related information. This information is required by the Visual Explain tool.

explicit hierarchical locking. In DB2 UDB for OS/390, locking that is used to make the parent-child relationship between resources known to IRLM. This kind of locking avoids global locking overhead when no inter-DB2 interest exists on a resource.

explicit privilege. A privilege that has a name and is held as the result of SQL GRANT and REVOKE statements, for example, the SELECT privilege. Contrast with *implicit privilege*.

export. To copy data from database manager tables to a file using formats such as PC/IXF, DEL, WSF, or ASC. Contrast with *import*.

exposed name. A correlation name, a table, or a view name specified in a FROM clause for which a correlation name is not specified.

expression. An SQL operand or a collection of operators and operands that yields a single value.

extended recovery facility (XRF). In an OS/390 environment, 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.

Extended UNIX Code (EUC). A protocol that can support sets of characters from 1 to 4 bytes in length. EUC is a means of specifying a collection of code pages rather than actually being a code page encoding scheme itself. This is the UNIX alternative to the PC double-byte (DBCS) code page encoding schemes.

extent. An allocation of space, within a container of a table space, to a single database object. This allocation consists of multiple pages.

extent map. A metadata structure stored within a table space that records the allocation of extents to each object in the table space.

external CCD table. In DB2 replication, a CCD table that can be subscribed to directly because it is a registered replication source. It has its own row in the register table, where it is referenced as SOURCE_OWNER and SOURCE_TABLE. Contrast with *internal CCD table*.

external function. In DB2 UDB for OS/390, 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 *sourced function* and *built-in function*.

external routine. In DB2 UDB for OS/390, a user-defined function or stored procedure that is based on code that is written in an external programming language.

F

fact table. In the OLAP Starter Kit, a table, or in many cases a set of four tables, in DB2 that contains all data values for a relational cube.

failed member state. In DB2 UDB for OS/390, a state of a member of a data sharing group. When a member fails, the XCF permanently records the failed member state. This state usually means that the member's task, address space, or MVS system terminated before the state changed from active to quiesced.

fallback. The process of returning to a previous release of DB2 UDB for OS/390 after attempting or completing migration to a current release.

false global lock contention. In DB2 UDB for OS/390, an indication of contention from the coupling facility when multiple lock names are hashed to the same indicator and when no real contention exists.

fast communication manager (FCM). A group of functions that provide internodal communication support.

federated database system. (1) A DB2 server and multiple data sources that the server sends queries to. In a federated database system, a client application can join data that is distributed across multiple database management systems using a single SQL statement and view the data as if it were local. (2) A distributed computing system that consists of:

- A DB2 server, called a *federated server.*
- Multiple data sources to which the federated server sends queries.

  Each data source consists of an instance of a relational database management system plus the database or databases that the instance supports.

The data sources are semi-autonomous. For example, the federated server can send queries to Oracle data sources at the same time that Oracle applications are accessing these data sources.

fenced. Pertaining to a type of user-defined function or stored procedure that is defined to protect the DBMS from modifications by the function. The DBMS is isolated from the function or stored procedure by a barrier. Contrast with *not-fenced*.

field procedure. In DB2 UDB for OS/390, a user-written exit routine that is designed to receive a single value and transform (encode or decode) it in any way that the user can specify.

file reference variable. A host variable that is used to indicate that data resides in a file on the client rather than in a client memory buffer.

file server. A workstation that runs the NetWare operating system software and acts as a network server. DB2 uses the file server to store DB2 server address information, which a DB2 client retrieves to establish an IPX/SPX client-server connection.

filter factor. In DB2 UDB for OS/390, a number between zero and one that estimates the proportion of rows in a table for which a predicate is true.

first failure service log. A file (`db2diag.log`) that contains diagnostic messages, diagnostic data, alert information, and related dump information. This file is used by database administrators.

fixed-length string. A character or graphic string whose length is specified and cannot be changed. Contrast with *varying-length string*.

flagger. A precompiler option that identifies SQL statements in applications that do not conform to selected validation criteria (for example, the ISO/ANSI SQL92 entry-level standard).

flat file interface. A set of Net.Data built-in functions that let you read and write data from plain-text files.

foreign update. An update that was applied to a target table and replicated to the local table.

forward log recovery. The third phase of restart processing during which DB2 UDB for OS/390 processes the log in a forward direction to apply all REDO log records.

forward recovery. A process used to roll forward a database or table space. It allows a restored database or table space to be rebuilt to a specified point in time by applying the changes recorded in the database log.

free space. In DB2 UDB for OS/390, the total amount of unused space in a page. The space that is not used to store records or control information is free space.

full outer join. The result of an SQL join operation that includes the matched rows of both tables that are being joined and preserves the unmatched rows of both tables. See *join*.

full refresh. In DB2 replication, a process in which all of the data of interest in a user table is copied to the target table, replacing existing data. Contrast with *differential refresh*.

fullselect. A subselect, a values-clause, or a number of both that are combined by set operators.

fully qualified LU name. See *network-qualified name*.

function. (1) A mapping, embodied as a program (the function body), that can be invoked by using zero or more input values (arguments) to a single value (the result). (2) In DB2 UDB for OS/390, a specific purpose of an entity or its characteristic action such as a column function or scalar function. Functions can be user-defined, built-in, or generated by DB2 UDB for OS/390.

function body. The piece of code that implements a function.

function definer. In DB2 UDB for OS/390, the authorization ID of the owner of the schema of the function that is specified in the CREATE FUNCTION statement.

function family. A set of functions with the same function name. The context determines whether the usage refers to a set of functions within a particular schema, or all the relevant functions with the same name within the current function path.

function implementer. In DB2 UDB for OS/390, the authorization ID of the owner of the function program and function package.

function invocation. The use of a function together with any argument values being passed to the function body. The function is invoked by its name.

function package. In DB2 UDB for OS/390, a package that results from binding the DBRM for a function program.

function package owner. In DB2 UDB for OS/390, the authorization ID of the user who binds the function program's DBRM into a function package.

function path. An ordered list of schema names that restricts the search scope for unqualified function invocations and provides a final arbiter for the function selection process.

function path family. All the functions of the given name in all the schemas identified (or used by default) in the user's function path.

function resolution. The process, internal to the DBMS, for which a particular function instance is selected for invocation. The function name, the data types of the arguments, and the function path are used to make the selection. Synonym for *function selection*.

function selection. See *function resolution*.

function shipping. The shipping of the subsections of a request to the specific node that contains the applicable data.

function signature. The logical concatenation of a fully qualified function name with the data types of all of its parameters. Each function in a schema must have a unique signature.

function template. In a federated database, a partial function that has no executable code. The user maps it to a data source function, so that the data source function can be invoked from the federated server.

G

gap. In DB2 replication, a situation in which the Capture program is not able to read a range of log or journal records, so there is potential loss of change data.

GBP. Group buffer pool.

GBP-dependent. In DB2 UDB for OS/390, the status of a page set or page set partition that is dependent on the group buffer pool. Either read/write interest is active among DB2 subsystems for this page set, or the page set has changed pages in the group buffer pool that are not yet cast out to DASD.

generalized trace facility (GTF). In an OS/390 environment, a service program that records significant system events such as I/O interrupts, SVC interrupts, program interrupts, or external interrupts.

generic resource name. In an OS/390 environment, 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 Parallel Sysplex environment.

getpage. An operation in which DB2 UDB for OS/390 accesses a data page.

GIMSMP. In an OS/390 environment, the load module name for the System Modification Program/Extended, a basic tool for installing, changing, and controlling changes to programming systems.

global lock. In DB2 UDB for OS/390, a lock that provides concurrency control within and among DB2 subsystems. The scope of the lock is across all DB2 subsystems of a data sharing group.

global lock contention. Conflicts on locking requests between different DB2 UDB for OS/390 members of a data sharing group when those members are trying to serialize shared resources.

global table lock. A table lock that is acquired on all nodes in a table's nodegroup.

global transaction. A unit of work in a distributed transaction processing environment in which multiple resource managers are required.

governor. See *resource limit facility*.

grant. To give a privilege or authority to an authorization ID.

graphic character. A DBCS character.

graphic string. A sequence of DBCS characters.

gross lock. In DB2 UDB for OS/390, the *shared*, *update*, or *exclusive* mode locks on a table, partition, or table space.

group. (1) A logical organization of users that have IDs according to activity or resource access authority. (2) In Satellite Edition, a collection of satellites that share characteristics such as database configuration and the application that runs on the satellite.

group buffer pool duplexing. In an OS/390 environment, the ability to write data to two instances of a group buffer pool structure: a *primary group buffer pool* and a *secondary group buffer pool*. OS/390 publications refer to these instances as the ″old″ (for primary) and ″new″ (for secondary) structures.

group name. In an OS/390 environment, the XCF identifier for a data sharing group.

group restart. In an OS/390 environment, a restart of at least one member of a data sharing group after the loss of either locks or the shared communications area.

group scope. See *command scope*.

GTF. See *generalized trace facility*.

GWAPI. Domino *Go* Web server API.

H

handle. (1) A variable that represents an internal structure within a software system. (2) A character string that is created by an extender that is used to represent an image, audio, or video object in a table. A handle is stored for an object in a user table and in administrative support tables. In this way, an extender can link the handle that is stored in a user table with information about the object that is stored in the administrative support tables. (3) A binary value that identifies a text document. A handle is created for each text document in a text column when that column is *enabled* for use by the Text Extender.

hash partitioning. A partitioning strategy in which a hash function is applied to the partitioning key value to determine the database partition to which the row is assigned.

hiperspace. In an OS/390 environment, 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, you bring the data into the address space in 4-KB blocks.

home address space. In an OS/390 environment, the area of storage that OS/390 currently recognizes as *dispatched*.

hop. In APPN, a portion of a route that has no intermediate nodes. A hop consists of a single transmission group connecting adjacent nodes.

host. In TCP/IP, any system that has at least one Internet address associated with it.

host computer. (1) In a computer network, a computer that provides services such as computation, database access, and network control functions. (2) The primary or controlling computer in a multiple computer installation.

host identifier. A name declared in the host program.

host language. Any programming language in which you can embed SQL statements.

host node. In SNA, a subarea node that contains a system services control point (SSCP), for example, an IBM System/390® computer with MVS and VTAM.

host program. A program written in a host language that contains embedded SQL statements.

host structure. In an application program, a structure that is referred to by embedded SQL statements.

host variable. In an application host program, a variable that is referred to by embedded SQL statements. Host variables are programming variables in the application program and are the primary mechanism for transmitting data between tables in the database and application program work areas.

HSM. In an OS/390 environment, hierarchical storage manager.

I

ICAPI. Internet Connection API.

ICF. In an OS/390 environment, integrated catalog facility.

IDCAMS. In an OS/390 environment, 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.

IDCAMS LISTCAT. In an OS/390 environment, a facility for obtaining information that is contained in the access method services catalog.

identify. A request that an attachment service program (in an address space that is separate from DB2 UDB for OS/390) issues through the MVS subsystem interface to inform DB2 UDB for OS/390 of its existence and to initiate the process of becoming connected to DB2.

IFCID. In DB2 UDB for OS/390, instrumentation facility component identifier.

IFI. In DB2 UDB for OS/390, instrumentation facility interface.

IFI call. In DB2 UDB for OS/390, an invocation of the instrumentation facility interface (IFI) by means of one of its defined functions.

IFP. In an OS/390 environment, IMS Fast Path.

ILU. See *independent logical unit*.

image copy. An exact reproduction of all or part of a table space. DB2 UDB for OS/390 provides utility programs to make full image copies (to copy the entire table space) or incremental image copies (to copy only those pages that were modified since the last image copy).

implicit privilege. A privilege that accompanies 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.

import. To copy data from an external file, using formats such as PC/IXF, DEL, WSF or ASC, into database manager tables. Contrast with *export*.

import metadata. The process of bringing metadata into the Data Warehouse Center, either dynamically (from the user interface) or in batch.

import utility. Transactional utility that inserts user-supplied record data into a table. Contrast with *load utility*.

IMS. Information Management System.

IMS attachment facility. A DB2 UDB for OS/390 subcomponent that uses OS/390 subsystem interface (SSI) protocols and cross-memory linkage to process requests from IMS to DB2 UDB for OS/390 and to coordinate resource commitment.

IMS DB. Information Management System Database.

IMS TM. Information Management System Transaction Manager.

in-abort. A status of a unit of recovery. If DB2 UDB for OS/390 fails after a unit of recovery begins to be rolled back, but before the process is completed, DB2 UDB for OS/390 continues to back out the changes during restart.

in-commit. A status of a unit of recovery. If DB2 UDB for OS/390 fails after beginning its two-phase commit processing, it “knows,” when restarted, that changes made to data are consistent.

incremental bind. 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. See also *bind*.

independent. In DB2 UDB for OS/390, an object (row, table, or table space) that is neither a parent nor a dependent of another object.

independent logical unit (ILU). A logical unit that is able to activate an LU-to-LU session without assistance from a system services control point (SSCP). An ILU does not have an SSCP-to-LU session. Contrast with *dependent logical unit*.

index. A set of pointers that are logically ordered by the values of a key. Indexes provide quick access to data and can enforce uniqueness on the rows in the table.

index file. A file that contains indexing information used by the Video Extender to find a *shot* or an individual frame in a video clip.

index key. The set of columns in a table used to determine the order of index entries.

index partition. The part of an index that is associated with a table partition at a given node. An index defined on a table is implemented by multiple index partitions, one per table partition.

index sargable predicates. SQL predicates that are applied to index entries in index leaf pages to reduce the number of index entries that qualify the SQL request. They help reduce the number of data rows accessed.

index space. In DB2 UDB for OS/390, a page set that is used to store the entries of one index.

index specification. In a federated database system, a set of metadata that pertains to a data source table. This metadata is made up of information that an index definition typically contains, for example, which column or columns to search in order to retrieve information quickly. The user might supply the federated server with this metadata if the table has no index or if it has an index that is unknown to the federated server. The purpose of the metadata is to facilitate retrieval of the table’s data.

indicator column. In DB2 UDB for OS/390, a 4-byte value that is stored in a base table in place of an LOB column.

indicator variable. A variable 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.

indoubt. A status of a unit of recovery. If DB2 UDB for OS/390 fails after it finishes its phase 1 commit processing and before it starts 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 UDB for OS/390 lacks the information that it needs to make this decision, the status of the unit of recovery is *indoubt* until DB2 UDB for OS/390 obtains this information from the coordinator. More than one unit of recovery can be indoubt at restart.

indoubt resolution. In DB2 UDB for OS/390, the process of resolving the status of an indoubt logical unit of work to either the committed or the rollback state.

indoubt transaction. A transaction in which one phase of a two-phase commit completes successfully but the system fails before a subsequent phase can complete.

inflight. A status of a unit of recovery. If DB2 UDB for OS/390 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 *inflight*.

information catalog. The database, managed by the Information Catalog Manager, that contains descriptive data (*business metadata*) that helps users identify and locate data and information that is available to them in the organization. The information catalog also contains some *technical metadata*.

Information Catalog Manager. An application for organizing, maintaining, finding, and using business information.

Glossary

inheritance. The passing of class resources or attributes from a parent class downstream in the class hierarchy to a child class.

initialization fullselect. The first fullselect in a recursive common table expression that gets the direct children of the initial value from the source table.

inner join. A join method in which a column that is not common to all of the tables being joined is dropped from the resultant table. Contrast with *outer join*.

inoperative package. A package that cannot be used because a function that it depends on has been dropped. Such a package must be explicitly rebound. Contrast with *invalid package*.

inoperative trigger. A trigger that depends on an object that has been dropped or made inoperative or on a privilege that has been revoked.

inoperative view. A view that is no longer usable because one of the following situations occurs:

- SELECT privilege on a table or view that the view is dependent on is revoked from the definer of the view.
- An object on which the view definition is dependent was dropped (or possibly made inoperative in the case of another view).

insert rule. A condition enforced by the database manager that must be met before a row can be inserted into a table.

insert trigger. In DB2 UDB for OS/390, a trigger that is defined with the triggering SQL operation INSERT.

install. The process of preparing a DB2 UDB for OS/390 subsystem to operate as an OS/390 subsystem.

installation verification scenario. A sequence of operations that exercises the main DB2 UDB for OS/390 functions and tests whether DB2 UDB for OS/390 was correctly installed.

instance. (1) See *database manager instance*. (2) A logical DB2 extender server environment. You can have several instances of DB2 extenders server on the same workstation, but only one instance for each DB2 instance.

instrumentation facility component identifier (IFCID). In DB2 UDB for OS/390, 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.

instrumentation facility interface (IFI). A programming interface that enables programs to obtain online trace data about DB2 UDB for OS/390, to submit DB2 UDB for OS/390 commands, and to pass data to DB2 UDB for OS/390.

