IBM® DB2 Universal Database™

Glossary

Version 8

IBM® DB2 Universal Database™

Glossary

Version 8

Before using this information and the product it supports, be sure to read the general information under *Notices.*

You can order IBM publications online or through your local IBM representative.

- To order publications online, go to the IBM Publications Center at www.ibm.com/shop/publications/order
- To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at www.ibm.com/planetwide

To order DB2 publications from DB2 Marketing and Sales in the United States or Canada, call 1-800-IBM-4YOU (426-4968).

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

DB2 Glossary

A

abend. See “abnormal end of task”.

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

abnormal end of task (abend). 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) An exit that is not under program control, such as a trap or a 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 backward slash (\) character).

access and storage method. Associates XML documents to a DB2 database through two major access and storage methods: XML columns and XML collections. See also “XML collection” on page 109 and “XML column” on page 109.

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. The method that is selected by the database manager 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.

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.

access token. In DB2 Data Links Manager, an encrypted key that is assigned by the database manager and that must be generated to access a file under the control of the Data Links Manager.

accounting string. User-defined accounting information that is sent to DRDA® servers by DB2 Connect. This information can be specified from the client workstation using the SQLESACT API or the DB2ACCOUNT environment variable, or the DB2 Connect workstation using the DFT_ACCOUNT_STR database manager configuration parameter.

active log. (1) The primary and secondary log files that are currently needed for recovery and rollback. (2) The portion of the DB2 for z/OS and OS/390 log to which log records are written as they are generated. The active log always contains the most recent log records. See also “archive log” on page 5.

Activity Monitor. The DB2 tool that assists DBAs in monitoring application performance and concurrency, resource consumption, and SQL statement usage of a database or database partition. It

provides predefined views of application, statement, and SQL cache activity. DBAs can choose to view live monitor data and can record monitor data for further analysis. The tool also gives recommendations to resolve resource utilization problems.

address space. (1) The actual memory that is used by an active program. See also “buffer pool” on page 9. (2) In DB2 for z/OS and OS/390, a range of virtual storage pages that is identified by a number (ASID) and a collection of segment and page tables that map the virtual pages to real pages of the computer’s memory.

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

administrative authority. SYSADM and DBADM authority levels having full privileges for instance resources and database resources respectively.

administration notification log. A list of user-friendly, national language messages that should help the administrator to resolve minor issues. Also known as the DB2 notify log.

administration notification messages. Errors, warnings, and informational messages that are written by DB2, the Capture and Apply programs, and user applications to a notification file or event log. Also, alarms, warnings, attentions, and informational messages that are written by the health monitor to a notification file or event log.

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*.

administrator. A person responsible for administrative tasks such as access authorization and content management. Administrators can also grant levels of authority to users. See also “user” on page 105.

ADSM. ADSTAR Distributed Storage Manager. See “Tivoli Storage Manager” on page 100.

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. See also “Systems Network Architecture” on page 96.

Advanced Program-to-Program Communication (APPC). A protocol that allows interconnected systems to communicate and share the processing of programs. APPC supports Systems Network Architecture and uses LU 6.2 protocol for establishing sessions between systems. See also “Common Programming Interface Communications” on page 16.

after-image. The updated content of a source-table column that is recorded in a change data (CD) table, or in a database log or journal. See also “before-image” on page 8.

after trigger. A trigger that is activated after the triggering SQL operation has competed. The triggering operation can be an INSERT, UPDATE, DELETE, fullselect, or SIGNAL SQLSTATE statement. See also “before trigger” on page 8 and “trigger” on page 102.

agent. (1) A separate process or thread that carries out all DB2 requests that are made by a particular client application. See also “warehouse agent” on page 108. (2) For z/OS and OS/390 environments, the structure that associates all processes that are involved in a unit of work. See also “system agent” on page 96, “coordinating agent” on page 20, and “allied agent” on page 3.

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

aggregate function. Synonym for “column function” on page 15.

aggregate table. A read-only replication target table that contains aggregations of data from the source table. This data is based on SQL column functions such as MIN, MAX, SUM, or AVG.

alert. A signal representing a state of an object (such as a database, table space, or instance). See “health monitor alert” on page 47.

The types of alerts are listed in order of severity and include:

- **attention**

  An informational alert indicating that an object is in a non-normal state.
- **warning**

  A non-critical condition that does not require immediate attention but might indicate a non-optimal system.
- **alarm**

  A critical condition requiring immediate action.

alert condition. A user-specified condition for monitoring replication, that, when met, causes the Replication Alert Monitor to send notification to an individual or group contact that an error occurred or an operational threshold was reached.

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 system or subsystem, or a remote DB2 system or subsystem.

alias chain. A series of table aliases that refer to one another in a sequential, non-repeating fashion.

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

allied agent. Synonym for “allied thread”.

allied thread. A thread that originates at the local DB2 for z/OS and OS/390 subsystem and can access data at a remote DB2 for z/OS and OS/390 subsystem. See also “thread” on page 99.

allocated cursor. A cursor that is defined for stored procedure result sets by using the SQL statement ALLOCATE CURSOR.

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

ambiguous cursor. (1) A cursor is ambiguous if all of the following conditions are true:

- The SELECT statement is dynamically prepared.
- The SELECT statement does not include either the FOR READ ONLY clause or the FOR UPDATE clause.
- The LANGLEVEL bind option is SAA1.

- The cursor otherwise satisfies the conditions of a deletable cursor.

An ambiguous cursor is considered read-only if the BLOCKING bind option is ALL; otherwise, it is considered deletable. (2) In DB2 for z/OS and 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. See also "unambiguous cursor" on page 103.

American National Standard Code for Information Interchange (ASCII). An encoding scheme that is used to represent character strings in many environments, typically on personal computers and workstations. See also "EBCDIC" on page 37 and "Unicode" on page 104.

anti-join. An answer set in which the returned rows do not meet the condition of the join predicate. See also "join" on page 55.

APF. See "authorized program facility" on page 7.

API. See "application programming interface".

APPC. See "Advanced Program-to-Program Communication " on page 2.

APPL. A VTAM® network definition statement that is used to define DB2 for z/OS and 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; some examples are payroll, inventory management, and word processing applications.

Application Development Client. 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.

application-directed connections. A connection that an application manages using the SQL CONNECT statement. See also "system-directed connection" on page 96.

application ID. A unique string that is generated when the application connects to the database, or when DB2 Connect receives a request to connect to a Distributed Relational Database Architecture database. An identifier 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 name. The name of the application running on the client that identifies it to the database manager or DB2 Connect. It is passed from the client to the server to establish the database connection.

application plan. The control structure that is produced during the bind process. DB2 for z/OS and 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). A functional interface that allows an application program written in a high-level language to use specific data or functions of the operating system or another program, such as a database management system. 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. In DB2, APIs enable most of the administrative functions from within an application program.

application requester. The component on a remote system that generates DRDA requests for data on behalf of an application. An application requester accesses a DB2 database server using the DRDA application-directed protocol. See also “application server”.

application server. In DB2 for z/OS and OS/390, the target of a request from a remote application. In the DB2 environment, the application server function is provided by the distributed data facility and is used to access DB2 data from remote applications. See also “application requester”.

Apply control server. A database that contains the Apply control tables, which store information about registered replication source tables and subscription sets. Contrast with “control server” on page 19.

Apply cycle. The interval of time during which data is replicated from a source table to a target table.

Apply latency. An approximate measurement of the time that replication requires to complete one cycle. See also “Capture latency” on page 10.

Apply program. A program that is used to refresh or update a replication target table, depending on the applicable source-to-target rules. See also “Capture program” on page 10 and “Capture trigger” on page 10.

Apply qualifier. A case-sensitive character string that identifies replication subscription sets that are unique to an instance of the Apply program.

Apply server. A system where the Apply program is running. Contrast with “Apply control server”.

APPN. See “Advanced Peer-to-Peer Networking” on page 2.

archive log. (1) The set of log files that is closed and is no longer needed for normal processing. These files are retained for use in rollfoward recovery. (2) The portion of the DB2 for z/OS and OS/390 log that contains log records that are copied from the active log. The archive log holds records that are older and no longer fit on the active log. See also “active log” on page 1.

ASCII. See “American National Standard Code” on page 4.

AST. See “automatic summary table” on page 7.

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

asynchronous. Without regular time relationship, unpredictable with respect to the processing of program instructions. See also “synchronous” on page 95.

asynchronous batched update. A process in which all changes to the source are recorded and applied to existing target data at specified intervals. See also “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. See also “asynchronous batched update”.

asynchronous I/O. The nonsequential processing of read and write requests across multiple disks.

Glossary

asynchronous replication. The process of copying data from a source table to a target table outside the scope of the original transaction that updated the source table. Contrast with “synchronous replication” on page 95.

attach. To remotely access objects at the instance level.

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

attachment relationship type. In the Information Catalog Center, the relationship type that is used to attach comments to other objects. Comments can contain additional information about the object that they are attached to. See also “relationship type” on page 80.