Interactive System Productivity Facility (ISPF). In an OS/390 environment, an IBM licensed program that provides interactive dialog services.

inter-DB2 R/W interest. In DB2 UDB for OS/390, a property of data in a table space, index, or partition that has been opened by more than one member of a data sharing group and that has been opened for writing by at least one of those members.

intermediate network node. In APPN, a node that is part of a route between an origin logical unit (OLU) and a destination logical unit (DLU) but that neither contains the OLU or the DLU nor serves as the network server for either the OLU or DLU.

internal CCD table. A CCD table that cannot be subscribed to directly. It does not have its own row in the register table; it is referenced as CCD_OWNER and CCD_TABLE in the row for the associated replication source. Contrast with *external CCD table*.

internal resource lock manager (IRLM). In an OS/390 environment, a subsystem that DB2 UDB for OS/390 uses to control communication and database locking.

Internet Protocol (IP). A protocol used to route data from its source to its destination in an Internet environment.

Internetwork Packet Exchange (IPX). A connectionless datagram protocol, used in a NetWare LAN environment, to transfer data to a remote node. IPX makes a best-effort attempt to send data packets, but does not guarantee reliable delivery of the data.

inter-partition parallelism. The ability to perform multiple database operations (such as index creation, database load, and SQL queries) at the same time across multiple partitions of a partitioned database. Contrast with *intra-partition parallelism*.

Inter-Process Communication (IPC). A mechanism of an operating system that allows processes to communicate with each other.

interval timing. In DB2 replication, the simplest method of controlling when to start a subscription cycle. You must specify a date and a time for a subscription cycle to start, and set a time interval that describes how frequently you want the subscription cycle to run. Contrast with *event timing* and *on-demand timing*.

intra-partition parallelism. The ability to perform multiple database operations (such as index creation, database load, SQL queries) at the same time within a single database partition. Contrast with *inter-partition parallelism.*

intra-query parallelism;. The ability to process parts of a single query at the same time using either intra-partition parallelism, inter-partition parallelism, or both.

invalid package. A package that becomes invalid when an object that the package depends on is dropped. (The object is of a type other than function, for example, index.) Such a package is implicitly rebound upon invocation. Contrast with *inoperative package*.

invariant character set. In DB2 UDB for OS/390, (1) a character set, such as the syntactic character set, whose code point assignments do not change from code page to code page; (2) a minimum set of characters that is available as part of all character sets.

I/O parallelism. See *parallel I/O*.

IP. See *Internet Protocol*.

IP address. A 4-byte value that uniquely identifies a TCP/IP host.

IPX. Internetwork Packet Exchange.

IRLM. In DB2 UDB for OS/390, internal resource lock manager.

ISAPI. Microsoft® Internet Server API.

isolation level. An attribute that defines the degree to which an application process is isolated from other concurrently executing application processes.

ISPF. In an OS/390 environment, Interactive System Productivity Facility.

ISPF/PDF. In an OS/390 environment, Interactive System Productivity Facility/Program Development Facility.

J

JCL. See *job control language*.

JES. See *Job Entry Subsystem*.

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 jobs.

job scheduler. A program used to automate certain tasks for running and managing database jobs.

join. An SQL relational operation that allows retrieval of data from two or more tables based on matching column values.

K

key. A column or an ordered collection of columns that are identified in the description of a table, index, or referential constraint.

key-sequenced data set (KSDS). In an OS/390 environment, a VSAM file or data set whose records are loaded in key sequence and controlled by an index.

key-value based partitioning strategy. A strategy for assigning rows in a table to database partitions. Rows are assigned based on the values of the partitioning key columns.

keyword. (1) One of the predefined words of a computer, command language, or an application. (2) A name that identifies an option used in an SQL statement.

KSDS. See *key-sequenced data set*.

L

labeled duration. A number that represents a duration of years, months, days, hours, minutes, seconds, or microseconds.

Language Environment®. A module that provides access from a Net.Data macro to an external data source, such as DB2, or to a programming language, such as Perl.

large object (LOB). A sequence of bytes with a length of up to 2 gigabytes. It can be any of three types: BLOB (binary), CLOB (single-byte character or mixed) or DBCLOB (double-byte character).

latch. A DB2 UDB for OS/390 internal mechanism for controlling concurrent events or the use of system resources.

LCID. In an OS/390 environment, log control interval definition.

LDS. See *linear data set*.

leaf page. In DB2 UDB for OS/390, a page that contains pairs of keys and RIDs and that points to actual data. Contrast with *nonleaf page*.

left outer join. In DB2 UDB for OS/390, 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 *join* and *right outer join*.

length attribute. A value associated with a string that represents the declared fixed length or maximum length of the string.

LEN node. See *low-entry networking node*.

linear data set (LDS). In an OS/390 environment, 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.

linkage editor. 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.

link-edit. In DB2 UDB for OS/390, the action of creating a loadable computer program using a linkage editor.

list prefetch. An access method that takes advantage of prefetching even in queries that do not access data sequentially. This is done by scanning the index and collecting RIDs in advance of accessing any data pages. These RIDs are then sorted, and data is prefetched using this list.

list structure. In an OS/390 environment, a coupling facility structure that lets data be shared and manipulated as elements of a queue.

Live Connection. A Net.Data component that consists of a Connection Manager and multiple cliettes. Live Connection manages the reuse of database and Java® virtual machine connections.

L-lock. See *logical lock*.

load copy. A backup image of data that was loaded at a previous time and can be restored during roll-forward recovery.

load module. A program unit that is suitable for loading into main storage for execution. The output of a linkage editor.

load utility. A nontransactional utility that performs block updates of table data. Contrast with *import utility*.

LOB. See *large object*.

LOB locator. A mechanism that allows an application program to manipulate a large object (LOB) value in the database system. An LOB locator is a simple token value that represents a single LOB value. An application program retrieves an LOB locator into a host variable and can then apply SQL functions to the associated LOB value using the locator.

LOB lock. In DB2 UDB for OS/390, a lock on an LOB value.

LOB table space. In DB2 UDB for OS/390, a table space that contains all the data for a particular LOB column in the related base table.

local. A way of referring to any object that the local subsystem maintains. In DB2 UDB for OS/390, for example, a local table is a table that is maintained by the local DB2 subsystem. Contrast with *remote*.

local database. A database that is physically located on the workstation in use. Contrast with *remote database*.

local database directory. A directory where a database physically resides. Databases that are displayed in the local database directory are located on the same node as the system database directory.

locale. In DB2 UDB for OS/390, the definition of a subset of a user's environment that combines characters that are defined for a specific language and country, and a CCSID.

local lock. A lock that provides intra-DB2 concurrency control, but not inter-DB2 concurrency control; its scope is a single DB2 UDB for OS/390 system.

local subsystem. The unique RDBMS to which the user or application program is directly connected (in the case of DB2 UDB for OS/390, by one of the DB2 UDB for OS/390 attachment facilities).

local table lock. A table lock that is acquired only on a single database partition.

local update. An update to the base table, not to the replica.

location name. The name by which DB2 UDB for OS/390 refers to a particular DB2 subsystem in a network of subsystems. Contrast with *LU name*.

location path. A subset of the abbreviated syntax of the location path defined by XPath. A sequence of XML tags to identify an XML element or attribute. It is used in extracting user-defined functions to identify the subject to be extracted, and it is used in the Text Extender's search user-defined functions to identify the search criteria.

locator. See *LOB locator*.

lock. (1) A means of serializing events or access to data. (2) A means of preventing uncommitted changes made by one application process from being perceived by another application process and for preventing one application process from updating data that is being accessed by another process. (3) A means of controlling concurrent events or access to data. DB2 UDB for OS/390 locking is performed by the IRLM.

lock duration. The interval over which a DB2 UDB for OS/390 lock is held.

lock escalation. In the database manager, the response that occurs when the number of locks issued for one agent exceeds the limit specified in the database configuration; the limit is defined by the MAXLOCKS configuration parameter. During a lock escalation, locks are freed by converting locks on rows of a table into one lock on a table. This is repeated until the limit is no longer exceeded.

locking. The mechanism used by the database manager to ensure the integrity of data. Locking prevents concurrent users from accessing inconsistent data.

lock mode. A representation for the type of access that concurrently running programs can have to a resource that a DB2 UDB for OS/390 lock is holding.

lock object. The resource that is controlled by a DB2 UDB for OS/390 lock.

lock parent. For explicit hierarchical locking in DB2 UDB for OS/390, a lock that is held on a resource that has child locks that are lower in the hierarchy; usually, the table space or partition intent locks are the parent locks.

lock promotion. The process of changing the size or mode of a DB2 UDB for OS/390 lock to a higher level.

lock size. The amount of data controlled by a DB2 UDB for OS/390 lock on table data; the value can be a row, a page, an LOB, a partition, a table, or a table space.

lock structure. In DB2 UDB for OS/390, a coupling facility data structure that is composed of a series of lock entries to support shared and exclusive locking for logical resources.

log. (1) A file used to record changes made in a system. (2) A collection of records that describe the events that occur during DB2 UDB for OS/390 execution and that indicate their sequence. The information thus recorded is used for recovery in the event of a failure during DB2 UDB for OS/390 execution. (3) See *database log*.

log head. The oldest written log record in the active log.

logical claim. In DB2 UDB for OS/390, a claim on a logical partition of a nonpartitioning index.

logical drain. In DB2 UDB for OS/390, a drain on a logical partition of a nonpartitioning index.

logical index partition. In DB2 UDB for OS/390, the set of all keys that reference the same data partition.

logical lock (L-lock). In DB2 UDB for OS/390, the lock type that transactions use to control intra-DB2 and inter-DB2 data concurrency between transactions. Contrast with *physical lock*.

logical node. A node on a processor that has more than one node assigned to it. See also *node*.

logical operator. A keyword that specifies how multiple search conditions are to be evaluated (AND, OR) or if the logical sense of a search condition is to be inverted (NOT).

logical page list (LPL). In DB2 UDB for OS/390, a list of pages that are in error and that cannot be referenced by applications until the pages are recovered. The page is in logical error, because the actual media (coupling facility or DASD) might not contain any errors. Usually a connection to the media has been lost.

logical partition. In DB2 UDB for OS/390, a set of key or RID pairs in a nonpartitioning index that are associated with a particular partition.

logical recovery pending (LRECP). In DB2 UDB for OS/390, the state in which the data and the index keys that refer to the data are inconsistent.

logical unit (LU). (1) In SNA, a port through which an end user accesses the SNA network to communicate with another end user. An LU is capable of supporting many sessions with other LUs. (2) In an OS/390 environment, an access point through which an application program accesses the SNA network in order to communicate with another application program. See also *LU name*.

logical unit 6.2 (LU 6.2). The LU type that supports sessions between two applications using APPC.

logical unit of work (LUW). The processing that a program performs between synchronization points.

logical unit of work identifier (LUWID). In an OS/390 environment, 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.

log initialization. The first phase of restart processing during which DB2 UDB for OS/390 attempts to locate the current end of the log.

log partition. The log file on each database partition that records database activity for that database partition.

log record. A record of an update to a database performed during a unit of work. This record is written after the log tail of the active log.

log record sequence number (LRSN). A number that DB2 UDB for OS/390 generates and associates with each log record. The LRSN is also used for page versioning. The LRSNs that a particular DB2 UDB for OS/390 data sharing group generates form a strictly increasing sequence for each DB2 log and a strictly increasing sequence for each page across the data sharing group.

log table. A table created by the Text Extender that contains information about which text documents are to be indexed.

log tail. The log record that was written most recently in an active log.

log truncation. In DB2 UDB for OS/390, 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.

long string. (1) A varying-length string whose maximum length is greater than 254 bytes. (2) In DB2 UDB for OS/390, 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.

long table space. A table space that can store only long string or large object (LOB) data.

low-entry networking node (LEN node). A type 2.1 node that supports independent LU protocols but does not support CP to CP sessions. It can be a peripheral node attached to a boundary node in a subarea network, an end node attached to an APPN network node in an APPN network, or a peer-connected node directly attached to another LEN node or APPN end node.

LPL. See *logical page list*.

LRECP. See *logical recovery pending*.

LRH. In DB2 UDB for OS/390, log record header.

LRSN. See *log record sequence number*.

LU. See *logical unit*.

LU name. In an OS/390 environment, the name by which VTAM refers to a node in a network. Contrast with *location name*.

LU 6.2. See *logical unit 6.2*.

LU type. The classification of a logical unit in terms of the specific subset of SNA protocols and options that it supports for a given session, specifically:

- The values allowed in the session activation request
- The usage of data stream controls, function management headers, request unit parameters, and sense data values
- Presentation services protocols such as those associated with function management headers

LUW. See *logical unit of work*.

LUWID. See *logical unit of work identifier*.

M

mapped conversation. In APPC, a conversation between two transaction programs (TPs) using the APPC mapped conversation API. In typical situations, end-user TPs use mapped conversation, and service TPs use basic conversations. Either type of program can use either type of conversation. Contrast with *basic conversation*.

masking character. A character used to represent optional characters at the front, middle, and end of a search term. Masking characters are normally used for finding variations of a term in a precise index.

Glossary

materialize. In DB2 UDB for OS/390, (1) The process of putting rows from a view or nested table expression into a work file for additional processing by a query.

(2) The placement of an LOB value into contiguous storage. Because LOB values can be very large, DB2 UDB for OS/390 avoids materializing LOB data until doing so becomes absolutely necessary.

MBCS. See *multi-byte character set*.

member. (1) For DB2, *subscription-set member*. (2) In the OLAP Starter Kit, a method of referencing data through three or more dimensions. An individual data value in a fact table is the intersection of one member from each dimension.

member name. The XCF identifier for a particular DB2 UDB for OS/390 subsystem in a data sharing group.

member scope. See *command scope*.

menu. In DB2 UDB for OS/390, a displayed list of available functions for selection by the operator. A menu is sometimes called a *menu panel*.

metadata. Data that describes the characteristics of stored data; descriptive data. For example, the metadata for a database table might include the name of the table, the name of the database that contains the table, the names of the columns in the table, and the column descriptions, either in technical terms or business terms.

metadata publication process. A process created by the Data Warehouse Center that contains all the steps created after publication to keep the published metadata synchronized with the original metadata.

migration. (1) The process of moving data from one computer system to another without converting the data. (2) Installation of a new version or release of a program to replace an earlier version or release. (3) The process of converting an existing DB2 UDB for OS/390 subsystem to an updated or current release. In this process, you can acquire the functions of the updated or current release without losing the data that you created on the previous release.

mixed-character string. A string containing a mixture of single-byte and multi-byte characters. Also called *mixed data string*.

mixed-data string. See *mixed-character string*.

mobile client. The node, usually a laptop computer, where the mobile enabler, replication source, and target tables used in a mobile environment are located. The mobile replication mode is invoked from the mobile client.

mobile replication enabler. A replication program that starts the mobile replication mode at the mobile client.

mobile replication mode. A mode of replication in which the Capture and Apply programs operate as needed rather than autonomously and continuously. This mode is invoked from the mobile client and allows data to be replicated when the mobile client is available for a connection to the source or target server.

mode. In the Data Warehouse Center, the stage of development of a step, such as development, test, or production.

MODEENT. In an OS/390 environment, a VTAM macro instruction that associates a logon mode name with a set of parameters that represent session protocols. A set of MODEENT macro instructions defines a logon mode table.

modeled statistics. Statistics for a database object that may or may not be referenced in an SQL statement, yet currently exist in an explain model. The object may or may not currently exist in the database.

mode name. (1) In APPC, the name used by the initiator of a session to designate the characteristics desired for the session, such as message length limits, sync point, class of service within the transport network, and session routing and delay characteristics. (2) In an OS/390 environment, a VTAM name for the collection of physical and logical characteristics and attributes of a session.

modify locks. In DB2 UDB for OS/390, an L-lock or P-lock with a MODIFY attribute. A list of these active locks is kept at all times in the coupling facility lock structure. If the requesting subsystem fails, that subsystem’s modify locks are converted to retained locks.

monitoring session. The act of monitoring a database manager or of playing back information from a previously monitored database manager. The DB2 Performance Monitor is used for creating a monitoring session and for selecting which database objects to monitor.

monitor switch. Database manager parameters manipulated by the user to control the type of information and the quantity of information returned in performance snapshots.