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.

auditing. To record information following the detection of monitored data access by applications or individuals.

audit facility. A utility that generates a trail of audit records for a series of predefined and monitored database events.

audit log file. Location of audit records generated from the audit facility.

audit trail. Data, in the form of a logical path that links a sequence of events, used for tracing the transactions that affected the contents of a record.

authentication. The process by which a system verifies a user’s identity. User authentication is completed by a security facility outside of DB2, often part of the operating system or a separate product.

authority. See “authority level”.

authority level. A user’s access and ability to perform high-level database management operations such as maintenance and utility operations. A user’s authority level is used with privileges to control access to the database and its database objects. See also “load authority” on page 57, “system authority” on page 96, and “privilege” on page 75.

authorization. The process by which DB2 obtains information about the authenticated user, such as the database operations that the user may perform and which data objects that the user may access. See also “privilege” on page 75 and “authority level”.

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 Universal Database and to which a set of privileges is applied. An authorization identifier can represent an individual, an organizational group, or a function, but DB2 Universal Database does not determine this representation.

authorization token. (1) A token associated with a transaction. (2) For DB2 for z/OS and OS/390, the correlation ID. (3) For DB2 Universal Database for iSeries, the job name of the job that caused a transaction.

authorized program facility (APF). In DB2 for z/OS and 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 configuration parameters. A set of configuration parameters whose values can be changed automatically by the database manager to reflect the current resource utilization.

automatic rebind. 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 that it requires is not valid. See also "bind" on page 8 and "rebind" on page 78.

automatic summary table (AST). A summary table defined such that changes made to the underlying tables are cascaded to the summary table immediately and without the need for a REFRESH TABLE statement. See also "summary table" on page 95 and "materialized query table" on page 61.

auxiliary index. In DB2 for z/OS and OS/390, an index on an auxiliary table in which each index entry refers to a LOB. See also "auxiliary table".

auxiliary table. A table that stores columns outside the table in which they are defined. See also "base table".

B

backout. The process of undoing uncommitted changes that an application process has made. A backout may be necessary in the event of a failure on the part of an application process, or as a result of a deadlock situation. See also "rollback" on page 84.

backout free interval. A set of log records that are not compensated if the transaction aborts. See also "backout".

backup. A copy of a database or table space that can be stored on a different medium and used to restore the database or table space in the event of failure or damage to the original.

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 for z/OS and OS/390 scans the log in a backward direction to apply UNDO log records for all aborted changes.

base aggregate table. A type of replication target table that contains data that is aggregated from a replication source table. It includes a timestamp to mark the time when the Apply program performed the aggregation. See also "change-aggregate table" on page 12.

base table. (1) A table created with the CREATE TABLE statement. Such a table has both its description and data stored in the database. (2) In DB2 for z/OS and OS/390, a table created with the CREATE TABLE statement that contains a LOB column definition. The actual LOB column data is not stored with this base table. The base table contains a row identifier for each row and an indicator column for each of its LOB columns. See also "declared temporary table" on page 30, "auxiliary table", "view" on page 107, "result table" on page 83, and "temporary table" on page 99.

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

basic conversation. An SNA LU 6.2 conversation between two transaction programs that uses the APPC basic conversation API. See also “mapped conversation” on page 61.

basic predicate. A predicate that compares two values.

basic sequential access method (BSAM). An access method that DB2 for z/OS and OS/390 uses for storing or retrieving data blocks in a continuous sequence, using either a sequential access or a direct access device. See also “queued sequential access method ” on page 77.

before-image. The content of a replication source-table column prior to an update by a transaction, as recorded in a change data (CD) table, or in a database log or journal. See also “after-image” on page 2.

before trigger. A trigger that is activated before the triggering SQL operation has competed. The triggering operation can be a fullselect, SIGNAL SQLSTATE, or SET transition variable SQL statement. See also “trigger” on page 102 and “after-trigger” on page 2.

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 from 0 bytes to 2 gigabytes less 1 byte. This string does not have an associated code page and character set. BLOBs can contain image, audio, and video data. See also “character large object” on page 12 and “double-byte character large object” on page 36.

binary string. A sequence of bytes that is not associated with a CCSID. For example, the BLOB data type is a binary string. See also “coded character set identifier” on page 14.

bind. To convert the output from the SQL compiler to a usable control structure, such as an access plan, application plan, or package. During the bind process, access paths to the data are selected and some authorization checking is performed. See also “rebind” on page 78, “automatic rebind” on page 7, “dynamic bind” on page 36, “incremental bind” on page 49, “static bind” on page 92.

bind file. A file that is produced by the precompiler when the PRECOMPILE command or the respective API is used with the BINDFILE option.

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. (1) A string of data elements that is recorded or transmitted as a unit. (2) A set of contiguous data pages in a buffer pool. (3) A set of consecutive pages on disk.

block based I/O. A database manager method of reading contiguous data pages from disk into contiguous portions of memory. See also “scattered read” on page 85.

block fetch. A function of DB2 that retrieves (or fetches) a large set of rows together. Using block fetch can significantly reduce the number of messages that are sent across the network. Block fetch applies only to cursors that do not update data.

block identifier (BID). An entry that is stored along with a key value in the leaf node of a block index. This identifier references a particular block in a multidimensional clustering table.

block index. An index that is structured in the same manner as a traditional record identifier (RID) index, except that at the leaf level, keys point to a block identifier (BID) instead of an RID.

block locks. The locking of a block within a multidimensional clustering environment.

block map. A bitmap that contains an array of block states, one for each block in the multidimensional clustering table. Each entry has eight bits, four of which are used:

- In use: set to 1 if the block is considered part of the table; 0 otherwise (that is, it is free).
- Load: set to 1 for newly loaded blocks; reset to 0 when the load utility completes.
- Constraint pending: set to 1 for newly loaded blocks; reset to 0 after constraints are checked.
- Refresh pending: set to 1 for newly loaded blocks; reset to 0 after automated summary table maintenance is completed.

block size. Specifies the number of pages in a block. It is equal to the extent size. Also known as block factor.

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. See also “block fetch” on page 8.

bootstrap data set (BSDS). A VSAM data set that contains name and status information for DB2 for z/OS and OS/390, as well as relative-byte address-range specifications for all active and archive log data sets. It also contains passwords for the DB2 for z/OS and 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 database partitions.

browse. To view information catalog objects that are grouped by subject. Contrast with *search*.

browser. (1) A Text Extender function that enables you to display text on a computer screen. (2) A program that lets users to look at data but not change it.

BSAM. See “basic sequential access method” on page 8.

BSDS. See “bootstrap data set ”.

buffer manipulators. The processes used in backup and restore operations to read from or write to the database. By default, a single buffer manipulators is used; however, this can be overridden by using the parallelism option of the BACKUP DATABASE or RESTORE DATABASE commands.

buffer pool. An area of memory into which database pages are read, modified, and held during processing.

built-in function. An SQL function that is provided by DB2 and appears in the SYSIBM schema. See also “user-defined function” on page 106.

business dimension. A category of data, such as products or time periods, that an organization might want to analyze. See also “dimension” on page 33 and “multidimensional analysis” on page 63.

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 that they

need. For example, business metadata for a program would contain a description of what the program does and which tables it uses. See also “technical metadata” on page 98.

business name. In the Data Warehouse Center, a descriptive name that can be associated with an object that also has a physical name. The object types that can have business names are tables, files, columns or fields. The business name can be used in a search. It is also passed to end-user tools through the warehouse metadata interchange facilities.

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

caching. The process of storing frequently used results from a request to memory for quick retrieval, until it is time to refresh the information. DB2 Universal Database provides many forms of caching, such as directory caching, package caching, file system caching, and LDAP caching.

CAF. See “call attachment facility”.

call attachment facility (CAF). A DB2 for z/OS and 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 that is an alternative to an embedded SQL API. In contrast to embedded SQL, the CLI does not require precompiling or binding to a database, but instead provides a standard set of functions to process SQL statements and related services at run time. See also “DB2 Call Level Interface ” on page 28.

Capture control server. (1) A database that contains the Capture control tables, which store information about registered replication source tables. (2) A system where the Capture program is running.

Capture latency. An approximate measurement of how recently the Capture program committed data to a CD table. See also “Apply latency” on page 5.

Capture program. A program that reads database log or journal records to capture changes made to DB2 source tables. See also “Apply program” on page 5 and “Capture trigger”.

Capture schema. The schema for the Capture control tables used by a particular instance of the Capture program.

Capture trigger. In DB2 replication, a mechanism that captures delete, update, and insert operations performed on non-DB2 source tables. See also “Capture program” and “Apply program” on page 5.

cardinality. The number of rows in a database table.

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 Universal Database 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. An expression that allows another expression to be selected based on the evaluation of one or more conditions.

case-insensitive search. A search result without consideration of the case of the string being searched.

cast function. A function that is used to convert instances of a source data type into instances of a different target data type. In general, a cast function has the name of the target data type and has one single argument whose type is the source data type. Its return type is the target data type.

catalog. A set of tables and views that are maintained by the database manager. These tables and views contain information about the database, such as descriptions of tables, views, and indexes. See “information catalog” on page 51, “database catalog” on page 23, and “RDBMS catalog” on page 77.

catalog node. See “catalog partition”.

catalog partition. In a partitioned database environment, the database partition where the catalog tables for the database resides. Each database in a partitioned database environment can have its catalog partition on a different database partition server. The catalog partition for a database is automatically created on the database partition server where the CREATE DATABASE command is run.

catalog table. A table in the DB2 Universal Database catalog that is automatically created in the catalog when the database is created. These tables contain information about the database and its objects For example, these tables contain information about the definitions of database objects such as user tables, views, and indexes, as well as security information about the authority that users have on these objects. You cannot explicitly create or drop a catalog table, but you can query and view its contents using the catalog views.

catalog view. (1) A SYSCAT or SYSSTAT view on the catalog table. (2) A view that contains information about the tables and column that are enabled for use by the Text Extender. Also a view of the system table that is created by the Text Extender for administration purposes.

CCD table. See “consistent-change-data table” on page 18.

CCSID. See “coded character set identifier” on page 14.

CDB. See “communications database” on page 16.

CDRA. See “Character Data Representation Architecture” on page 12.

CD table. See “change data table” on page 12.

CelDial sample catalog. A sample information catalog (ICCSAMP) available when you install the Information Catalog Center. An administrator initializes the catalog, and users can use the sample data to become familiar with the Information Catalog Center.

cell. A unique combination of dimension values. Physically, a cell is made up of blocks of pages whose records all share the same values for each clustering column.

central processor complex (CPC). In a z/OS or OS/390 environment, a physical collection of hardware (such as an ES/3090 system) that consists of main storage, one or more central processors, timers, and channels.

CFRM. See "Coupling Facility Resource Manager" on page 21.

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

change-aggregate table. A type of replication target table that contains data that is aggregated from a replication source table. It includes a timestamp to mark the time when the Apply program performed the aggregation. See also "base aggregate table" on page 7.

change-capture replication. The process of capturing changes made to a replication source table and copying them to a replication target table. Contrast with "full refresh" on page 42.

change-data table. A replication table at the Capture control 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 conversion. The process of changing data from one character coding representation to another.

character large object (CLOB). A sequence of characters (single-byte, multibyte, or both) with a size ranging from 0 bytes to 2 gigabytes less 1 byte. In general, character large object values are used whenever a character string might exceed the limits of the VARCHAR type. Also called character large object string. See also "binary large object" on page 8 and "double-byte character large object" on page 36.

character set. A defined set of characters. For example, 26 nonaccented letters A through Z.

character string. A sequence of bytes that represent bit data, single-byte characters, or a mixture of single-byte and multibyte-byte characters.

character string delimiter. The characters that are used to enclose character strings in delimited ASCII files that are imported or exported. See also "delimiter" on page 31.

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

check condition. A restricted form of search condition that is 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. See "table check constraint" on page 97.

check integrity. The condition that exists when each row in a table conforms to the check constraints that are defined on that table. Maintaining check integrity requires DB2 to enforce table check constraints on operations that add or change data.

check pending. A state 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 the database manager records internal status information on the log; the recovery process uses this information if the subsystem abnormally terminates.

CI. See "control interval" on page 19.

CICS. See "Customer Information Control System" on page 22.

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

CIDF. See "control interval definition field" on page 19.

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. See also "database log" on page 23 and "archive log" on page 5.

claim. In DB2 for z/OS and OS/390, a notification to the database manager 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" on page 36.

claim class. In DB2 for z/OS and OS/390, a specific type of object access that can be one of the following types: cursor stability, repeatable read , or write.

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

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

class word. A single word that indicates the nature of a data attribute.

clause. In SQL, a distinct part of a statement, such as a SELECT clause or a WHERE clause.

clean block index. An index such that every record in a block that is referenced by the index has the same key value for that index. A dimension block index is a clean block index.

cleanse. (1) To ensure that all values in a data set are consistent and correctly recorded. (2) To transform the data extracted from operational systems to make it usable by the data warehouse.

CLI. See "call level interface" on page 10.

client. Any program (or workstation that a program is running on) that communicates with and accesses a database server. See also "requester" on page 82.

client profile. A profile that is used to configure clients using the Import function in the Configuration Assistant. It can contain database connection information, client settings, CLI or ODBC common parameters, and configuration data for local APPC or NetBIOS communication subsystems. See also "server profile" on page 88.

CLIST. See "command list" on page 15.

CLOB. See "character large object" on page 12.

CLP. See "command line processor" on page 15.

CLPA. See "create link pack area" on page 21.

clustered index. (1) 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. Synonym for "partitioning index" on page 71. (2) In DB2 for z/OS and OS/390, an index that determines how rows are physically ordered in a table space.

clustering block index. A block index that is automatically created for a particular dimension when the dimension is defined on a multidimensional clustering table. This index is used to maintain the clustering of data along that dimension, together with the other dimensions that are defined on the table. Also known as a dimension block index.

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. A unique bit pattern that represents a character in a code page.

code set. International Organization for Standardization (ISO) term for code page. See "code page".

cold start. (1) The process of starting a system or program by using an initial program load procedure. (2) A process by which DB2 for z/OS and OS/390 restarts without processing any log records. See also "warm start" on page 108. (3) For DB2 replication, the process of starting the Capture program using empty Capture control tables. Contrast with "warm start" on page 108.

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

collection. (1) In DB2 for z/OS and OS/390, a group of packages that have the same qualifier. (2) In the Information Catalog Center, a container for objects. A collection contains objects that the user has privileges to see, similar to a personal folder of objects.

collocated join. The result of two tables being joined when the tables reside in a single-partition database partition group in the same database partition; or they are in the same database partition group and have the same number of partitioning columns, the columns are partition-compatible, and both tables use the same partitioning function, and pairs of the corresponding partitioning key columns participate in the equijoin predicates. See also "partition-compatible join" on page 71.

column data. The data store that is stored in a DB2 column. The type of data can be any data type that is supported by DB2.

column distribution value. Statistics that describe the most frequent values of some column or the quantile values. These values are used in the DB2 optimizer to help determine the best access plan.

column function. (1) An operation that derives its result using values from one or more rows. (2) A function that performs a computation on a set of values rather than on a single value. Synonym for "aggregate function" on page 3. See also "scalar function" on page 85 and "table function" on page 97.

column options. In a federated system, parameters of the ALTER NICKNAME statement that describe the values in certain columns of the data source object that a nickname references. This information is added to the global catalog and used by the DB2 query optimizer to develop better access plans. Column options provide a way to tell the data source wrapper to handle a column in a different way than it normally would.

come-from checking. An SNA LU 6.2 security option that defines a list of authorization identifiers that are allowed to connect to DB2 for z/OS and OS/390 from a partner LU.

command. A way to start database administration functions to access and maintain the database manager. See also "DB2 command" on page 28.

command line processor (CLP). A character-based interface for entering SQL statements and database manager commands.

command list. A language that DB2 for z/OS and OS/390 uses to perform TSO tasks.

command prefix. In DB2 for z/OS and 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 DB2 for z/OS and 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 for z/OS and OS/390 subsystems.

command scope. In DB2 for z/OS and 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 shareable resources, or affects shareable resources.

comments object type. An object type that annotates another object in the Information Catalog Center. For example, you can attach a comment to a chart object that contains notes about the data in the chart. The comments object type is predefined in the Information Catalog Center. You cannot add properties to it.

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.

commit point. A point in time when data is considered to be consistent.

committed phase. The second phase of the multisite 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.

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. See also "Advanced Program-to-Program Communication " on page 2.

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 (a qualified SQL identifier). The expression can be specified as a table name in any FROM clause in the fullselect that follows the WITH clause. See also "table expression" on page 97.

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

comparison operator. 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). See also "infix operator" on page 51.

compensation. In a federated system, the ability of DB2 to process SQL that is not supported by a data source. DB2 will not push down a query fragment if the data source cannot process it, or if DB2 can process it faster than the data source can. If the data source cannot process it, DB2 will process it instead. There are two basic ways that a federated server compensates for the loss of functionality at the data source: it will simulate the data source function, or it will return the set of data to the federated server and perform the function locally. See also "query optimizer" on page 77 and "push-down processing" on page 76.

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 initially contains all of the rows from the replication source table or view and any predicates from the source table or view. See also "noncomplete CCD table" on page 65 and "consistent-change-data table" on page 18.

compose. In the XML Extender, to generate XML documents from relational data in an XML collection.

composite block index. An index that contains only dimension key columns and is used to maintain the clustering of data over insert and update activity in a multidimensional clustering (MDC) table. See also "dimension block index" on page 33.

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 for z/OS and 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 that indicates 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 consistent-change-data 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. See also “consistent-change-data table” on page 18.

condition. A specification of either the criteria for selecting XML data, or the way to join the XML collection tables.

conditional restart . In DB2 for z/OS and OS/390, a restart that is directed by a user-defined conditional restart control record (CRCR).

conditional restart control record. In DB2 for z/OS and OS/390, a queue of records in the bootstrap data set (BSDS) that is associated with a conditional restart of DB2. Each element in the queue indicates the choices that were made when the record was created (through the change log inventory utility, DSNJU003), and the progress of the restart operation it controls.

configurable configuration parameters. A set of configuration parameters that contain information that can be changed. See also “configurable online configuration parameters” and “informational configuration parameter” on page 51.

configurable online configuration parameters. A set of configuration parameters whose values can be changed while the database manager is running.

configuration file. A File that contains the values specified for configuration parameters. These parameters specify the resources that are allocated to the DB2 products and to individual databases, and the diagnostic level. There are two types of configuration files: the database manager configuration file for each DB2 instance, and the database configuration file for each individual database. Database manager configuration parameters are stored in a file named db2systm. Database configuration parameters are stored in a file named SQLDBCON. In a partitioned database environment, each database partition has its own database configuration file, but the same database manager configuration file is used by all partitions that participate in the instance.

configuration parameter. A parameter whose value limits the resources that can be used by the database manager or database. Some configuration parameters are informational, and display characteristics about the environment that cannot be changed.

conflict detection. In update-anywhere replication configurations, conflict detection refers to either of the following processes:

- The process of detecting constraint errors.
- The process of detecting whether the same row was updated by users or application programs in both the source and target tables during the same replication cycle. When a conflict is detected, the transaction that caused the conflict is rejected.