MPP. (1) Massively parallel processing. (2) In an OS/390 environment with IMS, message processing program.

MSS. In an OS/390 environment, Mass Storage Subsystem.

MTO. In an OS/390 environment, master terminal operator.

multi-byte character set (MBCS). A set of characters in which each character is represented by 2 or more bytes. Character sets that use only two bytes are more commonly known as *double-byte character sets*.

multidimensional. In the OLAP Starter Kit, a method of referencing data through three or more dimensions. An individual data value in a fact table is the intersection of one member from each dimension.

multidimensional database. In the OLAP Starter Kit, a nonrelational database into which you copy relational data for OLAP analysis.

multi-site update. In DB2 UDB for OS/390, distributed relational database processing in which data is updated in more than one location within a single unit of work.

multitasking. A mode of operation that provides for concurrent performance or interleaved execution of two or more tasks.

must-complete. A state during DB2 UDB for OS/390 processing in which the entire operation must be completed to maintain data integrity.

MVS. Multiple Virtual Storage, which is part of OS/390.

MVS/ESA™. Multiple Virtual Storage/Enterprise Systems Architecture, which is part of OS/390.

N

NAU. See *network addressable unit*.

NDS. See *Network Directory Services*.

negotiable lock. In DB2 UDB for OS/390, a lock whose mode can be downgraded, by agreement among contending users, to be compatible to all. A physical lock is an example of a negotiable lock.

nested table expression. (1) A result table obtained directly or indirectly from one or more other tables through the evaluation of a fullselect that is specified in the FROM clause. (2) In DB2 UDB for OS/390, a subselect in a FROM clause (surrounded by parentheses).

NETID. Network identifier. See *network name*.

network address. An identifier for a node in a network.

network addressable unit (NAU). The origin or the destination of information transmitted by the path control network. An NAU may be a logical unit (LU), physical unit (PU), control point (CP), or system services control point (SSCP). See also *network name*.

Network Directory Services (NDS). A global, distributed, replicated database NetWare that maintains information about, and provides access to, every resource on the network. The NetWare Directory database organizes objects, independent of their physical location, in a hierarchical tree structure called the directory tree.

network identifier (NID). In an OS/390 environment, 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).

network name. In SNA, a symbolic name by which end users refer to a network addressable unit (NAU), a link station, or a link. Synonym for *NETID*.

network node (NN). In APPN, a node on the network that provides distributed directory services, topology database exchanges with other APPN network nodes, and session and routing services. Synonym for *APPN network node*.

network node server. An APPN network node that provides network services for its local logical units and adjacent end nodes.

network-qualified name. The name by which an LU is known throughout an interconnected SNA network. A network-qualified name consists of a network name identifying the individual subnetwork, and a network LU name. Network-qualified names are unique throughout an interconnected network. Also known as the *network-qualified LU name*, or *fully qualified LU name*.

network services. The services within network addressable units that control network operation through SSCP-to-SSCP, SSCP-to-PU, SSCP-to-LU, and CP-to-CP sessions.

nickname. (1) An identifier that a federated server uses to refer to a data source table or view. (2) A name that is defined in a DB2 DataJoiner database to represent a physical database object (such as a table or stored procedure) in a non-IBM database.

NID. See *network identifier*.

NN. See *network node*.

node. (1) In database partitioning, a synonym for database partition. (2) In hardware, a uniprocessor or symmetric multiprocessor (SMP) computer that is part of a clustered system or a massively parallel processing (MPP) system. For example, RS/6000® SP™ is an MPP system that consists of a number of nodes connected by a high-speed network. (3) In communications, an end point of a communications link, or a junction common to two or more links in a network. Nodes can be processors, communication controllers, cluster controllers, terminals, or workstations. Nodes can vary in routing and other functional capabilities.

node directory. A directory that contains information necessary to establish communications from a client workstation to all applicable database servers.

nodegroup. A named group of one or more database partitions.

noncomplete CCD table. In DB2 replication, a CCD table that is empty when it is created and has rows appended to it as changes are made to the source. Contrast with *complete CCD table*.

noncondensed attribute. A table attribute indicating that the table contains a history of changes to the data, not current data. A table that has this attribute set includes more than one row for each key value.

noncondensed CCD table. In DB2 replication, a CCD table that contains the history of changes to the values for a row. This type of table is useful for auditing purposes. Contrast with *condensed CCD table*.

nondelimited ASCII (ASC) format. A file format used to import data. Nondelimited ASCII is a sequential ASCII file with row delimiters used for data exchange with any ASCII product.

nonleaf page. In DB2 UDB for OS/390, 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. Contrast with *leaf page*.

nonpartitioning index. In DB2 UDB for OS/390, any index that is not a partitioning index.

normalization. In databases, the process of restructuring a data model by reducing its relations to their simplest forms.

not-deterministic function. In DB2 UDB for OS/390, a user-defined function whose result is not solely dependent on the values of the input arguments. Successive invocations with the same argument values can produce a different answer. This type of function is sometimes called a *variant* function. Contrast with a *deterministic function* (sometimes called a *not-variant function*), which always produces the same result for the same input.

not-fenced. A type of user-defined function or stored procedure that is defined to be run in the DBMS process. Contrast with *fenced*.

notification process. A process created by the Data Warehouse Center that contains all the steps created for notification when a step completes.

not-variant function. A user-defined function whose result is solely dependent on the values of the input arguments. Successive invocations with the same argument values always produce the same results. Contrast with *variant function*.

NRE. In an OS/390 environment, network recovery element.

NSAPI. Netscape API.

NUL. In C language, a single character that denotes the end of the string.

NULLIF. In DB2 UDB for OS/390, a scalar function that evaluates two passed expressions, returning either NULL if the arguments are equal or the value of the first argument if they are not.

null. In DB2 UDB for OS/390, a value that indicates the absence of information.

nullable. The condition in which a value for a column, function parameter, or result can have an absence of a value. For example, a field for a person's middle initial does not require a value and is considered nullable.

null value. A parameter position for which no value is specified.

NUL-terminated host variable. In DB2 UDB for OS/390, a varying-length host variable in which the end of the data is indicated by the presence of a NUL terminator.

NUL terminator. In C language, 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 an OS/390 environment with 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. In DB2 UDB for OS/390, data object identifier.

object. (1) Anything that can be created or manipulated with SQL—for example, tables, views, indexes, or packages. (2) In object-oriented design or programming, an abstraction consisting of data and operations associated with that data. (3) For NetWare, an entity that is defined on the network and thus given access to the file server.

object property. A property that identifies a category of information that is associated with an object. A NetWare bindery object can be assigned one or more properties. The DB2 server instance object has an object property, NET_ADDR, which denotes the location of the record within the object.

object type. (1) A 2-byte number that classifies an object in the bindery on a NetWare file server. 062B represents the DB2 database server object type. (2) A categorization or grouping of object instances that share similar behaviors and characteristics.

ODBC. See *Open Database Connectivity*.

ODBC driver. A driver that implements ODBC function calls and interacts with a data source.

offline backup. A backup of the database or table space that was made when the database or table space was not being accessed by applications. The Backup Database utility acquires exclusive use of the database until the backup is complete. Contrast with *online backup*.

offline restore. A restoration of a copy of a database or table space from a backup. The Backup Database utility has exclusive use of the database until the restore is completed. Contrast with *online restore*.

OLAP. See *online analytical processing*.

Glossary

on-demand timing. A method for controlling the timing of replication for occasionally connected systems. Requires that you use the ASNSAT program to operate the Capture and Apply programs. Contrast with *event timing* and *interval timing*.

online analytical processing (OLAP). In the OLAP Starter Kit, a multidimensional, multi-user, client server computing environment for users who need to analyze consolidated enterprise data in real time.

online backup. A backup of the database or table space that is made while the database or table space is being accessed by other applications. Contrast with *offline backup*.

online monitor. See *Performance Monitor*.

online restore. A restoration of a copy of a database or table space while the database or table space is being accessed by other applications. Contrast with *offline restore*.

Open Database Connectivity (ODBC). An API that allows access to database management systems using callable SQL, which does not require the use of an SQL preprocessor. The ODBC architecture allows users to add modules, called *database drivers*, that link the application to their choice of database management systems at run time. Applications do not need to be linked directly to the modules of all the supported database management systems.

operand. An entity on which an operation is performed.

optimized SQL text. SQL text, produced by the Explain facility, that is based on the query actually used by the optimizer to choose the access plan. This query is supplemented and rewritten by the various components of the SQL compiler during statement compilation. The text is reconstructed from its internal representation, and differs from the original SQL text. The optimized statement produces the same result as the original statement.

optimizer. A component of the SQL compiler that chooses an access plan for a data manipulation language statement by modeling the execution cost of many alternative access plans and choosing the one with the minimal estimated cost.

ordinary identifier. (1) In SQL, a letter followed by zero or more characters, each of which is a letter (a-z and A-Z), a symbol, a number, or the underscore character, used to form a name. (2) In DB2 UDB for OS/390, an *uppercase* letter followed by zero or more characters, each of which is an *uppercase* letter, a number, or the underscore character. An ordinary identifier must not be a reserved word.

ordinary token. A numeric constant, an ordinary identifier, a host identifier, or a keyword.

originating task. In DB2 UDB for OS/390, the primary agent in a parallel group that receives data from other execution units (referred to as *parallel tasks*) that are executing portions of the query in parallel.

outer join. (1) A join method in which a column that is not common to all of the tables being joined becomes part of the resultant table. Contrast with *inner join*. (2) In DB2 UDB for OS/390, 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*.

outline. In the OLAP Starter Kit, the structure that defines all elements of a database within the OLAP Starter Kit. For example, an outline contains definitions of dimensions, members, and formulas.

output file. A database or device file that is opened with the option to allow the writing of records.

overflow record. (1) On an indirectly addressed file, a record whose key is randomized to the address of a full track or to the address of a home record. (2) In DB2, an updated record that is too large to fit on the page it is currently stored in. The record is copied to a different page and its original location is replaced with a pointer to the new location. (3) In the Database Monitor, a record inserted in the event monitor data stream to indicate that records were discarded because the named pipe was full and records were not processed in time. An overflow record indicates how many records were discarded.

overloaded function name. A function name for which multiple functions exist within a function path or schema. Those within the same schema must have different signatures.

P

package. A control structure produced during program preparation that is used to execute SQL statements.

package list. In DB2 UDB for OS/390, an ordered list of package names that can be used to extend an application plan.

package name. In DB2 UDB for OS/390, 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.

packet. In data communication, a sequence of binary digits, including data and control signals, that is transmitted and switched as a composite whole.

page. (1) A block of storage within a table or index whose size is 4096 bytes (4 KB). (2) In DB2 UDB for OS/390, 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 an LOB table space, an LOB value can span more than one page, but no more than one LOB value is stored on a page.

page set. In an OS/390 environment, another way to refer to a table space or index space. Each page set consists of a collection of VSAM data sets.

page set recovery pending (PSRCP). In DB2 UDB for OS/390, a restrictive state of an index space in which the entire page set must be recovered. Recovery of a logical part is prohibited.

panel. In DB2 UDB for OS/390, a predefined display image that defines the locations and characteristics of display fields on a display surface (for example, a menu panel).

parallel group. In an OS/390 environment, a set of consecutive operations that execute in parallel and that have the same number of parallel tasks.

parallel I/O. The process of reading from or writing to two or more I/O devices at the same time to reduce response time.

parallel I/O processing. A form of I/O processing in which DB2 UDB for OS/390 initiates multiple concurrent requests for a single user query and performs I/O processing concurrently (in parallel) on multiple data partitions.

parallelism. The ability to perform multiple database operations at the same time (in parallel). See *inter-partition parallelism*, *intra-partition parallelism*, and *parallel I/O*.

parallel session. In SNA, two or more concurrently active sessions between the same two logical units. Each session can have different session parameters. See *session*.

Parallel Sysplex. A set of OS/390 systems that communicate and cooperate with each other through certain multisystem hardware components and software services.

parallel task. In an OS/390 environment, the execution unit that is dynamically created to process a query in parallel. It is implemented by an MVS service request block.

parameterized data type. A data type that can be defined with a specific length, scale, or precision. String and decimal data types are parameterized.

parameter marker. A question mark (?) that appears in a statement string of a dynamic SQL statement. The question mark can appear where a host variable might appear if the statement string was a static SQL statement.

parent key. A primary key or unique key that is used in a referential constraint. The values of a parent key determine the valid values of the foreign key in the constraint.

parent row. A row that has at least one dependent row.

parent table. A table that is a parent in at least one referential constraint.

parent table space. In DB2 UDB for OS/390, a table space that contains a parent table. A table space that contains a dependent of that table is a dependent table space.

Glossary

participant. In an OS/390 environment, an entity other than the commit coordinator that takes part in the commit process. Synonym for *agent* in SNA.

partition. In an OS/390 environment, 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.

partition compatible join. A join where all of the rows that are joined reside in the same database partition.

partitioned database. A database with two or more database partitions. Data in user tables can be located in one or more database partitions. When a table is on multiple partitions, some of its rows are stored in one partition and others are stored in other partitions. See *database partition*.

partitioned data set (PDS). In an OS/390 environment, 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. Synonym for *program library*.

partitioned page set. In an OS/390 environment, a partitioned table space or index space. Header pages, space map pages, data pages, and index pages refer to data only within the scope of the partition.

partitioned table space. In an OS/390 environment, a table space that is subdivided into parts (based on index key range), each of which can be processed independently by utilities.

partitioned function. A function that takes a partitioning key value of a row as input and produces a partition number as output.

partitioning key. (1) An ordered set of one or more columns in a given table. For each row in the table, the values in the partitioning key columns are used to determine on which database partition the row belongs. (2) In replication, an ordered set of one or more columns in a given table. For each row in the source table, the values in the partitioning key columns are used to determine in which target table the row belongs.

partitioning map. A vector of partition numbers that maps a partitioning map index to database partitions in the nodegroup.

partitioning map index. A number assigned to a hash partition or range partition.

partner logical unit (LU). (1) In SNA, the remote participant in a session. (2) An access point in the SNA network that is connected to the local DB2 UDB for OS/390 subsystem by way of a VTAM conversation.

pass-through. In a federated database system, a facility by which users can communicate with data sources in the SQL dialect of the data source.

path. See *SQL path*.

PCT. In CICS, program control table.

PDS. See *partitioned data set*.

peer-to-peer communication. Communication between two SNA logical units (LUs) that is not managed by a host; commonly used when referring to LU 6.2 nodes.

performance metrics. A collection of all performance variables belonging to the same database object.

Performance Monitor. A tool that lets database administrators use a graphical interface to monitor the performance of a DB2 system for tuning purposes. This tool can be accessed from the Control Center.

performance snapshot. Performance data for a set of database objects that is retrieved from the database manager at a point in time.

performance variable. A statistic derived from performance data obtained from the database manager. The expression for this variable can be user defined.

performance variable profile. A flat file that contains definitions of performance variables. This file can be edited, copied, and shared. Different profiles can be used by the same Performance Monitor so that different calculations can be performed.

persistence. In Net.Data, the state of keeping an assigned value for an entire transaction, where a transaction spans multiple Net.Data invocations. Only variables can be persistent. In addition, operations on resources affected by commitment control are kept active until an explicit commit or rollback is done, or when the transaction completes.

phantom row. A table row that can be read by application processes that are executing with any isolation level except repeatable read. When an application process issues the same query multiple times within a single unit of work, additional rows can appear between queries because of the data being inserted and committed by application processes that are running concurrently.

physical claim. In DB2 UDB for OS/390, a claim on an entire nonpartitioning index.

physical consistency. In DB2 UDB for OS/390, the state of a page that is not in a partially changed state.

physical drain. In DB2 UDB for OS/390, a drain on an entire nonpartitioning index.

physical lock (P-lock). A lock type that DB2 UDB for OS/390 acquires to provide consistency of data that is cached in different DB2 UDB for OS/390 subsystems. Physical locks are used only in data sharing environments. Contrast with *logical lock (L-lock)*.

physical lock contention. In DB2 UDB for OS/390, conflicting states of the requesters for a physical lock. See also *negotiable lock*.

physically complete. In DB2 UDB for OS/390, the state in which the concurrent copy process is completed and the output data set has been created.

physical unit (PU). The component that manages and monitors the resources (such as attached links and adjacent link stations) associated with a node, as requested by an SSCP through an SSCP-to-PU session. An SSCP activates a session with the PU in order to indirectly manage, through the PU, resources of the node such as attached links. This term applies to types 2.0, 4, and 5 nodes only. See also *control point*.

piece. In an OS/390 environment, a data set of a nonpartitioned page set.

plan. See *application plan*.

plan allocation. The process of allocating DB2 UDB for OS/390 resources to a plan in preparation to execute it.

plan name. In DB2 UDB for OS/390, the name of an application plan.

plan segmentation. In DB2 UDB for OS/390, the dividing of each plan into sections. When a section is needed, it is independently brought into the EDM pool.