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 for z/OS and OS/390 subsystems that are connected and communicating by way of a conversation).

connection concentrator. A process that allows applications to stay connected without any resources being consumed on the DB2 host server. Thousands of users can be active in applications, while only a few threads are active on the DB2 host server.

connection handle. The data object containing information that is associated with a connection that is managed by DB2 ODBC. This information includes general status information, transaction status, and diagnostic information. See also "statement handle" on page 92.

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

connection pooling. A process in which DB2 Connect drops the inbound connection with an application that requests disconnection, but keeps the outbound connection to the host in a pool. When a new application requests a connection, the DB2 Connect uses one from the existing pool. Using the already-present connection reduces the overall connection time, as well as the high processor connect cost on the host.

consistency token. A timestamp that is used to generate the unique identifier (version identifier in DB2 for z/OS and OS/390) for an application.

consistent-change-data (CCD) table. A type of replication target table that is used for storing history, for auditing, or for staging data. A CCD table can also be a replication source. See also "complete CCD table" on page 16, "condensed CCD table" on page 17, "external CCD table" on page 40, "internal CCD table" on page 53, "noncomplete CCD table" on page 65, and "noncondensed CCD table" on page 65.

constant. A language element that specifies an unchanging value. Constants are classified as string constants or numeric constants. See also "variable" on page 107.

constraint. A rule that limits the values that can be inserted, deleted, or updated in a table. See also "check constraint" on page 12, "referential constraints" on page 79, and "unique constraint" on page 104.

contact relationship type. In the Information Catalog Center, the relationship type that is used to identify contacts. A contact relationship type provides more information about an object. Such information might include the person who created the information that the object represents or the department that is responsible for maintaining the information. See also *relationship type*.

contains relationship type. In the Information Catalog Center, the relationship type that is used to identify Information Catalog Center objects that contain other objects. For example, use the contains relationship type to denote an object with a ″parent″ role, meaning that object can contain other objects. You can also use the contains relationship type to denote an object with a ″child″ role, meaning an object that can be contained in another object. See also *relationship type*.

container. (1) A physical storage location of the data. For example, a file, directory, or device. (2) See "table space container" on page 97.

contention. A situation in which a transaction attempts to lock a row or table that is already locked.

Control Center. The DB2 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 various tools, such as Replication Center, Health Center, Task Center, and Journal. See also "DataJoiner Replication Administration tool" on page 25 and "Satellite Administration Center" on page 85.

contracting conversion. A process that occurs when the length of a converted string is smaller than that of the source string. See also “expanding conversion” on page 38.

control interval. 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 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” on page 73 and “control point name”.

control point name. A network-qualified name of a control point that consists of a network identifier qualifier that identifies the network to which the control point node belongs. See also “control point”.

control privilege. The authority to completely control an object, which 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. A database server that contains replication control tables for the Capture, Apply, or Monitor program. See also “Apply control server” on page 5, “Capture control server” on page 10, and “Monitor control server” on page 63.

control table. See “replication control table” on page 82.

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 identifier or group identifier and password before establishing a connection.

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

coordinate. In DB2 Spatial Extender, a number that denotes a position that is relative to a point of reference. For example, in a map of the earth, a place can be referenced by both (a) a coordinate that denotes the place’s position relative to the equator and (b) a coordinate that denotes the place’s position relative to the Greenwich meridian.

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

coordinating agent. The agent that is started when the database manager receives a request from an application. The agent remains associated with the application during the life of the application. This agent initiates subagents that work for the application. See also "agent" on page 2 and "subagent" on page 93.

coordinator. In DB2 for z/OS and OS/390, the system component that initiates the commit or rollback of a unit of work that includes work that is done on one or more other systems.

coordinator node. See "coordinator partition".

coordinator partition. The database partition server 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. (1) A subquery that contains a correlated reference to a column of a table that is outside the subquery. (2) In DB2 for z/OS and OS/390, a subquery that is part of a WHERE or HAVING clause that is applied to a row or group of rows of a table or view that is named in an outer subselect statement.

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

correlation name. An identifier that designates a table or view within a single SQL statement. The name can be defined in any FROM clause or in the first clause of an UPDATE or DELETE statement.

cost. The estimated total resource usage that is necessary to run the access plan for a statement (or the elements of a statement). Cost is derived from a combination of processor cost (in number of instructions) and I/O (in numbers of seeks and page transfers).

cost category. A category into which DB2 for z/OS and 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:

- DB2 for z/OS and OS/390 has enough information to make a cost estimate without using default values.
- Some condition exists for which DB2 for z/OS and 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.

counter. A representation of information that is cumulative up until the sample is taken. The counter counts values that increase, such as the number of deadlocks. Counters are reset when you stop and restart an instance or database. See also "gauge" on page 44.

country / region code. The two–character representation for the country or region, that is used to establish monetary, date, and numeric formatting.

Coupling Facility Resource Manager. In a z/OS or OS/390 environment, the Coupling Facility Resource Manager manages all of the coupling facilities in a Sysplex.

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

CP. See "control point" on page 19.

CPC. See "central processing complex" on page 12.

CPI-C. See "Common Programming Interface Communications" on page 16.

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. See "control point name" on page 19.

crash recovery. The process of bringing a database back to a consistent and usable state after a failure. See also "version recovery" on page 107 and "forward recovery" on page 42.

CRC. See "command recognition character" on page 15.

CRCR. In DB2 for z/OS and OS/390, conditional restart control record. See also "conditional restart" on page 17.

create link pack area (CLPA). An option that is used during initial program load to initialize the link pack pageable area.

created temporary table. In DB2 for z/OS and OS/390, a table that holds temporary data and is defined with the SQL statement CREATE GLOBAL TEMPORARY TABLE. Information about created temporary tables is stored in the DB2 catalog, so this kind of table is persistent and can be shared across application processes. See "temporary table" on page 99. See also "declared temporary table" on page 30.

cross-memory linkage. In a 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 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" on page 22.

CSA. See "common service area" on page 16.

cumulative backup. See “incremental backup” on page 49.

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

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

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

current SQL ID. An identifier that, at a single point in time, holds the privileges that are exercised when certain dynamic SQL statements run. The current SQL ID can be a primary authorization ID or a secondary authorization ID.

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

cursor. A named control structure that is 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 blocking. A technique that reduces overhead by retrieving a block of rows in a single operation. These rows are cached while they are processed.

cursor sensitivity. The degree to which database updates are visible to the subsequent FETCH statements in a cursor. A cursor can be sensitive to changes that are made with positioned UPDATE and DELETE statements that specify the name of the cursor. A cursor can also be sensitive to changes that are made with searched UPDATE or DELETE statements, or with cursors other than this cursor. These changes can be made by this application process or by another application process.

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. See also “read stability” on page 78, “repeatable read ” on page 81, and “uncommitted read (UR)” on page 103.

Customer Information Control System (CICS). An IBM licensed program that provides online transaction-processing services and management for critical business applications. In DB2 for z/OS and OS/390 information, this term represents CICS Transaction Server for z/OS and OS/390, CICS/ESA, and CICS/MVS.

cycle. In DB2 for z/OS and 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. For example, a self-referencing table is a cycle with a single member.

cyclical referential constraint. A table that is a dependent of, or descendent of, another table.

D

DAD. See “Document Access Definition” on page 35.

daemon. A system process that provides a specific service to applications or users.

DARI. See “Database Application Remote Interface”.

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

database access thread. In DB2 for z/OS and OS/390, a thread that accesses data at the local subsystem on behalf of a remote subsystem. See also “allied thread” on page 3.

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

database agent. A representation for the physical process or thread that will do the actual work inside the database engine.

Database Application Remote Interface (DARI). Obsolete term for “stored procedure” on page 93.

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 that is used to access a database that is on a database server.

database configuration parameter. A parameter whose value limits the system resources that a database can use. See also “configuration parameter” on page 17 and “database manager” on page 24.

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

database descriptor (DBD). An internal representation of a DB2 for z/OS and OS/390 database definition, which reflects the data definition that is in the DB2 for z/OS and 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. See also “node directory” on page 65.

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

database function. The relationship between a set of input data and a set of result values. See also “built-in function” on page 9 and “user-defined function” on page 106.

database log. A set of primary and secondary log files that consist 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. See also “system-managed space table space ” on page 96.

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 configuration parameter. A configuration parameter that is established when the instance is created. Most database manager configuration parameters affect the amount of system resources that will be allocated to a single instance of the database manager, or they configure the setup of the database manager and the different communications subsystems based on environmental considerations. See also “configuration parameter” on page 17 and “database configuration parameter” on page 23.

database manager instance. (1) A logical database manager environment similar to an image of the actual database manager environment. It is possible have several instances of the database manager product on the same workstation. Use these instances to separate the development environment from the production environment, tune the database manager to a particular environment and protect sensitive information. (2) The DB2 code that manages data. An instance has its own databases (which other instances cannot access), and all its database partitions share the same system directories. It also has separate security from other instances on the same computer.

database name. The identifying name that a user provides as part of the CREATE DATABASE command or application programming interface. A database name must be unique within the location in which it is cataloged.

database node. See “database partition”.

database object. (1) An association within a database to anything that can be monitored. (2) Anything that can be created or manipulated with SQL. Tables, views, indexes, packages, triggers, and table spaces are database objects.

database object hierarchy. An arrangement of database objects into parent/child relationships. For example, a database is the child of its database instance parent.

database partition. In a partitioned database environment, a part of the database that consists of its own user data, indexes, configuration files, and transaction logs.

database partition group. In a partitioned database environment, a named set of one or more database partitions. This term replaces the term nodegroup.

database partition server. In a partitioned database environment, an occurrence of DB2 that is recorded in the db2nodes.cfg file.

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

database server. The target of a request from a local application or an intermediate database server. In the DB2 environment, the database server function is provided by the distributed data facility to access DB2 data from local applications or a remote database server that is acting as an intermediate database server.

database system monitor. A collection of APIs that collect information regarding the state of the database system at the instance, database, and application levels. This information is stored in data elements, which can be examined by taking point-in-time snapshots, or by using the event monitor to log system activity over a period of time.

data blocking. The process of replicating a specific number of minutes' worth of change data during an Apply cycle.

data currency. In DB2 for z/OS and 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.

data definition name (ddname). In DB2 for z/OS and 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".

data consolidation. A replication configuration that contains one read-only target database. The target table contains rows of data from one or more source databases.

data dictionary. A repository of information about an organization's application programs, databases, logical data models, users, and authorizations. A data dictionary can be manual or automated.

data distribution. A replication configuration that contains a single source database, from which changes are replicated to one or more read-only target tables. Before replication to the target tables can occur, the tables must contain a complete set of data from the source table.

data element. A data structure that is used by the system monitor to store information about the state of the database system. Data elements collect data for one or more logical data groups. Examples of data structures are counters, gauges, information, and timestamps. See also "logical data group" on page 59.

data interchange. The sharing of data between applications. XML supports data interchange without needing to go through the process of first transforming data from a proprietary format.

DataJoiner. A product that provides client applications that are 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. See also "federated server" on page 41.

DataJoiner Replication Administration (DJRA) tool. A database administration tool that you can use to perform various replication administration tasks. Unlike the DB2 Control Center, the DJRA tool can be used to administer replication for non-IBM databases. See also "Control Center" on page 18.

DATALINK. An SQL 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 Link Reconcile Not Possible (DRNP). The state of a DB2 table in which one or more DATALINK type columns contain file references whose integrity is violated (for example, as the result of restoring a database without the ability to restore the files that the database refers to).

Data Link Reconcile Pending (DRP). The state of a DB2 table in which one or more DATALINK type columns contain file references whose integrity might be in doubt (for example, as the result of restoring a database without rolling forward through the database logs).

Data Links Filesystem Filter (DLFF). A DB2 Data Links Manager component. A file system filter program that enforces data integrity by ensuring valid and controlled access to linked files. See also “linked file” on page 56.

Data Links File Manager (DLFM). A component of the DB2 Data Links Manager that enables a DB2 database to manage files that are outside of the database.

Data Links File System (DLFS). A file system that is under the control of the Data Links Filesystem Filter (DLFF).

Data Links Manager Administrator. The person and the user ID that is responsible for administering the DB2 Data Links Manager and its associated environment. Sometimes also referred to as *DLFM User*, because when DB2 Data Links Manager is installed, an account with the default user ID of dlfm is set up for use by the Data Links Manager Administrator.

The Data Links Manager Administrator user ID also owns all of the resources used by the DLFM component, for example:

- The DB2 instance containing the DLFM_DB database
- Linked files referenced in a READ PERMISSION DB DATALINK type column
- The user ID under which the DLFM Server itself runs

See also “dlmadmin” on page 35 and “superuser” on page 95.

Data Links server. A computer that contains these DB2 Data Links Manager components: a Data Links File Manager (DLFM), a Data Links Filesystem Filter (DLFF) controlling a Data Links File System (DLFS), and a DB2 database (used as the Logging Manager).

Data Manager Application (DMAPP). A component of the DB2 Data Links Manager configuration in a DCE-DFS environment that provides controlled access to the files in the DFS cell.

data manipulation language (DML). A subset of SQL statements that is used to manipulate data. Most applications primarily use DML SQL statements, which are supported by the DB2 Connect program. SELECT, INSERT, UPDATE, and DELETE statements are similar across the IBM relational database products. See also “Structured Query Language” on page 93 and “data definition language” on page 25.

data mart. A subset of a data warehouse that contains data that is tailored and optimized for the specific reporting needs of a department or team. A data mart can be a subset of a warehouse for an entire organization, such as data that is contained in online analytical processing (OLAP) tools.

data mining. The process of collecting critical business information from a data warehouse, correlating the information and uncovering associations, patterns, and trends.

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

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

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

data sharing member. (1) A local or remote relational or nonrelational data manager that is capable of supporting data access via an ODBC driver that supports the ODBC APIs. (2) In a federated system, typically a relational DBMS instance and one or more databases that are supported by that instance. You can also include other types of data sources in a federated system, such as flat-file databases and table-structured files.

data source. In a federated system, typically a relational DBMS instance and one or more databases supported by that instance. However, there are other types of data sources that you can include in your federated system, such as flat-file databases and table-structured files.

data source objects. In a federated system, objects at the data source on which you want to perform operations. Examples include a database table, a database view, or a spreadsheet list. You create nicknames on the federated server to identify the data source objects. See also “nickname” on page 65.

data space. In DB2 for z/OS and OS/390, space ranging in size from 0 bytes 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 type mapping. In a federated system, the mapping of a data type used at a data source to a DB2 data type. For example, the Oracle type FLOAT maps by default to the DB2 type DOUBLE. DB2 supplies default mappings for most kinds of data types; the default mappings are in the wrappers.

data warehouse. (1) A subject-oriented nonvolatile collection of data that is 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 data marts within an enterprise and delivers a common view of enterprise data. (2) A central repository for all or significant parts of the data that an organization’s business systems collect. Also known as an *information warehouse*. See also “data mart” on page 26.

Data Warehouse Center. The component of DB2 Universal Database that provides the 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 computers 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 are Data Warehouse Center programs.

Data Warehouse Center property. An attribute that applies across sessions of the Data Warehouse Center, such as the tools catalog that contains the technical metadata. See also “property” on page 75.

date. A three-part value that designates a day, month, and year. For example, YYYY-MM-DD.

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 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 DB2 Control Center.

DBA. See "database administrator" on page 23.

DBCLOB. See "double-byte character large object" on page 36.

DBCS. See "double-byte character set" on page 36.

DBD. See "database descriptor" on page 23.

DBID. In DB2 for z/OS and OS/390, a database identifier.

DBMS. See "database management system" on page 24.

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

DBRM. See "database request module" on page 24.

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

DB2 Call Level Interface (CLI). An application that uses a standard set of functions to process SQL statements and related services at run time. It does not have to be precompiled or bound.

DB2 client. Allows access to a remote database without knowing its physical location. The DB2 client determines the location of the database, manages the transmission of requests to the database server, and returns the results.

DB2 command. An instruction to the operating system to access and maintain the database manager. For example, DB2 commands allow a user to start or stop a database, display information on current users and the status of databases.

DB2 Connect. A product that enables client applications to read and update data that is stored on host or iSeries servers.

DB2 control server. A DB2 Universal Database system that contains the satellite control database, SATCTLDB.

DB2 Data Links Manager. A separately orderable DB2 feature that enables your applications to manipulate data residing in both unstructured files and in the relational database management system (RDBMS). DB2 Data Links Manager enables DB2 Universal Database to manage unstructured files as though they were directly stored in the database and provides the integration between the RDBMS and the external file systems through extensions to DB2 Universal Database.

DB2 DataPropagator. A product that provides DB2 replication for OS/390, z/OS, OS/400, z/VM, VM, and VSE operating-system environments. For UNIX and Windows operating-system environments, replication is integrated with DB2 and does not require a separate license. See also "replication" on page 82.

DB2DC. See "Development Center" on page 33.

DB2 Download Tool. A tool that performs high-speed data transfers between an MVS and an SP system.

DB2DT. See "DB2 Download Tool".

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.

DB2 host. In a DB2 Data Links Manager configuration, a DB2 database, on a DB2 server, that contains a DATALINK column.

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

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

DB2 Life Sciences Data Connect. A database middleware system that allows you to run a single query on a virtual database, whose underlying data can be stored in multiple life sciences industry data sources.

DB2 Net Search Extender. A program that provides full-text retrieval through a DB2 stored procedure. The Net Search Extender is primarily optimized for performance. Using DB2 Net Search Extender can be particularly advantageous in applications where search performance on large indexes and scalability according to concurrent queries are important factors.

DB2 PM. DATABASE 2 Performance Monitor for z/OS and OS/390.

db2_recon_aid utility. A utility that identifies database tables containing DATALINK type columns, and optionally runs the DB2 Reconcile utility on these tables.

DB2 Relational Connect. A product that is used in a federated system to query and retrieve data that is located in other database managers, such as Oracle, Informix, Sybase, Microsoft SQL Server, and members of the DB2 Universal Database family, such as DB2 for z/OS and OS/390, DB2 Universal Database for iSeries, and DB2 Universal Database for Windows.