P-lock. See *physical lock*.

PLT. In CICS, program list table.

point-in-time table. In DB2 replication, a type of target table whose content matches all or part of a source table, with an added system column that identifies the approximate time when the particular row was inserted or updated at the source system.

point of consistency. A point in time when all the recoverable data a program accesses is consistent. The point of consistency occurs when updates, inserts, and deletions are either committed to the physical database or rolled back. Synonym for *commit point* and *sync point*.

policy. See *CFRM policy*.

postponed abort UR. In DB2 UDB for OS/390, a unit of recovery that was inflight or in-abort, was interrupted by system failure or cancellation, and did not complete backout during restart.

PPT. (1) In CICS, processing program. (2) In OS/390, program properties table.

precision. In numeric data types, the total number of binary or decimal digits, excluding the sign.

precompile. To process programs that contain SQL statements before they are compiled. SQL statements are replaced with statements that will be recognized by the host language compiler. The output from a precompile process includes source code that can be submitted to the compiler and used in the bind process.

Glossary

predicate. An element of a search condition that expresses or implies a comparison operation.

prefetch. To read data ahead of, and in anticipation of, its use.

prepare. (1) To convert an SQL statement from text form to an executable form, by submitting it to the SQL compiler. (2) In DB2 UDB for OS/390, the first phase of a two-phase commit process in which all participants are requested to prepare for commit.

prepared SQL statement. In SQL, a named object that is the executable form of an SQL statement that has been processed by the PREPARE statement.

primary authorization ID. The authorization ID used to identify the application process to DB2 UDB for OS/390.

primary group buffer pool. For a duplexed group buffer pool, the DB2 UDB for OS/390 structure that is used to maintain the coherency of cached data. This structure is used for page registration and cross-invalidation. The OS/390 equivalent is *old* structure. Compare with *secondary group buffer pool*.

primary index. In DB2 UDB for OS/390, an index that enforces the uniqueness of a primary key.

primary key. A unique key that is part of the definition of a table. A primary key is the default parent key of a referential constraint definition.

primary log. A set of one or more log files used to record changes to a database. Storage for these files is allocated in advance. Contrast with *secondary log*.

principal. In an OS/390 environment, an entity that can communicate securely with another entity. In the DCE, principals are represented as entries in the DCE registry database and include users, servers, computers, and others.

principal name. In an OS/390 environment, the name by which a principal is known to the DCE security services.

private connection. A communications connection that is specific to DB2 UDB for OS/390.

private protocol access. A method of accessing distributed data by which you can direct a query to another DB2 system. Contrast with *DRDA access*.

private protocol connection. A DB2 private connection of the application process. See also *private connection*.

privilege. (1) The right to access a specific database object in a specific way. These rights are controlled by users with SYSADM (system administrator) authority or DBADM (database administrator) authority or by creators of objects. Privileges include rights such as creating, deleting, and selecting data from tables. (2) In DB2 UDB for OS/390, the capability of performing a specific function, sometimes on a specific object. See also *explicit privilege* and *implicit privilege*.

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 UDB for OS/390 catalog.

procedure. See *stored procedure*.

process. (1) In the Data Warehouse Center, a series of steps, which commonly operates on source data, that changes data from its original form into a form conducive to decision support. A Data Warehouse Center process commonly consists of one or more sources, one or more steps, and one or more targets. (2) In DB2 UDB for OS/390, the unit to which DB2 UDB for OS/390 allocates resources and locks. 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. Synonym for *application process*.

property. In the Data Warehouse Center, a characteristic or attribute that describes a unit of information. Each object type has a set of associated properties. For each object, a set of values is assigned to the properties.

protected conversation. In an OS/390 environment, a VTAM conversation that supports two-phase commit flows.

protocol.ini. A file that contains LAN configuration and binding information for all the protocol and medium-access control (MAC) system modules.

PSRCP. In DB2 UDB for OS/390, page set recovery pending.

PU. See *physical unit*.

public authority. The authority for an object granted to all users.

PU type. In SNA, the classification of a physical unit according to the type of node on which it resides.

Q

QSAM. Queued sequential access method.

quantified predicate. A predicate that compares a value with a set of values.

query. (1) A request for information from the database based on specific conditions, for example, a request for a list of all customers in a customer table whose balance is greater than $1000. (2) In DB2 UDB for OS/390, a component of certain SQL statements that specifies a result table.

query block. In DB2 UDB for OS/390, the part of a query that is represented by one of the FROM clauses. Each FROM clause can have multiple query blocks, depending on how DB2 UDB for OS/390 internally processes the query.

Query by Image Content (QBIC). A capability that is provided by the Image Extender that allows users to search images by their visual characteristics, such as average color and texture.

query CP parallelism. In DB2 UDB for OS/390, parallel execution of a single query, which is accomplished by using multiple tasks. Compare with *Sysplex query parallelism*.

query I/O parallelism. In DB2 UDB for OS/390, 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 is formed of input data blocks that are awaiting processing or of output data blocks that are awaiting transfer to auxiliary storage or to an output device.

quiesce. To end a process by allowing operations to complete normally, while rejecting any new requests for work.

quiesced member state. In DB2 UDB for OS/390, a state of a member of a data sharing group. An active member becomes quiesced when a STOP DB2 command takes effect without a failure. If the member task, address space, or OS/390 system fails before the command takes effect, the member state is failed.

quoted name. See *delimited identifier*.

R

RACF®. In an OS/390 environment, Resource Access Control Facility.

RAMAC®. In an OS/390 environment, the IBM family of enterprise disk storage system products.

RBA. See *relative byte address*.

RCT. In DB2 UDB for OS/390 with the CICS attachment facility, the resource control table.

RDB. See *relational database*.

RDBMS. See *relational database management system*.

RDBNAM. See *relational database name*.

RDF. In DB2 UDB for OS/390, record definition field.

read stability (RS). An isolation level that locks only the rows that an application retrieves within a transaction. Read stability ensures that any qualifying row that is read during a transaction is not changed by other application processes until the transaction is completed, and that any row changed by another application process is not read until the change is committed by that process. Read stability allows more concurrency than repeatable read, and less than cursor stability.

rebind. To create a package for an application program that was previously bound. For example, if an index is added for a table that is accessed by a program, the package must be rebound for it to take advantage of the new index.

record. The storage representation of a single row of a table or other data.

record identifier (RID). A number that is used internally by DB2 to uniquely identify a record in a table. The RID contains enough information to address the page in which the record is stored. Compare with *row ID*.

record identifier (RID) pool. In DB2 UDB for OS/390, an area of main storage above the 16-MB line that is reserved for sorting record identifiers during list prefetch processing.

recording. The information from performance snapshots that can be viewed at a later time.

recoverable log. A database log in which all log records are retained so that, in the event of a failure, lost data can be recovered during forward recovery. Contrast with *circular log*.

recovery. (1) The act of resetting a system, or data that is stored in a system, to an operable state following damage. (2) The process of rebuilding databases by restoring a backup and rolling forward the logs associated with it.

recovery log. See *database log*.

recovery pending. A state of the database or table space. A database or table space is put in recovery pending state when it is restored from a backup. While the database or table space is in this state, its data cannot be accessed.

recovery token. In DB2 UDB for OS/390, an identifier for an element that is used in recovery (for example, *NID* or *URID*).

RECP. In DB2 UDB for OS/390, recovery pending.

recursion cycle. The cycle that occurs when a fullselect within a common table expression includes the name of the common table expression in a FROM clause.

recursive common table expression. A common table expression that refers to itself in a FROM clause from the fullselect. Recursive common table expressions are used to write recursive queries.

recursive query. A fullselect that uses a recursive common table expression.

redo. In DB2 UDB for OS/390, a state of a unit of recovery that indicates that changes are to be reapplied to the DASD media to ensure data integrity.

referential constraint. The referential integrity rule that the nonnull values of the foreign key are valid only if they also appear as values of a parent key.

referential integrity. (1) The state of a database in which all values of all foreign keys are valid. (2) The condition that exists when all intended references from data in one column of a table to data in another column of the same or a different table are valid. Maintaining referential integrity requires that DB2 UDB for OS/390 enforce referential constraints on all LOAD, RECOVER, INSERT, UPDATE, and DELETE operations.

referential structure. In DB2 UDB for OS/390, 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.

refresh. A process in which all of the data of interest in a user table is copied to the target table, replacing existing data. See also *full refresh* and *differential refresh*.

registration. See *replication source*.

registration process. In DB2 replication, the process of defining a replication source. Contrast with *subscription process*.

registry database. In an OS/390 environment, a database of security information about principals, groups, organizations, accounts, and security policies. The DCE security component maintains the registry database.

regular table space. A table space that can store any nontemporary data.

rejected transaction. In DB2 replication, a transaction that contains one or more updates from replica tables that are out of date in comparison to the source table.

relational cube. A set of data and metadata that together define a multidimensional database. A relational cube is the portion of a multidimensional database that is stored in a relational database. See also *multidimensional database*.

relational database. A database that can be perceived as a set of tables and manipulated in accordance with the relational model of data.

relational database management system (RDBMS). In DB2 UDB for OS/390, a collection of hardware and software that organizes and provides access to a relational database.

relational database name (RDBNAM). A unique identifier for an RDBMS within a network. In DB2 UDB for OS/390, this must be the value in the LOCATION column of table SYSIBM.LOCATIONS in the CDB. DB2 UDB for OS/390 publications refer to the name of another RDBMS as a LOCATION value or a location name.

relationship. In DB2 UDB for OS/390, 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.

relative byte address (RBA). In an OS/390 environment, 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.

remigration. The process of returning to a current release of DB2 UDB for OS/390 following a fallback to a previous release. This procedure constitutes another migration process.

remote. In DB2 UDB for OS/390, any object that is maintained by a remote DB2 subsystem. A remote view, for example, is a view that is maintained by a remote DB2 subsystem. Contrast with *local*.

remote attach request. In DB2 UDB for OS/390, a request made by a remote location to attach to the local DB2 subsystem. Specifically, the request that is sent is an SNA Function Management Header 5.

remote database. A database that is physically located on a workstation other than the one in use. Contrast with *local database*.

remote subsystem. In DB2 UDB for OS/390, any RDBMS, except the *local subsystem*, 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 OS/390 system.

remote unit of work (RUOW). A unit of work that allows for the remote preparation and execution of SQL statements.

reoptimization. The DB2 UDB for OS/390 process of reconsidering the access path of an SQL statement at run time; during reoptimization, DB2 UDB for OS/390 uses the values of host variables, parameter markers, or special registers.

REORG pending (REORP). In DB2 UDB for OS/390, a condition that restricts SQL access and most utility access to an object that must be reorganized.

REORP. See *REORG pending*.

repeatable read (RR). An isolation level that locks all the rows in an application that are referenced within a transaction. When a program uses repeatable read protection, rows referenced by the program cannot be changed by other programs until the program ends the current transaction.

replica. A type of target table that can be updated locally and receives updates from a user table through a subscription definition. It can be a source for updating the user table or read-only target tables.

replica target table. A replication table at the target server that is a type of update-anywhere target table.

replication. The process of maintaining a defined set of data in more than one location. It involves copying designated changes for one location (a source) to another (a target), and synchronizing the data in both locations.

replication administrator. The user responsible for defining replication sources and subscriptions. This user can also run the Capture and Apply programs.

replication source. A database table or view that can accept copy requests and is the source table in a subscription set. See also *subscription set*.

replication subscription. A specification for copying changed data from replication sources to target tables at a specified time and frequency, with the option of enhancing data. It defines all of the information that is required by the Apply program to copy data.

request commit. In DB2 UDB for OS/390, the vote that is submitted to the prepare phase if the participant has modified data and is prepared to commit or roll back.

requester. In DB2 UDB for OS/390, the source of a request to a remote RDBMS, the system that requests the data. Synonym for *application requester*.

reserved word. (1) A word used in a source program to describe an action to be taken by the program or compiler. It must not appear in the program as a user-defined name or a system name. (2) A word that has been set aside for special use in the SQL standard.

resource. In DB2 UDB for OS/390, the object of a lock or claim, which could be a table space, an index space, a data partition, an index partition, or a logical partition.

resource allocation. In DB2 UDB for OS/390, the part of plan allocation that deals specifically with the database resources.

resource control table (RCT). In DB2 UDB for OS/390 with CICS, a construct of the CICS attachment facility, created by site-provided macro parameters, that defines authorization and access attributes for transactions or transaction groups.

resource definition online. In an OS/390 environment with CICS, a feature that you use to define CICS resources online without assembling tables.

resource limit facility (RLF). A portion of DB2 UDB for OS/390 code that prevents dynamic manipulative SQL statements from exceeding specified time limits. Synonym for *governor*.

resource limit specification table. In DB2 UDB for OS/390, a site-defined table that specifies the limits to be enforced by the resource limit facility.

restart pending (RESTP). In DB2 UDB for OS/390, 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 RECOVER POSTPONED command or the automatic online backout, which DB2 UDB for OS/390 invokes after restart if the system parameter LBACKOUT=AUTO.

RESTP. See *restart pending*.

restore. To return a backup copy to the active storage location for use.

restore set. A backup copy of a database or table space plus zero or more log files that, when restored and rolled forward, bring the database or table space back to a consistent state.

result set. The set of rows that a stored procedure returns.

result set locator. A 4-byte value that DB2 UDB for OS/390 uses to uniquely identify a query result set that a stored procedure returns.

result table. The set of rows produced by the evaluation of a SELECT statement.

retained lock. A MODIFY lock that a DB2 UDB for OS/390 subsystem was holding at the time of a subsystem failure. The lock is retained in the coupling facility lock structure across a DB2 UDB for OS/390 failure.

revoke. To remove a privilege or authority from an authorization ID.

RID. See *record identifier*.

RID pool. See *record identifier pool*.

right outer join. In DB2 UDB for OS/390, 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 *join*.

RLF. See *resource limit facility*.

RO. In DB2 UDB for OS/390, read-only access.

rollback. The process of restoring data changed by SQL statements to the state at its last commit point. See *point of consistency*.

roll-forward. The process of updating the data in a restored database by applying changes recorded in the database log. See *forward recovery*.

root page. In DB2 UDB for OS/390, 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.

routine. In DB2 UDB for OS/390, a user-defined function or a stored procedure.

row. The horizontal component of a table consisting of a sequence of values, one for each column of the table.

ROWID. See *row identifier*.

row identifier (ROWID). In DB2 UDB for OS/390, a value that uniquely identifies a row. This value is stored with the row and does not change.

row lock. In DB2 UDB for OS/390, a lock on a single row of data.

row-replica. In DB2 replication, a type of update-anywhere replica maintained by DataPropagator for Microsoft Jet without transaction semantics.

row-replica conflict detection. In DB2 replication, conflict detection that is performed row by row, not transaction by transaction, as is done for DB2 replicas.

row trigger. In DB2 UDB for OS/390, a trigger that is defined with the trigger granularity FOR EACH ROW.

RR. See *repeatable read*.

RRE. In an OS/390 environment with IMS, residual recovery entry.

RS. See *read stability*.

RRSAF. Recoverable Resource Manager Services attachment facility, which is a DB2 UDB for OS/390 subcomponent that uses OS/390 Transaction Management and Recoverable Resource Manager Services to coordinate resource commitment between DB2 UDB for OS/390 and all other resource managers that also use OS/390 RRS in an OS/390 system.

RUOW. See *remote unit of work*.