DB2 SDK. See "DB2 Application Development Client" on page 28.

db2setup utility. A utility that guides users through the installation process with a graphical interface and online help. You can use this utility to create or assign groups and user IDs, create a DB2 instance, and install product messages. Default values are provided for all required installation parameters.

DB2 Spatial Extender. A program that is used to create a geographic information system .

DB2 Text Extender. A full text retrieval system integrated in DB2 Universal Database that provides powerful search features enhanced by additional rich linguistic functionality for applications with highly

structured documents where the information need is complex, and the quality and precision of the search result are key issues over and above system response times.

DB2 tools catalog. A set of tables or files that is maintained by the database tools (Data Warehouse Center, Control Center, Task Center, Information Catalog Center) and contains information about the processes and tasks that DB2 runs, such as loads, reorgs, database maintenance processes, data movement processes, and the associated schedules, logs, and dependencies.

DB2 tools metadata. The information about the processes and tasks that DB2 runs, such as loads, reorgs, database maintenance processes, data movement processes, and the associated schedules, logs, and dependencies. The DB2 tools metadata is contained in the DB2 tools catalog.

DB2 XML Extender. A program that is used to store and manage XML documents in DB2 tables. Well-formed and validated XML documents can be generated from existing relational data, stored as column data, and the content of XML elements and attributes can be stored in DB2 tables.

DCE-DFS. See "Distributed Computing Environment" on page 34.

DCLGEN. See "declarations generator".

DDF. See "distributed data facility" on page 34.

DDL. See "data definition language" on page 25.

ddname. See "data definition name (ddname)" on page 25.

deadlock. A condition under which a transaction cannot proceed because it is dependent on exclusive resources that are locked by another transaction, which in turn is dependent on exclusive resources that are 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 transaction can proceed.

decision-support system. In the Information Catalog Center, a system of applications that help users make decisions. This kind of system allows users to work with information that is presented in meaningful ways; for example, spreadsheets, charts, and reports.

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

declared temporary table. A table that holds temporary data and is defined with the SQL statement DECLARE GLOBAL TEMPORARY TABLE. Information about declared temporary tables is not stored in the DB2 catalog, so this kind of table is not persistent and can be only used by the application process that issued the DECLARE statement. See also "base table" on page 7, "created temporary table" on page 21, and "temporary table" on page 99.

decompose. In XML Extender, to separate XML documents into a collection of relational tables in an XML collection.

default subsystem name (DSN). (1) In the z/OS or OS/390 environment, the name of the TSO command processor of DB2. (2) The name of the DB2 subsystem that can connect to the control server (the default subsystem name is DSN). (3) In the z/OS or OS/390 environment, the first three characters of DB2 module and macro names.

default view. In XML Extender, a representation of data in which an XML table and all of its related side tables are joined.

deferred embedded SQL. 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.

deferred write. In DB2 for z/OS and OS/390, the process of asynchronously writing changed data pages to disk.

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. The number of concurrently executed operations that are initiated to process a query.

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

delete history. In the Information Catalog Center, a log of delete activity, the capture of which is turned on and off by the Information Catalog Center administrator. The log can be transferred to a tag language file.

delete hole. A row for a SELECT statement of a cursor that no longer has a corresponding row in the base table because the row was deleted. A delete hole is created when a row in the base table is deleted while a cursor is open whose SELECT statement result contains the row that is deleted. Such a row is no longer accessible though the cursor. See also “hole” on page 47 and “update hole” on page 105.

delete rule. A rule that is 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. A trigger that is defined with the triggering SQL operation DELETE. See also “trigger” on page 102.

delimited identifier. A sequence of characters enclosed within 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. See also “ordinary identifier” on page 68.

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.

delta backup. A copy of all database data that has changed since the last successful backup (full, incremental, or delta) of the table space in question. A delta backup is also known as a differential, or noncumulative, backup image. The predecessor of a delta backup image is the most recent successful backup that contains a copy of each of the table spaces in the delta backup image.

denormalization. The intentional duplication of columns in multiple tables whose consequence is increased data redundancy. Denormalization is sometimes necessary to minimize performance problems and is a key step in designing a physical relational database design. See also "normalization" on page 66.

dependent. In SQL, an object (row, table, or table space) that has at least one parent. See also "parent row" on page 70, "parent table" on page 71, and "parent table space" on page 71.

dependent foreign key table. A dependent foreign key table of a given table is a table that has at least one foreign key constraint referencing the given table.

dependent immediate materialized query table. A dependent materialized query table that is defined with the REFRESH IMMEDIATE option.

dependent logical unit (DLU). A logical unit that requires assistance from a system services control point (SSCP) to instantiate an LU-to-LU session. See "independent logical unit" on page 50.

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. See also "parent row" on page 70.

dependent materialized query table. A materialized query table that references a given table directly or indirectly (For example, from a view) in its materialized query table definition.

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

dependent table space. A table space that contains a dependent of a parent table. See also "parent table space" on page 71.

derived data. In the Information Catalog Center, data that is copied or enhanced (perhaps by summarizing the data) from operational data sources into an informational database.

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

descendent immediate materialized query table. A descendent immediate materialized query table is a materialized query table defined with the REFRESH IMMEDIATE option that directly refers to a descendent immediate materialized query table in its materialized query table definition.

descendent materialized query table. A materialized query table that references in its materialized query table definition a descendent materialized query table directly or indirectly.

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 has a parent of another table, or a descendent of a dependent table.

descriptive data. See "metadata" on page 62.

DETERMINISTIC 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 "not-deterministic function" on page 66.

Development Center. A component of a DB2 that provides a graphical interface for building, testing, and deploying stored procedures and user-defined functions. Features include a server view, an integrated SQL debugger, export and import wizards, and an editor.

Development Center project. A file that is created by the Development Center to manage development work. A project that contains information about database connections and routines that are being developed with the Development Center.

device name. A name reserved by the system or a device driver that refers to a specific device. For example, the DOS device name for the parallel port is LPT1.

DFP. In a z/OS or 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.

dictionary relationship type. In the Information Catalog Center, the relationship type that is used to associate a glossary entry object type with another object. A glossary entry object type can be used to define terminology that is associated with the object. See also “relationship type” on page 80.

differential backup image. See “delta backup” on page 31.

differential refresh. See “change-capture replication” on page 12.

dimension. A data category, such as time, accounts, products, or markets. The elements of a dimension are referred to as members. Dimensions offer a very concise, simple way of organizing and selecting data for retrieval, exploration, and analysis. Dimensions also represent the highest consolidation level in a multidimensional database outline. See also “business dimension” on page 9, “multidimensional analysis” on page 63, and “dimension table”.

dimension block index. In multidimensional clustering, a block index that is automatically created for a particular dimension when the dimension is defined on an MDC table. This index is used to maintain the clustering of data along that dimension, together with the other dimensions defined on the table.

dimension table. The representation of a dimension in a star schema. Each row in a dimension table represents all of the attributes for a particular member of the dimension. See also “dimension” and “star schema” on page 92.

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 one table participate in the equijoin predicates, the other table is rehashed; otherwise (if there is at least one equijoin predicate), both tables are rehashed. See “join” on page 55.

directory. The DB2 for z/OS and 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 that were created during the enabling process.

disaster recovery. The activities that are required to restore the database in the event of a fire, earthquake, vandalism, or other catastrophic events. Typically, disaster recovery requires that you restore the entire database, so when a major disaster occurs, a full database backup is needed on a standby site.

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 - Distributed File System (DCE-DFS). A transarc product that provides the security, performance, scalability, and ease-of-use of a native file system in a distributed, networked environment. This file system environment is supported by DB2 Data Links Manager in the AIX operating environment.

distributed data facility (DDF). A set of DB2 for z/OS and OS/390 components through which DB2 for z/OS and 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 installation. A process by which DB2 products can be installed using systems management software, such as Microsoft Systems Management Server (SMS) on Windows NT or Windows 2000, or simply with a shared CD-ROM drive or shared network hard drive using response files. Also known as a silent installation or unattended installation.

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 transaction. A transaction that updates data in more than one database. See also “two-phase commit” on page 103.

distributed unit of work. 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. See also “Control Center” on page 18.

DLC. See “data link control” on page 25.

DLFF. See “Data Links Filesystem Filter (DLFF)” on page 26.

DLFM. See “Data Links File Manager (DLFM)” on page 26.

DLFM_ASNCOPYD file-copy daemon (Data Links Manager Replication daemon). The DLFM process which enables replication of DB2 Data Links Manager files (in conjunction with the associated DB2 relational data) in support of data replication.

DLFM_DB database. A DB2 database that acts as a logging manager for the Data Links server.

DLFS. See "Data Links File System (DLFS)" on page 26.

DLM. See "DB2 Data Links Manager" on page 28.

dlmadmin account. In DB2 Data Links Manager, an account that is created at install time on Windows NT and Windows 2000 environments only. By default, the account is named dlmadmin. This account has advanced user privileges, and is intended to be equivalent to the root user in UNIX environments. Its purpose is to act as a superuser to perform any necessary advanced administration operations by both the DLFM component and the Data Links Manager Administrator on the Data Links server. However, unlike the DLFM User account, dlmadmin does not own any of the DLFM resources. The Data Links Manager Administrator uses this account in addition to the ″DLFM User″ account.

DLU. See "dependent logical unit" on page 32.