S

sargable. A predicate that can be evaluated as a search argument.

satellite. An occasionally connected client that has a DB2 server that synchronizes with its group at the satellite control database.

Satellite Administration Center. A user interface that provides centralized administrative support for satellites.

satellite control server. A DB2 Universal Database system that contains the satellite control database, SATCTLDB.

SBCS. See *single-byte character set*.

SCA. In DB2 UDB for OS/390, the shared communications area.

scalar fullselect. A fullselect that returns a single value—one row of data that consists of exactly one column.

scalar function. An SQL operation that produces a single value from another value and is expressed as a function name followed by a list of arguments enclosed in parentheses. Contrast with *column function*.

scale. The number of digits in the fractional part of a number.

schema. (1) A collection of database objects such as tables, views, indexes, or triggers. A database schema provides a logical classification of database objects. (2) In DB2 UDB for OS/390, 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. (3) In the Data Warehouse Center, a collection of warehouse target tables and the relationships between the warehouse target table columns, where the target tables can come from one or more warehouse targets.

SDK. See *Software Developer's Kit*.

SDWA. In an OS/390 environment, the system diagnostic work area.

search condition. A criterion for selecting rows from a table. A search condition consists of one or more predicates.

secondary authorization ID. In DB2 UDB for OS/390, an authorization ID that is associated with a primary authorization ID by an authorization exit routine.

secondary group buffer pool. For a duplexed group buffer pool in a DB2 UDB for OS/390 environment, the structure that is used to back up changed pages that are written to the primary group buffer pool. No page registration or cross-invalidation occurs using the secondary group buffer pool. The OS/390 equivalent is *new* structure. Compare to *primary group buffer pool*.

secondary log. A set of one or more log files used to record changes to a database. Storage for these files is allocated as needed when the primary log is full. Contrast with *primary log*.

section. In DB2 UDB for OS/390, 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.

segmented table space. In DB2 UDB for OS/390, 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.

self-referencing constraint. In DB2 UDB for OS/390, a referential constraint that defines a relationship in which a table is a dependent of itself.

self-referencing row. A row that is a parent of itself.

self-referencing subquery. A subselect or fullselect within a DELETE, INSERT, or UPDATE statement that refers to the same table that is the object of the SQL statement.

self-referencing table. A table that is both a parent and a dependent table in the same referential constraint.

sequential data set. A non-DB2 UDB for OS/390 data set whose records are organized on the basis of their successive physical positions, such as on magnetic tape. Several of the DB2 UDB for OS/390 database utilities require sequential data sets.

sequential prefetch. In DB2 UDB for OS/390, 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.

server. (1) In a network, a node that provides facilities to other stations, for example, a file server, a printer server, a mail server. (2) In a federated database system, a unit of information that identifies a data source to a federated server. This information can include the server's name, its type, its version, and the name of the wrapper that the federated server uses to communicate with and retrieve data from the data source. (3) A functional unit that provides services to one or more clients over a network. In the DB2 UDB for OS/390 environment, a server is the target for a request from a remote RDBMS and is the RDBMS that provides the data. See also *application server*.

service class. In DB2 UDB for OS/390, an 8-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.

service name. A name that provides a symbolic method of specifying the port number to be used at a remote node. The TCP/IP connection requires the address of the remote node and the port number to be used on the remote node to identify an application.

session. A logical connection between two stations or SNA network addressable units (NAUs) that allows the two stations or NAUs to communicate.

session limit. In SNA, the maximum number of concurrently active logical unit to logical unit (LU-to-LU) sessions that a particular logical unit (LU) can support.

session partner. In SNA, one of the two network addressable units (NAUs) participating in an active session.

session protocols. In DB2 UDB for OS/390, the available set of SNA communication requests and responses.

session security. For LU 6.2, partner LU verification and session data encryption. A Systems Network Architecture (SNA) function that allows data to be transmitted in encrypted form.

set operator. The SQL operators UNION, EXCEPT, and INTERSECT corresponding to the relational operators union, difference, and intersection. A set operator derives a result table by combining two other result tables.

shadowing. A recovery technique in which current page contents are never overwritten. Instead, new pages are allocated and written while the pages whose values are being replaced are retained as shadow copies until they are no longer needed to support the restoration of the system state due to a transaction rollback.

shared communications area (SCA). A coupling facility list structure that a DB2 UDB for OS/390 data sharing group uses for inter-DB2 communication.

shared lock. A lock that limits concurrently executing application processes to read-only operations on database data. Contrast with *exclusive lock*.

shift-in character. A special control character (X'0F') that is used in EBCDIC systems to denote that the subsequent bytes represent SBCS characters. Contrast with *shift-out character*.

shift-out character. A special control character (X'0E') that is used in EBCDIC systems to denote that the subsequent bytes, up to the next shift-in control character, represent DBCS characters. Contrast with *shift-in character*.

short string. (1) A fixed-length string or a variable-length string whose maximum length is less than or equal to 254 bytes. (2) In DB2 UDB for OS/390, a string whose actual length, or a variable-length string whose maximum length, is 255 bytes (or 127 double-byte characters) or less. Regardless of length, an LOB string is not a short string.

sign-on. A request that is made on behalf of an individual CICS or IMS application process by an attachment facility to enable DB2 UDB for OS/390 to verify that it is authorized to use DB2 UDB for OS/390 resources.

simple page set. In DB2 UDB for OS/390, a nonpartitioned page set. A simple page set initially consists of a single data set (page set piece). If 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 UDB for OS/390 considers the data sets to be a single contiguous linear address space that contains a maximum of 64 GB. Data is stored in the next available location within this address space without regard to any partitioning scheme.

simple table space. In DB2 UDB for OS/390, a table space that is neither partitioned nor segmented.

single-byte character set (SBCS). A character set in which each character is represented by a one-byte code.

single-precision floating point number. A 32-bit approximate representation of a real number.

SMF. In an OS/390 environment, system management facility.

SMS. In an OS/390 environment, Storage Management Subsystem.

SMS table space. See *system-managed space table space*.

Glossary

SNA. See *Systems Network Architecture*.

SNA network. The part of the user application network that conforms to the formats and protocols of Systems Network Architecture (SNA). It enables reliable transfer of data among users and provides protocols for controlling the resources of various network configurations. The SNA network consists of network addressable units (NAUs), gateway function, intermediate session routing function components, and the transport network.

snapshot. See *performance snapshot* and *explain snapshot*.

socket. A callable TCP/IP programming interface that is used by TCP/IP network applications to communicate with remote TCP/IP partners.

soft checkpoint. The process of writing some information to the log file header; this information is used to determine the starting point in the log in case a database restart is required.

Software Developer's Kit (SDK). An application development product that allows applications to be developed on a client workstation to access remote database servers including host relational databases through the DB2 Connect products.

source. In the Data Warehouse Center, a table, view, or file that is input to a step.

source function. A user-defined function (UDF) that is used to implement one or more other UDFs.

sourced function. In DB2 UDB for OS/390, 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 *external function* and *built-in function*.

source program. A set of host language statements and SQL statements that is processed by an SQL precompiler.

source server. In DB2 replication, the database location of the replication source and the Capture program.

source table. In DB2 replication, a table that contains the data that is to be copied to a target table. The source table can be a replication source table, a change data table, or a consistent-change-data table. Contrast with *target table*.

source type. An existing type that is used to internally represent a distinct type.

special register. A storage area that is defined for an application process by the database manager and is used to store information that can be referenced in SQL statements. Examples are USER and CURRENT DATE.

specific function name. (1) The name that uniquely identifies a function to the system. (2) In DB2 UDB for OS/390, a particular user-defined function that is known to the database manager by its specific name. Many specific user-defined functions can have the same function name. When a user-defined function is defined to the database, every function is assigned a specific name that is unique within its schema. Either the user can provide this name, or a default name is used.

spill file. In DB2 replication, a temporary file created by the Apply program that is used as the source for updating data to multiple target tables.

Spreadsheet Add-in. In the OLAP Starter Kit, software that merges with Microsoft Excel and Lotus 1-2-3 to allow multidimensional analysis of data. The software library appears as a menu add-in to the spreadsheet and provides such multidimensional analysis features as connect, zoom-in, and calculate.

SPUFI. In DB2 UDB for OS/390, SQL Processor Using File Input.

SQL. See *Structured Query Language*.

SQL authorization ID (SQL ID). In DB2 UDB for OS/390, the authorization ID that is used for checking dynamic SQL statements in some situations.

SQLCA. See *SQL communication area*.

SQL communication area (SQLCA). A set of variables that provides an application program with information about the execution of its SQL statements or its requests from the database manager.

SQL connection. In DB2 UDB for OS/390, an association between an application process and a local or remote application server.

SQLDA. See *SQL descriptor area*.

SQL descriptor area (SQLDA). (1) A set of variables that is used in the processing of certain SQL statements. The SQLDA is intended for dynamic SQL programs. (2) A structure that describes input variables, output variables, or the columns of a result table.

SQL escape character. In DB2 UDB for OS/390, the symbol that is used to enclose an SQL delimited identifier. This symbol is the double quotation mark ("). Compare to *escape character*.

SQL ID. See *SQL authorization ID*.

SQL path. In DB2 UDB for OS/390, an ordered list of schema names that are used in the resolution of unqualified references to user-defined functions, distinct types, and stored procedures. In dynamic SQL, the current path is found in the CURRENT PATH special register. In static SQL, it is defined in the PATH bind option.

SQL processing conversation. Any conversation that requires access of DB2 UDB for OS/390 data, either through an application or by dynamic query requests.

SQL Processor Using File Input (SPUFI). In DB2 UDB for OS/390, 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.

SQL return code. Either SQLCODE or SQLSTATE.

SQL routine. In DB2 UDB for OS/390, a user-defined function or stored procedure that is based on code that is written in SQL.

SQL string delimiter. In DB2 UDB for OS/390, a symbol that is used to enclose an SQL string constant. The SQL string delimiter is the apostrophe ('), except in COBOL applications, where the user assigns the symbol, which is either an apostrophe or a double quotation mark (").

SSCP. See *system services control point*.

SSI. In an OS/390 environment, subsystem interface.

SSM. In DB2 UDB for OS/390, subsystem member.

stack. An area in memory that stores temporary register information, parameters, and return addresses of subroutines.

staging table. In DB2 replication, a CCD table that can be used as the source for updating data to multiple target tables.

stand-alone. An attribute of a program that means it is capable of executing separately from DB2 UDB for OS/390, without using DB2 UDB for OS/390 services.

standard conflict detection. Conflict detection in which the Apply program searches for conflicts in rows that are already captured in the change data tables of the replica or user table. See also *conflict detection*, *enhanced conflict detection*, and *row-replica conflict detection*.

star schema. The type of relational database schema used by the OLAP Starter Kit, often created in the Data Warehouse Center.

statement. An instruction in a program or procedure.

statement handle. In CLI, a handle that refers to the data object that contains information about an SQL statement. This includes information such as dynamic arguments, bindings for dynamic arguments and columns, cursor information, result values, and status information. Each statement handle is associated with a connection handle.

statement string. For a dynamic SQL statement in a DB2 UDB for OS/390 environment, the character string form of the statement.

Glossary

static bind. A process by which SQL statements are bound after they are precompiled. All static SQL statements are prepared for execution at the same time. See also *bind*.

statement trigger. In DB2 UDB for OS/390, a trigger that is defined with the trigger granularity FOR EACH STATEMENT.

static SQL. SQL statements that are embedded within a program, and are prepared during the program preparation process before the program is executed. After being prepared, a static SQL statement does not change, although values of host variables specified by the statement can change.

status. In the Data Warehouse Center, the work-in-progress processing condition of a step, such as scheduled, populating, or successful.

step. In the Data Warehouse Center, a single operation on data in a warehouse process. In most cases, a step includes a warehouse source, a description of the transformation or movement of data, and a target. A step can be run according to a schedule, or it can cascade from another step.

step edition. In the Data Warehouse Center, a snapshot of the data in a warehouse source at a particular time.

storage group. A named set of DASD volumes on which DB2 UDB for OS/390 data can be stored.

stored procedure. (1) A block of procedural constructs and embedded SQL statements that is stored in a database and can be called by name. Stored procedures allow an application program to be run in two parts. One part runs on the client and the other on the server. This allows one call to produce several accesses to the database. Synonym for *procedure*. (2) In DB2 UDB for OS/390, a user-written application program that can be started through the use of the SQL CALL statement.

Stored Procedure Builder. A tool for creating stored procedures, building stored procedures on local and remote DB2 servers, modifying and rebuilding existing stored procedures, and testing and debugging the execution of installed stored procedures using a graphical interface. This tool is standalone or can be accessed from various integrated development environments.

Stored Procedure Builder project. A file that is created by the Stored Procedure Builder that contains connection information and stored procedure objects that have not been successfully built in the database.

storyboard. A visual summary of a video. The Video Extender includes features that can be used to identify and store video frames that are representative of the shots in a video. These representative frames can be used to build a storyboard.

string. In programming languages, the form of data used for storing and manipulating text.

strong typing. In DB2 UDB for OS/390, a process that guarantees that only user-defined functions and operations that are defined on a distinct type can be applied to that type. For example, you cannot directly compare two currency types, such as Canadian dollars and US dollars. But you can provide a user-defined function to convert one currency to the other and then do the comparison.

Structured Query Language (SQL). A standardized language for defining and manipulating data in a relational database.

subagent. A type of agent that works on subrequests. A single application can make many requests, and each request can be broken into many subrequests. Therefore, there can be multiple subagents working on behalf of the same application. All subagents working for the application are coordinated by the coordinating agent for that application.

subcomponent. A group of closely related DB2 UDB for OS/390 modules that work together to provide a general function.

subject area. In the Data Warehouse Center, a set of processes that create warehouse data for a particular logical business area. Processes in a subject area operate on data for a particular subject to create the detail data, data summaries, and cubes needed by that subject.

subordinate agent. See *subagent*.

subpage. In DB2 UDB for OS/390, the unit into which a physical index page can be divided.

subquery. A SELECT statement within the WHERE or HAVING clause of another SQL statement; a nested SQL statement.

subscription. See *subscription set*.

subscription cycle. In DB2 replication, a process in which the Apply program retrieves changed data for a given subscription set, replicates the changes to the target table, and updates the appropriate replication control tables to reflect the progress it made.

subscription process. In DB2 replication, a process in which you define subscription sets and subscription-set members. Contrast with *registration process*.

subscription set. In DB2 replication, the specification of a group of source tables, target tables, and the control information that governs the replication of changed data. See also *subscription-set member*.

subscription-set member. In DB2 replication, a member of a subscription set. There is one member for each source-target pair. Each member defines the structure of the target table and which rows and columns will be replicated from the source table.

subselect. That form of a query that does not include an ORDER BY clause, an UPDATE clause, or UNION operators.

substitution character. In SQL, a unique character that is substituted during character conversion for any characters in the source program that do not have a match in the target coding representation.

subsystem. In DB2 UDB for OS/390, a distinct instance of a relational database management system (RDBMS).

symbolic destination name. Specifies the name of a remote partner. The name corresponds to an entry in the CPI Communications side information table that contains the necessary information (partner LU name, mode name, partner TP name) for the client to set up an APPC connection to the server.

synchronization level. In APPC, the specification indicating whether the corresponding transaction programs exchange confirmation requests and replies.

synchronous. Pertaining to two or more processes that depend upon the occurrences of specific events, such as a common timing signal. Contrast with *asynchronous*.

sync point. See *point of consistency*.

synonym. In DB2 UDB for OS/390, an alternative name, in SQL, for a table or view. Synonyms can only be used to refer to objects at the subsystem in which the synonym is defined.

syntactic character set. A set of 81 graphic characters that are registered in the IBM registry as character set 00640. This set was originally recommended to the programming language community to be used for syntactic purposes toward maximizing portability and interchangeability across systems and country boundaries. It is contained in most of the primary registered character sets, with a few exceptions. Compare to *invariant character set*.

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 UDB for OS/390 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 UDB for OS/390 creates internally, such as prefetch processing, deferred writes, and service tasks.

system catalog. See *catalog*.

system conversation. The conversation that two DB2 UDB for OS/390 subsystems must establish to process system messages before any distributed processing can begin.

system database directory. A directory that contains entries for every database that can be accessed using the database manager. It is created when the first database is created or cataloged on the system.

system diagnostic work area (SDWA). In an OS/390 environment, the data that is recorded in a SYS1.LOGREC entry that describes a program or hardware error.

system-managed space (SMS) table space. A table space whose space is managed by the operating system. This storage model is based on files created under subdirectories, and managed by the file system. Contrast with *database managed space (DMS) table space*.

system services control point (SSCP). The control point in a SNA network that provides network services for dependent nodes.

Systems Network Architecture (SNA). The description of the logical structure, formats, protocols, and operational sequences for transmitting information units through the networks and also the operational sequences for controlling the configuration and operation of networks.

SYS1.DUMPxx data set. In an OS/390 environment, a data set that contains a system dump.

SYS1.LOGREC. In an OS/390 environment, a service aid that contains important information about program and hardware errors.

T

table. A named data object consisting of a specific number of columns and some unordered rows. See also *base table*.

table check constraint. In DB2 UDB for OS/390, a user-defined constraint that specifies the values that specific columns of a base table can contain.

table designator. A column name qualifier that designates a specific object table.

table function. In DB2 UDB for OS/390, a function that receives a set of arguments and returns a table to the SQL statement that refers to the function. A table function can be referenced only in the FROM clause of a subselect.

table locator. In DB2 UDB for OS/390, a mechanism that allows access to trigger transition tables in the FROM clause of SELECT statements, the subselect of INSERT statements, or from within user-defined functions. A table locator is a fullword integer value that represents a transition table.

table queue. A mechanism for transferring rows between database nodes. Table queues are distributed row streams with simplified rules for the insertion and removal of rows. Table queues can also be used to deliver rows between different processes in the serial database.

table space. (1) An abstraction of a collection of containers into which database objects are stored. A table space provides a level of indirection between a database and the tables stored within the database. A table space:

- Has space on media storage devices assigned to it.
- Has tables created within it. These tables use space in the containers that belong to the table space. The data, index, long field, and LOB portions of a table can be stored in the same table space, or can be individually broken out into separate table spaces.

(2) In DB2 UDB for OS/390, a page set that is used to store the records in one or more tables.

table space container. A generic term describing an allocation of space to a table space. Depending on the table space type, the container can be a directory, device, or file.

table space set. In DB2 UDB for OS/390, a set of table spaces and partitions that should be recovered together for one of these reasons:

- Each of them contains a table that is a parent or descendent of a table in one of the others.
- The set contains a base table and associated auxiliary tables.

A table space set can contain both types of relationships.

target. In the Data Warehouse Center, a table, view, or file that is produced or populated by a step; the output of a step.

target server. In DB2 replication, the database location of the target table. Normally this is also the location of the Apply program.

target table. In DB2 replication, the table on the target server to which data is copied. It can be a user copy table, a point-in-time table, a base aggregate table, a change aggregate table, a consistent-change-data table, or a replica table.

task control block (TCB). A control block that is used to communicate information about tasks within an address space that are connected to DB2 UDB for OS/390. An address space can support many task connections (as many as one per task), but only one address space connection.

TCB. See *task control block*.

TCP/IP. See *Transmission Control Protocol/Internet Protocol*.

TCP/IP port. A 2-byte value that identifies an end user or a TCP/IP network application within a TCP/IP host.

technical metadata. In the Data Warehouse Center, data that describes the technical aspects of the data, such as its database type and length. Technical metadata includes information about where the data comes from and the rules used to extract, clean, and transform the data. Much of the metadata in the Data Warehouse Center is technical metadata. Contrast with *business metadata*.

temporary table. A table created during the processing of an SQL statement to hold intermediate results. Contrast with *result table*.

temporary table space. A table space that can store only temporary tables.

territory. A portion of the POSIX locale that is mapped to the country code for internal processing by the database manager.

thread. (1) In some operating systems, the smallest unit of operation to be performed in a process. (2) The DB2 UDB for OS/390 structure that describes an application's connection, traces its progress, processes resource functions, and delimits its accessibility to DB2 UDB for OS/390 resources and services. Most DB2 UDB for OS/390 functions execute under a thread structure. Compare to *allied thread* and *database access thread*.

three-part name. The full name of a table, view, or alias. It consists of a location name, authorization ID, and an object name, separated by periods.

threshold trigger. An event that occurs when the value of a performance variable exceeds or falls below a user-defined threshold value. The action that occurs as a result of a threshold trigger can be:

- Logging information in an alert log file.
- Displaying information in an alert log window.
- Generating an audio alarm.
- Issuing a message window.
- Invoking a predefined command or program.

time. A three-part value that designates a time of day in hours, minutes, and seconds.

time duration. A DECIMAL(6,0) value that represents a number of hours, minutes, and seconds.

timeron. A unit of measurement used to give a rough relative estimate of the resources, or cost, required by the database server to execute two plans for the same query. The resources calculated in the estimate include weighted processor and I/O costs.

timeout. Abnormal termination of either the DB2 UDB for OS/390 subsystem or of an application because of the unavailability of resources. Installation specifications are set to determine both the amount of time DB2 UDB for OS/390 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.

timestamp. A seven-part value that consists of a date and time expressed in years, months, days, hours, minutes, seconds, and microseconds.

timestamp duration. A DECIMAL(20,6) value that represents a number of years, months, days, hours, minutes, seconds, and microseconds.

Glossary

Tivoli Storage Manager (TSM). A client/server product that provides storage management and data access services in a heterogeneous environment. TSM supports various communication methods, provides administrative facilities to manage the backup and storage of files, and provides facilities for scheduling backup operations.

TM Database. See *Transaction Manager Database*.

TMP. In an OS/390 environment, Terminal Monitor Program.

to-do. A state of a unit of recovery that indicates that the changes by the unit of recovery to recoverable DB2 UDB for OS/390 resources are indoubt and must be either applied to the DASD media or backed out, as determined by the commit coordinator.

token. The basic syntactic unit of a computing language. A token consists of one or more characters, excluding the blank character and excluding characters within a string constant or delimited identifier.

topology and routing services (TRS). An APPN control point component that manages the topology database and computes routes.

TP. See *transaction program*.

trace. A DB2 UDB for OS/390 facility that provides the ability to monitor and collect DB2 UDB for OS/390 monitoring, auditing, performance, accounting, statistics, and serviceability (global) data.

transaction. (1) An exchange between a workstation and a program, two workstations, or two programs that accomplish a particular action or result. An example is the entry of a customer's deposit and the update of the customer's balance. Synonym for *unit of work*. (2) One Net.Data invocation. If persistent Net.Data is used, then a transaction can span multiple Net.Data invocations.

transaction compensation. A process that restores rows that are affected by a committed transaction that is rejected. When a committed transaction is rejected, the rows are restored to the state that they were in before the transaction was committed.

transaction lock. In DB2 UDB for OS/390, a lock that is used to control concurrent execution of SQL statements.

transaction manager. A function that assigns identifiers to transactions, monitors their progress, and takes responsibility for transaction completion and failure recovery.

Transaction Manager Database (TM Database). A database that is used to log transactions when a two-phase commit (SYNCPOINT TWOPHASE) is used with DB2 databases. In the event of transaction failure, the TM Database information can be accessed to resynchronize databases involved in the failed transaction.

transaction program (TP). An application program that uses APPC to communicate with a partner application program.

transaction program name. In SNA LU 6.2 conversations, the name of the program at the remote logical unit that is to be the other half of the conversation.

transformation. In the Data Warehouse Center, an operation performed on data. Pivot and cleanse are types of transformations.

transformer. A program that operates on warehouse data. The Data Warehouse Center provides two types of transformers: statistical transformers, which provide statistics about the data in one or more tables; and warehouse transformers, which prepare the data for analysis. Each step has a type that corresponds to the transformer used in a process that performs types of data manipulation. For example, a clean step uses the Clean transformer.

transition table. A named temporary table that contains the transition values for each row affected by the triggering modification. An old transition table contains the values of affected rows before the modification is applied, and a new transition table contains the values of the affected rows after the modification is applied.

transition variable. A variable that is valid only in FOR EACH ROW triggers. It allows access to the transition values for the current row. An old transition variable is the value of the row before the modification is applied, and the new transition variable is the value of the row after the modification is applied.

Transmission Control Protocol/Internet Protocol (TCP/IP). A set of communications protocols that provide peer-to-peer connectivity functions for both local and wide area networks.

trigger. (1) In DB2, an object in a database that is invoked indirectly by the database manager when a particular SQL statement is run. (2) A set of SQL statements that are stored in a DB2 UDB for OS/390 database and executed when a certain event occurs in a DB2 UDB for OS/390 table.

trigger activation. In DB2 UDB for OS/390, the process that occurs when the trigger event that is defined in a trigger definition is executed. Trigger activation consists of the evaluation of the triggered action condition and conditional execution of the triggered SQL statements.

trigger activation time. In DB2 UDB for OS/390, an indication in a trigger definition of whether the trigger should be activated before or after the triggered event.

trigger body. In DB2 UDB for OS/390, the set of SQL statements that is executed when a trigger is activated and its triggered action condition evaluates to true.

trigger cascading. In DB2 UDB for OS/390, the process that occurs when the triggered action of a trigger causes the activation of another trigger.

triggered action. (1) The action that is executed when the trigger event occurs. (2) In DB2 UDB for OS/390, the SQL logic that is performed when a trigger is activated. The triggered action consists of an optional triggered action condition and a set of triggered SQL statements that are executed only if the condition evaluates to true.

triggered-action condition. (1) The search condition that controls the execution of the SQL statements within the triggered action. (2) In DB2 UDB for OS/390, an optional part of the triggered action. This Boolean condition appears as a WHEN clause and specifies a condition that DB2 evaluates to determine if the triggered SQL statements should be executed.

triggered SQL statements. In DB2 UDB for OS/390, the set of SQL statements that is executed when a trigger is activated and its triggered action condition evaluates to true. Triggered SQL statements are also called the *trigger body*.

trigger event. In a trigger definition, an update operation (INSERT, UPDATE, or DELETE statement) that causes the trigger to be run.

trigger granularity. In DB2 UDB for OS/390, a characteristic of a trigger, which determines whether the trigger is activated:

- Only once for the triggering SQL statement.
- Once for each row that the SQL statement modifies.

trigger package. In DB2 UDB for OS/390, a package that is created when a CREATE TRIGGER statement is executed. The package is executed when the trigger is activated.

triggering event. In DB2 UDB for OS/390, the specified operation in a trigger definition that causes the activation of that trigger. The triggering event is comprised of a triggering operation (INSERT, UPDATE, or DELETE) and a triggering table on which the operation is performed.

triggering SQL operation. In DB2 UDB for OS/390, the SQL operation that causes a trigger to be activated when performed on the triggering table.

triggering table. In DB2 UDB for OS/390, the table for which a trigger is created. When the defined triggering event occurs on this table, the trigger is activated.

truncation. The process of discarding part of a result from an operation when it exceeds memory or storage capacity.

TSO. In an OS/390 environment, Time-Sharing Option.

TSO attachment facility. A DB2 UDB for OS/390 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.

tuning parameters table. A table at the source server that contains timing information used by the Capture program. The information includes:

- How long to keep rows in the change data table.
- How much time can elapse before changes are stored in a database log or journal.

- How often to commit changed data to the unit of work tables.

two-phase commit. A two-step process by which recoverable resources and an external subsystem are committed. During the first step, the database manager subsystems are polled to ensure that they are ready to commit. If all subsystems respond positively, the database manager instructs them to commit.

typed parameter marker. A parameter marker that is specified along with its target data type. It has the general form:

```
CAST(? AS data-type)
```

type 1 indexes. Indexes that were created by a release of DB2 before DB2 for MVS/ESA Version 4 or that are specified as type 1 indexes in Version 4. Contrast with *type 2 indexes*. As of DB2 UDB for OS/390 Version 7, type 1 indexes are no longer supported.

type 2 indexes. Indexes that are created on a release of DB2 after DB2 for OS/390 Version 6 or that are specified as type 2 indexes in Version 4 or Version 6. Contrast with *type 1 indexes*.

U

UDF. See *user-defined function*.

UDT. See *user-defined type*.

unambiguous cursor. A cursor that allows a relational database to determine whether blocking can be used with the answer set. A cursor defined FOR FETCH ONLY or FOR READ ONLY can be used with blocking, whereas a cursor defined FOR UPDATE cannot.

unbind session (UNBIND). A request to deactivate a session between two logical units (LUs).

uncommitted read (UR). An isolation level that allows an application to access uncommitted changes of other transactions. The application does not lock other applications out of the row it is reading, unless the other application attempts to drop or alter the table.

uncoordinated transaction. A transaction that accesses more than one resource, but its commit or rollback is not being coordinated by a transaction manager.

underlying view. In DB2 UDB for OS/390, the view on which another view is directly or indirectly defined.

undo. A state of a unit of recovery that indicates that the changes that the unit of recovery made to recoverable DB2 UDB for OS/390 resources must be backed out.

Unicode. An international character encoding scheme that is a subset of the ISO 10646 standard. Each character supported is defined using a unique 2-byte code.

unique constraint. The rule that no two values in a primary key or key of a unique index can be the same. Also referred to as *uniqueness constraint*.

unique index. An index that ensures that no identical key values are stored in a table.

unique key. A key that is constrained so that no two of its values are equal.

unit of recovery. A recoverable sequence of operations within a single resource manager, such as an instance of DB2 UDB for OS/390. Contrast with *unit of work*.

unit of work. A recoverable sequence of operations within an application process. At any time, an application process is a single unit of work, but the life of an application process can involve many units of work as a result of commit or rollback operations. In a DB2 UDB for OS/390 *multi-site update* operation, a single unit of work can include several *units of recovery*. Synonym for *transaction*.

unit-of-work table. A replication control table at the source server that contains commit records read from the database log or journal. The records include a unit-of-recovery ID that can be used to join the unit-of-work table and the change data table to produce transaction-consistent change data. For DB2, the unit-of-work table optionally includes the correlation ID, which can be useful for auditing purposes.

unlock. The act of releasing an object or system resource that was previously locked and returning it to general availability within DB2 UDB for OS/390.

untyped parameter marker. A parameter marker that is specified without its target data type. It has the form of a single question mark.

update rule. A condition enforced by the database manager that must be met before a column can be updated.

update trigger. In DB2 UDB for OS/390, a trigger that is defined with the triggering SQL operation UPDATE.

upstream. In DB2 UDB for OS/390, the node in the syncpoint tree that is responsible, in addition to other recovery or resource managers, for coordinating the execution of a two-phase commit.

UR. See *uncommitted read*.

URE. In DB2 UDB for OS/390, unit of recovery element.

URID (unit of recovery ID). In DB2 UDB for OS/390, 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 copy table. In DB2 replication, a target table whose content matches all or part of a source table and contains only user data columns.

user-defined data type (UDT). See *distinct type*.

user-defined distinct type. See *distinct type*.

user-defined function (UDF). A function that is defined to the database management system and can be referred to in SQL queries. It can be one of the following functions:

- An external function, in which the body of the function is written in a programming language whose arguments are scalar values and a scalar result is produced for each invocation.
- A sourced function, implemented by another built-in or user-defined function already known to the DBMS. This function can be either a scalar function or column (aggregating) function, and returns a single value from a set of values (for example, MAX or AVG).

user-defined performance variable. A performance variable created by a user and added to the performance variable profile.

user-defined program. A program that a user supplies and defines to the Data Warehouse Center, as contrasted with supplied programs, which are included with and defined automatically in the Data Warehouse Center.

user-defined type (UDT). A data type that is not native to the database manager and was created by a user. In DB2 UDB for OS/390, the term *distinct type* is used instead of user-defined type.

user mapping. An association between the authorization under which a user connects to a federated server and the authorization under which the user connects to a data source.

user table. In DB2 replication, a table created for and used by an application before it is defined as a replication source. It is used as the source for updates to read-only target tables, consistent-change-data tables, replicas, and row-replica tables.

UT. In DB2 UDB for OS/390, utility-only access.

UTC. See *Coordinated Universal Time*.

V

value. (1) The smallest unit of data manipulated in SQL. (2) A specific data item at the intersection of a column and a row.

variable. A data element that specifies a value that can be changed.

variant function. A user-defined function whose result is dependent on its input parameter values as well as other factors. Successive invocations with the same parameter values might produce different results. Contrast with *not-variant function*.

varying-length string. A character, graphic, or binary string whose length is not fixed but can range within set limits. Also referred to as a *variable-length string*.

version. In DB2 UDB for OS/390, a member of a set of similar programs, DBRMs, packages, or LOBs.