DMAPP. See "Data Manager Application (DMAPP)" on page 26.

DML. See "data manipulation language" on page 26.

DMS table space. See "database-managed space table space" on page 23.

DNS. See "domain name server (DNS)".

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.

Document Type Definition (DTD). A set of declarations for XML elements and attributes. The DTD defines what elements are used in the XML document, in what order they can be used, and which elements can contain other elements. You associate a DTD with a document access definition (DAD) file to validate XML documents.

domain. A part of a network that is administered as a unit with a common protocol.

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. For example, www.ibm.com.

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.

Domino™ *Go* Web server. The Web server 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, with a size ranging from 0 bytes to 2 gigabytes. A data type that can be used to store large double-byte text objects. Such a string always has an associated code page. See also “binary large object” on page 8 and “character large object” on page 12.

double-byte character set (DBCS). A set of characters in which each character is represented by two bytes. These character sets are commonly used by national languages, such as Japanese and Chinese, that have more symbols than can be represented by a single byte. See also “single-byte character set” on page 89 and “multibyte character set” on page 63.

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

drain. In DB2 for z/OS and OS/390, the act of acquiring a locked resource by quiescing access to that object. See also “claim” on page 13.

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

DRDA. See “Distributed Relational Database Architecture” on page 34.

DRDA access. An open method of accessing distributed data by which you can connect to another database server (by location), using an SQL statement, to execute packages that have been previously bound at that location. The SQL CONNECT statement or a three-part name SQL statement is used to identify the server. See also “private protocol access” on page 74.

DRNP. See “Data Link Reconcile Not Possible (DRNP)” on page 26.

DRP. See “Data Link Reconcile Pending (DRP)” on page 26.

DSN. See “default subsystem name” on page 31.

DTD. See “Document Type Definition” on page 35.

DTD reference table. A table that consists of DTDs, which are used to validate XML documents and to help applications to define a DAD. This table is created when a database is enabled for XML. Users can insert their own DTDs into the DTD_REF table.

DUOW. See “distributed unit of work” on page 34.

dual log path. A secondary log path that is used to maintain duplicate copies of online archived files and the active log.

duration. In SQL, a number that represents an interval of time. See “date duration” on page 28, “labeled duration” on page 55, and “time duration” on page 99.

dynamic bind. A process by which SQL statements are bound as they are entered. See “bind” on page 8. See also “static bind” on page 92.

dynamic SQL. SQL statements that are prepared and executed at run time. In dynamic SQL, the SQL statement is contained as a character string in a host variable and is not precompiled. See also “embedded SQL” on page 37 and “static SQL” on page 93.

E

EA-enabled table space. In DB2 for z/OS and 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 gigabytes.

EBCDIC. A coded character set of 256 8-bit characters developed for the representation of textual data, typically used on zSeries and iSeries servers. See also "ASCII" on page 5 and "Unicode" on page 104.

edition. See "step edition" on page 93.

EDM. Electronic data management.

EID. Event identifier.

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

element. See "XML element" on page 109.

embedded SQL. SQL statements coded within an application program. See "static SQL" on page 93.

EN. See "end node ".

enable. (1) To prepare a database, a text table, or a text column for use by the Text Extender or the XML Extender. (2) To turn on or activate.

enclave. In Language Environment (which is used by DB2 for z/OS and 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. In APPN, a node that supports sessions between its local control point and the control point in an adjacent network node.

end-to-end latency. See "Apply latency" on page 5.

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 "conflict detection" on page 17.

entity. (1) A person, object, or concept about which you want to store information. In a relational database, entities are represented as tables. A database includes information about the entities in an organization or business, and their relationships to each other. (2) A unit of data that can be classified and have stated relationships to other entities within that database.

enumerated list. In DB2 for z/OS and OS/390, a set of DB2 objects that are defined with a LISTDEF utility control statement in which pattern-matching characters (*, %, _ , or ?) are used.

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 operation in which the join-condition has the form *expression = expression*.

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

escape character. See "SQL escape character" on page 91.

ESDS. In a z/OS or OS/390 environment, entry sequenced data set.

ESMT. See "external subsystem module table" on page 40.

EUC. See "Extended UNIX Code (EUC) encoding scheme" on page 39.

event analyzer. A database object that provides information about the database events that have taken place. An event analyzer is used with the event monitor file to assess and record performance information.

event monitor. A database object for monitoring and collecting data on database activities over a period of time. For example, starting the database might be an event that causes an event monitor to track the number of users on the system by taking an hourly snapshot of authorization IDs using the database.

event timing. The most precise method of controlling when to start a replication subscription cycle. To use event timing, you must specify an event name and the time that you want the event to be processed. See also "interval timing" on page 53.

exception table. (1) A user-created table that reflects the definition of the table being loaded. (2) A table that holds rows that violate referential constraints or check constraints that the CHECK DATA utility finds.

exclusive lock. A lock that prevents running executing application processes from accessing data. See also "shared lock" on page 88.

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 to perform specific functions.

expanding conversion. A process that occurs when the length of the converted string is greater than that of the source string. See also "contracting conversion" on page 19.

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.

explain snapshot. (1) A collection of information that is compressed when an SQL statement is explained. (2) A capture of the current internal representation of an SQL query and related information. This information is required by the Visual Explain tool.

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 referred to in an SQL statement at the time that the statement was explained.

explicit hierarchical locking. In DB2 for z/OS and 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. See “privilege” on page 75. See also “implicit privilege” on page 49.

export. (1) To copy data from database tables to a file using formats such as PC/IXF, DEL, WSF, or ASC. See also “import” on page 49. (2) In the Information Catalog Center, to populate a tag language file with information catalog contents for use with another program.

export utility. A transaction utility that extracts data from a table. See also “import utility” on page 49 and “load utility” on page 57.

exposed name. A correlation name, a table, or a view name that is 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 binary-coded decimal interchange code (EBCDIC). See “EBCDIC” on page 37.

extended recovery facility (XRF). In a z/OS or 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) encoding scheme. An encoding scheme that defines a set of encoding rules that can support one to four character sets. The encoding rules are based on the ISO2022 definition for the encoding of 7-bit and 8-bit data. The EUC encoding scheme uses control characters to identify some of the character sets.

Extensible Markup Language (XML). A text-based tag language that is used for document processing and for publishing information on the Web.

extensible stylesheet language (XSL). A language for expressing style sheets. It is a file that describes how to display an XML document of a given type. XSL consists of two parts: a language for transforming XML documents, and an XML vocabulary for specifying formatting semantics.

extensible stylesheet language transformation (XSLT). XSLT is used as a general purpose XML processing language. It is widely used for purposes other than XSL, such as generating HTML Web pages from XML data.

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 that is 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 consistent-change-data 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 identified by the SOURCE_OWNER and SOURCE_TABLE columns. See “consistent-change-data table” on page 18. See also “internal CCD table” on page 53.

external function. 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. See also “sourced function” on page 90, “built-in function” on page 9, and “SQL function” on page 91.

external name. The name of an executable file for a stored procedure or user-defined function that is written in a host programming language.

external procedure. An application program that is written in a host language, possibly containing SQL statements, that can be started with the SQL CALL statement. See also “SQL procedure” on page 91.

external routine. A function, method, or procedure written in a host language and that possibly contains SQL statements.

external subsystem module table (ESMT). In the OS/390 environment, a table that specifies the name of the external subsystem module table, which specifies which attachment modules must be loaded by Information Management System.

extract control file. A file that contains statements that control the operation of an extractor utility program.

extract program. In the Information Catalog Center, a utility program that copies metadata from a metadata source (such as an *RDBMS catalog*), translates the metadata into tag language, and places this output into a tag language file.

F

fact table. (1) In DB2 OLAP Server, a table, or in many cases a set of tables, that contains all data values for a relational cube. (2) A relational table that contains facts, such as units sold or cost of goods, and foreign keys that link the fact table to each dimension table.

FAT. File allocation table. A table that is used to allocate space on a disk for a file and to locate the file.

failed member state. In DB2 for z/OS and 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. (1) The process by which a database server, after failure causes it to run on another computer, returns automatically to run on the original computer when it becomes available. (2) The process of returning to a previous release of DB2 for z/OS and OS/390 after attempting or completing migration to a current release.

false global lock contention. In DB2 for z/OS and 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 interpartition communication support.

federated database. In a federated system, the database that is within the federated server. Users and applications interface with the federated database. To these clients, the data sources appear as a single collective database in DB2.

federated system. A special type of distributed database management system (DBMS). A federated system allows you to query and manipulate data located on other servers. The data can be in database managers such as Oracle, Sybase, Informix, and Microsoft SQL Server, or it can be in lists or stores such as a spreadsheet, Web site, or data mart.

A federated system consists of a DB2 instance that will operate as a server, a database that will serve as the federated database, one or more data sources, and clients (users and applications) who will access the database and data sources.

federated server. The DB2 server in a federated system. Any number of DB2 instances can be configured to function as federated servers. You can use existing DB2 instances as your federated server, or you can create new ones specifically for the federated system.

fenced. Pertaining to a type of user-defined function or stored procedure that is defined to protect the database manager from modifications by the function. The database manager is isolated from the function or stored procedure by a barrier. See also “not-fenced” on page 66.

fetch. An SQL action that positions a cursor on the next row of its result table and assigns the values of that row to host variables.

fetch orientation. The specification of the desired placement of the cursor as part of a FETCH statement (for example BEFORE, AFTER, NEXT, PRIOR, CURRENT, FIRST, LAST, ABSOLUTE, and RELATIVE). See also “scrollability” on page 86.

fetch sensitivity. The specification that a FETCH statement has visibility to all changes made by this cursor, as well as changes made by other cursors, or other application processes. Fetch sensitivity results in always fetching the rows from the base table of the SELECT statement of the cursor.

field procedure. In DB2 for z/OS and 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 access token. See “read token” on page 78.