- A version of a program 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).
- A version of a DBRM 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.
- A version of a package 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.
- A version of a LOB 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.

view. A logical table that consists of data that is generated by a query. Contrast with *base table*.

view check option. In DB2 UDB for OS/390, an option that specifies whether every row that is inserted or updated through a view must conform to the definition of that view. A view check option can be specified with the WITH CASCADED CHECK OPTION, WITH CHECK OPTION, or WITH LOCAL CHECK OPTION clauses of the CREATE VIEW statement.

Virtual Storage Access Method (VSAM). An access method for direct or sequential processing of fixed-length 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.

Virtual Telecommunications Access Method (VTAM). In an OS/390 environment, an IBM licensed program that controls communication and the flow of data in an SNA network.

Visual Explain. A tool that lets database administrators and application programmers use a graphical interface to display and analyze detailed information on the access plan of a given SQL statement. The tasks provided by this tool can be accessed from the Control Center.

VSAM. See *Virtual Storage Access Method*.

VTAM. See *Virtual Telecommunication Access Method*.

W

warehouse. A subject-oriented nonvolatile collection of data used to support strategic decision making. The warehouse is the central point of data integration for business intelligence. It is the source of data for datamarts within an enterprise and delivers a common view of enterprise data.

warehouse agent. In the Data Warehouse Center, a run-time process that manages data movement and transformation.

warehouse control database. The Data Warehouse Center database that contains the control tables that are required to store Data Warehouse Center metadata.

warehouse program group. In the Data Warehouse Center, a container (folder) that holds program objects.

warehouse source. A subset of tables and views from a single database, or a set of files, that have been defined to the Data Warehouse Center.

warehouse target. A subset of tables, indexes, and aliases from a single database that are managed by the Data Warehouse Center.

warm start. (1) A restart that allows reuse of previously initialized input and output work queues. Contrast with *cold start*. (2) In DB2 replication, a start of the Capture program that allows reuse of previously initialized input and output work queues.

well known address. An address used to uniquely identify a particular node in the network to establish connections between nodes. The well known address is a combination of the network address and the port used on the logical node.

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 UDB for OS/390 stored procedure runs.

work file. In DB2 replication, a temporary file used by the Apply program when processing a subscription set.

wrapper. In a federated database system, the mechanism by which the federated server invokes routines to communicate with, and retrieve data from, a data source. The routines are contained in a library called a *wrapper module*.

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. See *write to operator*.

WTOR. A write to operator (WTO) with reply.

X

XCF. See *cross-system coupling facility*.

XID. Exchange station ID.

XRF. See *extended recovery facility*.

Appendix A. Using the DB2 Library

The DB2 Universal Database library consists of online help, books (PDF and HTML), and sample programs in HTML format. This section describes the information that is provided, and how you can access it.

To access product information online, you can use the Information Center. For more information, see “Accessing Information with the Information Center” on page 72. You can view task information, DB2 books, troubleshooting information, sample programs, and DB2 information on the Web.

DB2 PDF Files and Printed Books

DB2 Information

The following table divides the DB2 books into four categories:

DB2 Guide and Reference Information
These books contain the common DB2 information for all platforms.

DB2 Installation and Configuration Information
These books are for DB2 on a specific platform. For example, there are separate *Quick Beginnings* books for DB2 on OS/2, Windows, and UNIX-based platforms.

Cross-platform sample programs in HTML
These samples are the HTML version of the sample programs that are installed with the Application Development Client. The samples are for informational purposes and do not replace the actual programs.

Release notes
These files contain late-breaking information that could not be included in the DB2 books.

The installation manuals, release notes, and tutorials are viewable in HTML directly from the product CD-ROM. Most books are available in HTML on the product CD-ROM for viewing and in Adobe Acrobat (PDF) format on the DB2 publications CD-ROM for viewing and printing. You can also order a printed copy from IBM; see “Ordering the Printed Books” on page 69. The following table lists books that can be ordered.

On OS/2 and Windows platforms, you can install the HTML files under the `sqllib\doc\html` directory. DB2 information is translated into different languages; however, all the information is not translated into every language. Whenever information is not available in a specific language, the English information is provided

On UNIX platforms, you can install multiple language versions of the HTML files under the `doc/%L/html` directories, where *%L* represents the locale. For more information, refer to the appropriate *Quick Beginnings* book.

You can obtain DB2 books and access information in a variety of ways:
- “Viewing Information Online” on page 71
- “Searching Information Online” on page 74
- “Ordering the Printed Books” on page 69
- “Printing the PDF Books” on page 68

Table 1. DB2 Information

Name	**Description**	**Form Number** **PDF File Name**	**HTML Directory**
	DB2 Guide and Reference Information		
Administration Guide	*Administration Guide: Planning* provides an overview of database concepts, information about design issues (such as logical and physical database design), and a discussion of high availability.	SC09-2946 db2d1x70	db2d0
	Administration Guide: Implementation provides information on implementation issues such as implementing your design, accessing databases, auditing, backup and recovery.	SC09-2944 db2d2x70	
	Administration Guide: Performance provides information on database environment and application performance evaluation and tuning.	SC09-2945 db2d3x70	
	You can order the three volumes of the *Administration Guide* in the English language in North America using the form number SBOF-8934.		
Administrative API Reference	Describes the DB2 application programming interfaces (APIs) and data structures that you can use to manage your databases. This book also explains how to call APIs from your applications.	SC09-2947 db2b0x70	db2b0
Application Building Guide	Provides environment setup information and step-by-step instructions about how to compile, link, and run DB2 applications on Windows, OS/2, and UNIX-based platforms.	SC09-2948 db2axx70	db2ax
APPC, CPI-C, and SNA Sense Codes	Provides general information about APPC, CPI-C, and SNA sense codes that you may encounter when using DB2 Universal Database products. Available in HTML format only.	No form number db2apx70	db2ap

Table 1. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
Application Development Guide	Explains how to develop applications that access DB2 databases using embedded SQL or Java (JDBC and SQLJ). Discussion topics include writing stored procedures, writing user-defined functions, creating user-defined types, using triggers, and developing applications in partitioned environments or with federated systems.	SC09-2949 db2a0x70	db2a0
CLI Guide and Reference	Explains how to develop applications that access DB2 databases using the DB2 Call Level Interface, a callable SQL interface that is compatible with the Microsoft ODBC specification.	SC09-2950 db2l0x70	db2l0
Command Reference	Explains how to use the Command Line Processor and describes the DB2 commands that you can use to manage your database.	SC09-2951 db2n0x70	db2n0
Connectivity Supplement	Provides setup and reference information on how to use DB2 for AS/400, DB2 for OS/390, DB2 for MVS, or DB2 for VM as DRDA application requesters with DB2 Universal Database servers. This book also details how to use DRDA application servers with DB2 Connect application requesters. Available in HTML and PDF only.	No form number db2h1x70	db2h1
Data Movement Utilities Guide and Reference	Explains how to use DB2 utilities, such as import, export, load, AutoLoader, and DPROP, that facilitate the movement of data.	SC09-2955 db2dmx70	db2dm
Data Warehouse Center Administration Guide	Provides information on how to build and maintain a data warehouse using the Data Warehouse Center.	SC26-9993 db2ddx70	db2dd
Data Warehouse Center Application Integration Guide	Provides information to help programmers integrate applications with the Data Warehouse Center and with the Information Catalog Manager.	SC26-9994 db2adx70	db2ad
DB2 Connect User's Guide	Provides concepts, programming, and general usage information for the DB2 Connect products.	SC09-2954 db2c0x70	db2c0
DB2 Query Patroller Administration Guide	Provides an operational overview of the DB2 Query Patroller system, specific operational and administrative information, and task information for the administrative graphical user interface utilities.	SC09-2958 db2dwx70	db2dw
DB2 Query Patroller User's Guide	Describes how to use the tools and functions of the DB2 Query Patroller.	SC09-2960 db2wwx70	db2ww
Glossary	Provides definitions for terms used in DB2 and its components. Available in HTML format and in the *SQL Reference*.	No form number db2t0x70	db2t0

Table 1. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
Image, Audio, and Video Extenders Administration and Programming	Provides general information about DB2 extenders, and information on the administration and configuration of the image, audio, and video (IAV) extenders and on programming using the IAV extenders. It includes reference information, diagnostic information (with messages), and samples.	SC26-9929 dmbu7x70	dmbu7
Information Catalog Manager Administration Guide	Provides guidance on managing information catalogs.	SC26-9995 db2dix70	db2di
Information Catalog Manager Programming Guide and Reference	Provides definitions for the architected interfaces for the Information Catalog Manager.	SC26-9997 db2bix70	db2bi
Information Catalog Manager User's Guide	Provides information on using the Information Catalog Manager user interface.	SC26-9996 db2aix70	db2ai
Installation and Configuration Supplement	Guides you through the planning, installation, and setup of platform-specific DB2 clients. This supplement also contains information on binding, setting up client and server communications, DB2 GUI tools, DRDA AS, distributed installation, the configuration of distributed requests, and accessing heterogeneous data sources.	GC09-2957 db2iyx70	db2iy
Message Reference	Lists messages and codes issued by DB2, the Information Catalog Manager, and the Data Warehouse Center, and describes the actions you should take. You can order both volumes of the Message Reference in the English language in North America with the form number SBOF-8932.	Volume 1 GC09-2978 db2m1x70 Volume 2 GC09-2979 db2m2x70	db2m0
OLAP Integration Server Administration Guide	Explains how to use the Administration Manager component of the OLAP Integration Server.	SC27-0787 db2dpx70	n/a
OLAP Integration Server Metaoutline User's Guide	Explains how to create and populate OLAP metaoutlines using the standard OLAP Metaoutline interface (not by using the Metaoutline Assistant).	SC27-0784 db2upx70	n/a
OLAP Integration Server Model User's Guide	Explains how to create OLAP models using the standard OLAP Model Interface (not by using the Model Assistant).	SC27-0783 db2lpx70	n/a
OLAP Setup and User's Guide	Provides configuration and setup information for the OLAP Starter Kit.	SC27-0702 db2ipx70	db2ip
OLAP Spreadsheet Add-in User's Guide for Excel	Describes how to use the Excel spreadsheet program to analyze OLAP data.	SC27-0786 db2epx70	db2ep
OLAP Spreadsheet Add-in User's Guide for Lotus 1-2-3	Describes how to use the Lotus 1-2-3 spreadsheet program to analyze OLAP data.	SC27-0785 db2tpx70	db2tp

Table 1. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
Replication Guide and Reference	Provides planning, configuration, administration, and usage information for the IBM Replication tools supplied with DB2.	SC26-9920 db2e0x70	db2e0
Spatial Extender User's Guide and Reference	Provides information about installing, configuring, administering, programming, and troubleshooting the Spatial Extender. Also provides significant descriptions of spatial data concepts and provides reference information (messages and SQL) specific to the Spatial Extender.	SC27-0701 db2sbx70	db2sb
SQL Getting Started	Introduces SQL concepts and provides examples for many constructs and tasks.	SC09-2973 db2y0x70	db2y0
SQL Reference, Volume 1 and Volume 2	Describes SQL syntax, semantics, and the rules of the language. This book also includes information about release-to-release incompatibilities, product limits, and catalog views. You can order both volumes of the *SQL Reference* in the English language in North America with the form number SBOF-8933.	Volume 1 SC09-2974 db2s1x70 Volume 2 SC09-2975 db2s2x70	db2s0
System Monitor Guide and Reference	Describes how to collect different kinds of information about databases and the database manager. This book explains how to use the information to understand database activity, improve performance, and determine the cause of problems.	SC09-2956 db2f0x70	db2f0
Text Extender Administration and Programming	Provides general information about DB2 extenders and information on the administration and configuring of the text extender and on programming using the text extenders. It includes reference information, diagnostic information (with messages) and samples.	SC26-9930 desu9x70	desu9
Troubleshooting Guide	Helps you determine the source of errors, recover from problems, and use diagnostic tools in consultation with DB2 Customer Service.	GC09-2850 db2p0x70	db2p0
What's New	Describes the new features, functions, and enhancements in DB2 Universal Database, Version 7.	SC09-2976 db2q0x70	db2q0
	DB2 Installation and Configuration Information		
DB2 Connect Enterprise Edition for OS/2 and Windows Quick Beginnings	Provides planning, migration, installation, and configuration information for DB2 Connect Enterprise Edition on the OS/2 and Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2953 db2c6x70	db2c6

Table 1. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
DB2 Connect Enterprise Edition for UNIX Quick Beginnings	Provides planning, migration, installation, configuration, and task information for DB2 Connect Enterprise Edition on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2952 db2cyx70	db2cy
DB2 Connect Personal Edition Quick Beginnings	Provides planning, migration, installation, configuration, and task information for DB2 Connect Personal Edition on the OS/2 and Windows 32-bit operating systems. This book also contains installation and setup information for all supported clients.	GC09-2967 db2c1x70	db2c1
DB2 Connect Personal Edition Quick Beginnings for Linux	Provides planning, installation, migration, and configuration information for DB2 Connect Personal Edition on all supported Linux distributions.	GC09-2962 db2c4x70	db2c4
DB2 Data Links Manager Quick Beginnings	Provides planning, installation, configuration, and task information for DB2 Data Links Manager for AIX and Windows 32-bit operating systems.	GC09-2966 db2z6x70	db2z6
DB2 Enterprise - Extended Edition for UNIX Quick Beginnings	Provides planning, installation, and configuration information for DB2 Enterprise - Extended Edition on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2964 db2v3x70	db2v3
DB2 Enterprise - Extended Edition for Windows Quick Beginnings	Provides planning, installation, and configuration information for DB2 Enterprise - Extended Edition for Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2963 db2v6x70	db2v6
DB2 for OS/2 Quick Beginnings	Provides planning, installation, migration, and configuration information for DB2 Universal Database on the OS/2 operating system. This book also contains installation and setup information for many supported clients.	GC09-2968 db2i2x70	db2i2
DB2 for UNIX Quick Beginnings	Provides planning, installation, migration, and configuration information for DB2 Universal Database on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2970 db2ixx70	db2ix
DB2 for Windows Quick Beginnings	Provides planning, installation, migration, and configuration information for DB2 Universal Database on Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2971 db2i6x70	db2i6
DB2 Personal Edition Quick Beginnings	Provides planning, installation, migration, and configuration information for DB2 Universal Database Personal Edition on the OS/2 and Windows 32-bit operating systems.	GC09-2969 db2i1x70	db2i1

Table 1. DB2 Information (continued)

Name	Description	Form Number	HTML Directory
		PDF File Name	
DB2 Personal Edition Quick Beginnings for Linux	Provides planning, installation, migration, and configuration information for DB2 Universal Database Personal Edition on all supported Linux distributions.	GC09-2972 db2i4x70	db2i4
DB2 Query Patroller Installation Guide	Provides installation information about DB2 Query Patroller.	GC09-2959 db2iwx70	db2iw
DB2 Warehouse Manager Installation Guide	Provides installation information for warehouse agents, warehouse transformers, and the Information Catalog Manager.	GC26-9998 db2idx70	db2id
Cross-Platform Sample Programs in HTML			
Sample programs in HTML	Provides the sample programs in HTML format for the programming languages on all platforms supported by DB2. The sample programs are provided for informational purposes only. Not all samples are available in all programming languages. The HTML samples are only available when the DB2 Application Development Client is installed. For more information on the programs, refer to the *Application Building Guide*.	No form number	db2hs
Release Notes			
DB2 Connect Release Notes	Provides late-breaking information that could not be included in the DB2 Connect books.	See note #2.	db2cr
DB2 Installation Notes	Provides late-breaking installation-specific information that could not be included in the DB2 books.	Available on product CD-ROM only.	
DB2 Release Notes	Provides late-breaking information about all DB2 products and features that could not be included in the DB2 books.	See note #2.	db2ir

Notes:

1. The character *x* in the sixth position of the file name indicates the language version of a book. For example, the file name `db2d0e70` identifies the English version of the *Administration Guide* and the file name `db2d0f70` identifies the French version of the same book. The following letters are used in the sixth position of the file name to indicate the language version:

Language	Identifier
Brazilian Portuguese	b
Bulgarian	u
Czech	x
Danish	d
Dutch	q
English	e
Finnish	y
French	f
German	g
Greek	a

Hungarian	h
Italian	i
Japanese	j
Korean	k
Norwegian	n
Polish	p
Portuguese	v
Russian	r
Simp. Chinese	c
Slovenian	l
Spanish	z
Swedish	s
Trad. Chinese	t
Turkish	m

2. Late breaking information that could not be included in the DB2 books is available in the Release Notes in HTML format and as an ASCII file. The HTML version is available from the Information Center and on the product CD-ROMs. To view the ASCII file:
   - On UNIX-based platforms, see the `Release.Notes` file. This file is located in the `DB2DIR/Readme/`*`%L`* directory, where *`%L`* represents the locale name and `DB2DIR` represents:
     - `/usr/lpp/db2_07_01` on AIX
     - `/opt/IBMdb2/V7.1` on HP-UX, PTX, Solaris, and Silicon Graphics IRIX
     - `/usr/IBMdb2/V7.1` on Linux.
   - On other platforms, see the `RELEASE.TXT` file. This file is located in the directory where the product is installed. On OS/2 platforms, you can also double-click the **IBM DB2** folder and then double-click the **Release Notes** icon.

Printing the PDF Books

If you prefer to have printed copies of the books, you can print the PDF files found on the DB2 publications CD-ROM. Using the Adobe Acrobat Reader, you can print either the entire book or a specific range of pages. For the file name of each book in the library, see Table 1 on page 62.

You can obtain the latest version of the Adobe Acrobat Reader from the Adobe Web site at http://www.adobe.com.

The PDF files are included on the DB2 publications CD-ROM with a file extension of PDF. To access the PDF files:

1. Insert the DB2 publications CD-ROM. On UNIX-based platforms, mount the DB2 publications CD-ROM. Refer to your *Quick Beginnings* book for the mounting procedures.
2. Start the Acrobat Reader.
3. Open the desired PDF file from one of the following locations:
   - On OS/2 and Windows platforms:

     `x``:\doc\`*`language`* directory, where *x* represents the CD-ROM drive and *language* represent the two-character country code that represents your language (for example, EN for English).
   - On UNIX-based platforms:

`/cdrom/doc/%L` directory on the CD-ROM, where */cdrom* represents the mount point of the CD-ROM and *%L* represents the name of the desired locale.

You can also copy the PDF files from the CD-ROM to a local or network drive and read them from there.

Ordering the Printed Books

You can order the printed DB2 books either individually or as a set (in North America only) by using a sold bill of forms (SBOF) number. To order books, contact your IBM authorized dealer or marketing representative, or phone `1-800-879-2755` in the United States or `1-800-IBM-4YOU` in Canada. You can also order the books from the Publications Web page at http://www.elink.ibmlink.ibm.com/pbl/pbl.

Two sets of books are available. SBOF-8935 provides reference and usage information for the DB2 Warehouse Manager. SBOF-8931 provides reference and usage information for all other DB2 Universal Database products and features. The contents of each SBOF are listed in the following table:

Table 2. Ordering the printed books

SBOF Number	Books Included	
SBOF-8931	• Administration Guide: Planning • Administration Guide: Implementation • Administration Guide: Performance • Administrative API Reference • Application Building Guide • Application Development Guide • CLI Guide and Reference • Command Reference • Data Movement Utilities Guide and Reference • Data Warehouse Center Administration Guide • Data Warehouse Center Application Integration Guide • DB2 Connect User's Guide • Installation and Configuration Supplement • Image, Audio, and Video Extenders Administration and Programming • Message Reference, Volumes 1 and 2	• OLAP Integration Server Administration Guide • OLAP Integration Server Metaoutline User's Guide • OLAP Integration Server Model User's Guide • OLAP Integration Server User's Guide • OLAP Setup and User's Guide • OLAP Spreadsheet Add-in User's Guide for Excel • OLAP Spreadsheet Add-in User's Guide for Lotus 1-2-3 • Replication Guide and Reference • Spatial Extender Administration and Programming Guide • SQL Getting Started • SQL Reference, Volumes 1 and 2 • System Monitor Guide and Reference • Text Extender Administration and Programming • Troubleshooting Guide • What's New
SBOF-8935	• Information Catalog Manager Administration Guide • Information Catalog Manager User's Guide • Information Catalog Manager Programming Guide and Reference	• Query Patroller Administration Guide • Query Patroller User's Guide

DB2 Online Documentation

Accessing Online Help

Online help is available with all DB2 components. The following table describes the various types of help.

Type of Help	Contents	How to Access...
Command Help	Explains the syntax of commands in the command line processor.	From the command line processor in interactive mode, enter: `? command` where *command* represents a keyword or the entire command. For example, `? catalog` displays help for all the CATALOG commands, while `? catalog database` displays help for the CATALOG DATABASE command.
Client Configuration Assistant Help *Command Center Help* *Control Center Help* *Data Warehouse Center Help* *Event Analyzer Help* *Information Catalog Manager Help* *Satellite Administration Center Help* *Script Center Help*	Explains the tasks you can perform in a window or notebook. The help includes overview and prerequisite information you need to know, and it describes how to use the window or notebook controls.	From a window or notebook, click the **Help** push button or press the **F1** key.

Message Help	Describes the cause of a message and any action you should take.	From the command line processor in interactive mode, enter:
		`? XXXnnnnn`
		where *XXXnnnnn* represents a valid message identifier.
		For example, `? SQL30081` displays help about the SQL30081 message.
		To view message help one screen at a time, enter:
		`? XXXnnnnn | more`
		To save message help in a file, enter:
		`? XXXnnnnn > filename.ext`
		where *filename.ext* represents the file where you want to save the message help.

Type of Help	Contents	How to Access...
SQL Help	Explains the syntax of SQL statements.	From the command line processor in interactive mode, enter: `help statement` where *statement* represents an SQL statement. For example, `help SELECT` displays help about the SELECT statement. **Note:** SQL help is not available on UNIX-based platforms.
SQLSTATE Help	Explains SQL states and class codes.	From the command line processor in interactive mode, enter: `? sqlstate or ? class code` where *sqlstate* represents a valid five-digit SQL state and *class code* represents the first two digits of the SQL state. For example, `? 08003` displays help for the 08003 SQL state, while `? 08` displays help for the 08 class code.

Viewing Information Online

The books included with this product are in Hypertext Markup Language (HTML) softcopy format. Softcopy format enables you to search or browse the information and provides hypertext links to related information. It also makes it easier to share the library across your site.

You can view the online books or sample programs with any browser that conforms to HTML Version 3.2 specifications.

To view online books or sample programs:

- If you are running DB2 administration tools, use the Information Center.
- From a browser, click **File —>Open Page**. The page you open contains descriptions of and links to DB2 information:
  - On UNIX-based platforms, open the following page:

    ```
    INSTHOME/sqllib/doc/%L/html/index.htm
    ```

    where *%L* represents the locale name.
  - On other platforms, open the following page:

    ```
    sqllib\doc\html\index.htm
    ```

    The path is located on the drive where DB2 is installed.

    If you have not installed the Information Center, you can open the page by double-clicking the **DB2 Information** icon. Depending on the system you are using, the icon is in the main product folder or the Windows Start menu.

Installing the Netscape Browser

If you do not already have a Web browser installed, you can install Netscape from the Netscape CD-ROM found in the product boxes. For detailed instructions on how to install it, perform the following:

1. Insert the Netscape CD-ROM.
2. On UNIX-based platforms only, mount the CD-ROM. Refer to your *Quick Beginnings* book for the mounting procedures.

3. For installation instructions, refer to the `CDNAV`*`nn`*`.txt` file, where *nn* represents your two character language identifier. The file is located at the root directory of the CD-ROM.

Accessing Information with the Information Center

The Information Center provides quick access to DB2 product information. The Information Center is available on all platforms on which the DB2 administration tools are available.

You can open the Information Center by double-clicking the Information Center icon. Depending on the system you are using, the icon is in the Information folder in the main product folder or the Windows **Start** menu.

You can also access the Information Center by using the toolbar and the **Help** menu on the DB2 Windows platform.

The Information Center provides six types of information. Click the appropriate tab to look at the topics provided for that type.

Tasks
Key tasks you can perform using DB2.

Reference
DB2 reference information, such as keywords, commands, and APIs.

Books
DB2 books.

Troubleshooting
Categories of error messages and their recovery actions.

Sample Programs
Sample programs that come with the DB2 Application Development Client. If you did not install the DB2 Application Development Client, this tab is not displayed.

Web
DB2 information on the World Wide Web. To access this information, you must have a connection to the Web from your system.

When you select an item in one of the lists, the Information Center launches a viewer to display the information. The viewer might be the system help viewer, an editor, or a Web browser, depending on the kind of information you select.

The Information Center provides a find feature, so you can look for a specific topic without browsing the lists.

For a full text search, follow the hypertext link in the Information Center to the **Search DB2 Online Information** search form.

The HTML search server is usually started automatically. If a search in the HTML information does not work, you may have to start the search server using one of the following methods:

On Windows
Click **Start** and select **Programs —> IBM DB2 —> Information —> Start HTML Search Server**.

On OS/2
Double-click the **DB2 for OS/2** folder, and then double-click the **Start HTML Search Server** icon.

Refer to the release notes if you experience any other problems when searching the HTML information.

Note: The Search function is not available in the Linux, PTX, and Silicon Graphics IRIX environments.

Using DB2 Wizards

Wizards help you complete specific administration tasks by taking you through each task one step at a time. Wizards are available through the Control Center and the Client Configuration Assistant. The following table lists the wizards and describes their purpose.

Note: The Create Database, Create Index, Configure Multisite Update, and Performance Configuration wizards are available for the partitioned database environment.

Wizard	Helps You to...	How to Access...
Add Database	Catalog a database on a client workstation.	From the Client Configuration Assistant, click **Add**.
Back up Database	Determine, create, and schedule a backup plan.	From the Control Center, right-click the database you want to back up and select **Backup —> Database Using Wizard**.
Configure Multisite Update	Configure a multisite update, a distributed transaction, or a two-phase commit.	From the Control Center, right-click the **Databases** folder and select **Multisite Update**.
Create Database	Create a database, and perform some basic configuration tasks.	From the Control Center, right-click the **Databases** folder and select **Create —> Database Using Wizard**.
Create Table	Select basic data types, and create a primary key for the table.	From the Control Center, right-click the **Tables** icon and select **Create —> Table Using Wizard**.
Create Table Space	Create a new table space.	From the Control Center, right-click the **Table Spaces** icon and select **Create —> Table Space Using Wizard**.
Create Index	Advise which indexes to create and drop for all your queries.	From the Control Center, right-click the **Index** icon and select **Create —> Index Using Wizard**.
Performance Configuration	Tune the performance of a database by updating configuration parameters to match your business requirements.	From the Control Center, right-click the database you want to tune and select **Configure Performance Using Wizard**. For the partitioned database environment, from the Database Partitions view, right-click the first database partition you want to tune and select **Configure Performance Using Wizard**.
Restore Database	Recover a database after a failure. It helps you understand which backup to use, and which logs to replay.	From the Control Center, right-click the database you want to restore and select **Restore —> Database Using Wizard**.

Setting Up a Document Server

By default, the DB2 information is installed on your local system. This means that each person who needs access to the DB2 information must install the same files. To have the DB2 information stored in a single location, perform the following steps:

1. Copy all files and subdirectories from `\sqllib\doc\html` on your local system to a Web server. Each book has its own subdirectory that contains all the necessary HTML and GIF files that make up the book. Ensure that the directory structure remains the same.
2. Configure the Web server to look for the files in the new location. For information, refer to the NetQuestion Appendix in the *Installation and Configuration Supplement*.
3. If you are using the Java version of the Information Center, you can specify a base URL for all HTML files. You should use the URL for the list of books.
4. When you are able to view the book files, you can bookmark commonly viewed topics. You will probably want to bookmark the following pages:
   - List of books
   - Tables of contents of frequently used books
   - Frequently referenced articles, such as the ALTER TABLE topic
   - The Search form

For information about how you can serve the DB2 Universal Database online documentation files from a central machine, refer to the NetQuestion Appendix in the *Installation and Configuration Supplement*.

Searching Information Online

To find information in the HTML files, use one of the following methods:

- Click **Search** in the top frame. Use the search form to find a specific topic. This function is not available in the Linux, PTX, or Silicon Graphics IRIX environments.
- Click **Index** in the top frame. Use the index to find a specific topic in the book.
- Display the table of contents or index of the help or the HTML book, and then use the find function of the Web browser to find a specific topic in the book.
- Use the bookmark function of the Web browser to quickly return to a specific topic.
- Use the search function of the Information Center to find specific topics. See “Accessing Information with the Information Center” on page 72 for details.

Appendix B. Notices

IBM may not offer the products, services, or features discussed in this document in all countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Canada Limited
Office of the Lab Director
1150 Eglinton Ave. East
North York, Ontario
M3C 1H7
CANADA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information may contain sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

Trademarks

The following terms, which may be denoted by an asterisk(*), are trademarks of International Business Machines Corporation in the United States, other countries, or both.

ACF/VTAM
AISPO
AIX
AIX/6000
AIXwindows
AnyNet
APPN
AS/400
BookManager
CICS
C Set++
C/370
DATABASE 2
DataHub
DataJoiner
DataPropagator
DataRefresher
DB2
DB2 Connect
DB2 Extenders
DB2 OLAP Server
DB2 Universal Database
Distributed Relational Database Architecture
DRDA
eNetwork
Extended Services
FFST
First Failure Support Technology
IBM
IMS
IMS/ESA
LAN DistanceMVS
MVS/ESA
MVS/XA
Net.Data
OS/2
OS/390
OS/400
PowerPC
QBIC
QMF
RACF
RISC System/6000
RS/6000
S/370
SP
SQL/DS
SQL/400
System/370
System/390
SystemView
VisualAge
VM/ESA
VSE/ESA
VTAM
WebExplorer
WIN-OS/2

The following terms are trademarks or registered trademarks of other companies:

Microsoft, Windows, and Windows NT are trademarks or registered trademarks of Microsoft Corporation.

Java or all Java-based trademarks and logos, and Solaris are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Tivoli and NetView are trademarks of Tivoli Systems Inc. in the United States, other countries, or both.

UNIX is a registered trademark in the United States, other countries or both and is licensed exclusively through X/Open Company Limited.

Other company, product, or service names, which may be denoted by a double asterisk(**) may be trademarks or service marks of others.

Appendix C. Contacting IBM

If you have a technical problem, please review and carry out the actions suggested by the *Troubleshooting Guide* before contacting DB2 Customer Support. This guide suggests information that you can gather to help DB2 Customer Support to serve you better.

For information or to order any of the DB2 Universal Database products contact an IBM representative at a local branch office or contact any authorized IBM software remarketer.

If you live in the U.S.A., then you can call one of the following numbers:

- 1-800-237-5511 for customer support
- 1-888-426-4343 to learn about available service options

Product Information

If you live in the U.S.A., then you can call one of the following numbers:

- 1-800-IBM-CALL (1-800-426-2255) or 1-800-3IBM-OS2 (1-800-342-6672) to order products or get general information.
- 1-800-879-2755 to order publications.

http://www.ibm.com/software/data/
The DB2 World Wide Web pages provide current DB2 information about news, product descriptions, education schedules, and more.

http://www.ibm.com/software/data/db2/library/
The DB2 Product and Service Technical Library provides access to frequently asked questions, fixes, books, and up-to-date DB2 technical information.

Note: This information may be in English only.

http://www.elink.ibmlink.ibm.com/pbl/pbl/
The International Publications ordering Web site provides information on how to order books.

http://www.ibm.com/education/certify/
The Professional Certification Program from the IBM Web site provides certification test information for a variety of IBM products, including DB2.

ftp.software.ibm.com
Log on as anonymous. In the directory `/ps/products/db2`, you can find demos, fixes, information, and tools relating to DB2 and many other products.

comp.databases.ibm-db2, bit.listserv.db2-l
These Internet newsgroups are available for users to discuss their experiences with DB2 products.

On Compuserve: GO IBMDB2
Enter this command to access the IBM DB2 Family forums. All DB2 products are supported through these forums.

For information on how to contact IBM outside of the United States, refer to Appendix A of the *IBM Software Support Handbook*. To access this document, go to the following Web page: http://www.ibm.com/support/, and then select the IBM Software Support Handbook link near the bottom of the page.

Note: In some countries, IBM-authorized dealers should contact their dealer support structure instead of the IBM Support Center.

Index

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.