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 System Migrator (FSM). The virtual file system whose space usage is controlled by the Tivoli Space Manager. DB2 Data Links Manager supports the use of this file system in the AIX operating environment.

file update operations. All actions that are involved when a file is changed, especially in the case where the file is referenced in a DATALINK type column and is under the control of a DB2 Data Links Manager. See also "linked file" on page 56.

filter factor. In DB2 for z/OS and OS/390, a number between zero and one that estimates the proportion of rows in a table for which a predicate is true. Those rows are said to qualify by that predicate. Filter factors affect the choice of access paths by estimating the number of rows qualified by a set of predicates.

fixed-length string. A character or graphic string whose length is specified and cannot be changed. See also "variable-length string" on page 107.

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).

foreign key. A column or set of columns that refers to a parent key. In a relational database, a key in one table that references the primary key in another table.

foreign server. In a federated system, another term for data source that is used most often in the context of the SQL/MED standard. See also "data source" on page 27.

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

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

forward-only cursor. See "nonscrollable cursor" on page 66.

forward recovery. A process used to rebuild a restored database or table space to a specified point in time by applying the changes recorded in the database log.

fragmentation. The separation of the index into pieces as a result of inserts and deletions in the index.

free space. The total amount of unused space in a page. The space that is not used to store records or control information is free space.

free space control record (FSCR). A record containing available space approximations for each of the next 500 pages. In multidimensional clustering (MDC) tables, there is one FSCR for each block. It is stored on the first page of that block and covers only the pages in that block.

FSM. See "File System Migrator (FSM)".

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 also "join" on page 55, "outer join" on page 69, "left outer join" on page 56, and "right outer join" on page 84.

full refresh. The process in which all of the data that matches registration and subscription-set predicates for a replication source table is copied to the target table. A full refresh replaces all existing

data in the target table. In a data distribution configuration, a full refresh must be complete before any other data is replicated. See also “change-capture replication” on page 12.

fullselect. A subselect, a values-clause, or a number of both that are combined by set operators. Fullselect specifies a result table. If UNION is not used, the result of the fullselect is the result of the specified subselect.

fully qualified LU name. See “network-qualified name” on page 65.

function. A mapping, embodied as a program (the function body) and invoked by means of zero or more input values (arguments). Functions can be user-defined, built-in, or created by DB2. See also “column function” on page 15, “scalar function” on page 85, “table function” on page 97, “SQL function” on page 91, and “row function” on page 85.

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

function definer. In DB2 for z/OS and OS/390, the authorization identifier 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 for z/OS and OS/390, the authorization identifier of the owner of the function program and function package.

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

function mapping. In a federated system, a mapping between a data source function and an existing DB2 function. DB2 supplies default mappings between existing built-in data source functions and built-in DB2 functions; the default mappings are in the wrapper. The DB2 counterpart function can be either a complete function or a function template. Function mappings are created with the CREATE FUNCTION MAPPING statement. See “function template” on page 44.

function mapping options. In a federated system, parameters of the CREATE FUNCTION MAPPING statement to which you can assign values that pertain to the mapping being created or to the data source function within the mapping. Such values, for example, can include estimated statistics on the overhead that will be consumed when the data source function is invoked. The query optimizer uses these estimates to decide if the function should be invoked by the data source or by DB2, when the data is returned from the data source. See “function mapping”.

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

function package owner. In DB2 for z/OS and OS/390, the authorization identifier 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 that are identified (or used by default) in the user’s function path.

function resolution. The process, internal to the database manager, 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 process of sending the subsections of a request to the specific database partition 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. A DB2 function that you create for the purpose of invoking a function on a data source. A federated server can recognize a data source function only if there is a mapping between the data source function and a counterpart function at the federated database. When no counterpart exists, or when you want to force the federated server to use the data source function, you can create a function template to act as the counterpart. However, unlike a regular function, a function template has no executable code. After you create a function template, you must then create the function mapping between the template and the data source function.

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.

gauge. An indicator for the current value for an item. See also “counter” on page 20.

GBP. See “group buffer pool” on page 46.

GBP-dependent. In DB2 for z/OS and 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 disk.

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

generated column. A column that is derived from an expression that involves one or more columns in the table.

generic resource name. In a 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.

geocoder. In DB2 Spatial Extender, a geocoder is a scalar function that translates existing data (the function’s input) into data that you can understand in spatial terms (the function’s output). Typically, the existing data is relational data that describes a location. For example, a geocoder supplied by Spatial Extender translates United States addresses into instances of a spatial data type. Another geocoder might translate the identifier of a shelf in a warehouse into data that denotes the location of that shelf in the warehouse.

geographic coordinate system. In DB2 Spatial Extender, a reference system that uses latitude and longitude to define the locations of points on the surface of a sphere or spheroid.

geographic information system (GIS). A complex of objects, data, and applications that allows you to generate and analyze spatial information about geographic features, including objects that comprise the earth's surface (for example: rivers, forests, hills, deserts) and objects that occupy it (for example: cities, residences, office buildings, landmarks).

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

ghost index. An invisible index within the existing index object, created during the index create. It is not visible to users until it is fully created. See also "shadow index" on page 88.

GIMSMP. In a z/OS and 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.

GIS. See "geographic information system ".

global catalog. In a federated system, the database system catalog. The catalog contains information about objects in the federated database and information about objects at the data source. The catalog also contains information about the entire federated system.The information in the global catalog is used by the DB2 query optimizer to plan the best way to process SQL statements.

global lock. In DB2 for z/OS and 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 for z/OS and OS/390 members of a data sharing group when those members are trying to serialize shared resources.

global optimizer. In a federated system, a feature of the DB2 SQL Compiler that analyzes the distributed queries and determines the most efficient way to run the query. The global optimizer evaluates queries based on resource cost. See "push-down processing" on page 76.

global record. The row in the register table that defines global replication characteristics for a particular instance of the Capture program.

global table lock. A table lock that is acquired on all partitions of a table's database partition group.

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

governor. See "resource limit facility" on page 83.

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

graphic character. A DBCS character.

graphic string. A sequence of DBCS characters.

gross lock. In DB2 for z/OS and 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 a satellite environment, a collection of satellites that share characteristics such as database configuration and the application that runs on the satellite.

group buffer pool (GBP). A coupling facility cache structure that is used by a data sharing group to cache data and to ensure that the data is consistent for all members. See also “cache structure” on page 10.

group buffer pool duplexing. In a 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. DB2 for OS/390 Publications refer to these instances as the ″old″ (for primary) and ″new″ (for secondary) structures.

grouping task. A task in the Task Center that contains other tasks. You use a task to define task actions that depend on the results of the tasks that the grouping task contains.

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

group restart. In a 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” on page 15.

GTF. See “generalized trace facility” on page 44.

GWAPI. Domino *Go* Web server API.

H

HACMP. See “High Availability Cluster Multiprocessor (HACMP)” on page 47.

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.

health. The general condition or state of the database environment.

Health Center. The DB2 graphical interface that shows the overall state of the database environment and all current alerts. From the Health Center, you can get details about alerts and recommended resolution actions.

health indicator. A measure of some aspect of the health of an object. Criteria are applied to the measurement to determine healthiness, where the criteria applied depends on the type of health indicator as follows:

- Threshold-based: The measurement represents a statistic of the behavior of the object. Warning and alarm threshold values set boundaries on the value of the statistic to define normal, warning, and alarm ranges.

- State-based: The measurement represents two or more states, one of which is normal; and all others are considered non-normal.

health monitor. An instance-level monitor that creates alerts based on a health indicator exceeding a threshold or being in non-normal state. The monitor sends notifications to the notification log, and also sends emails and pages to contacts on its notification list.

health monitor alert. An alert, generated by the Health Monitor, that is based on the type of health indicator:

- Threshold-based: The health indicator value exceeds or falls below warning or alarm thresholds
- State-based: The health indicator value is a non-normal state.

health snapshot. Health data, retrieved from the database manager at a point in time, for a set of database objects

heterogeneous replication. Replication between DB2 and non-DB2 relational databases. See also "federated system" on page 41.

hierarchical relationship category. In the Information Catalog Center, a category of relationship types that are used to connect objects that have a hierarchical relationship.

High Availability Cluster Multiprocessor (HACMP). Any hardware environment with multiple processor nodes that supports the takeover of operations on one processor by another. In a DB2 Data Links Manager configuration, both the DB2 host server and the Data Links server can be configured in an HACMP environment.

hiperspace. In a OS/390 environment, a storage space that contains up to 2 gigabytes of contiguous virtual storage addresses that a program can use as a data 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.

hole. A row for a SELECT statement of a cursor that no longer has a corresponding row in the base table because the row was deleted or updated. A hole is created when a row in the base table no longer qualifies to be in the result set while a cursor is open whose SELECT statement result contains the row that no longer qualifies. Such a row is no longer accessible though the cursor. See also "delete hole" on page 31 and "update hole" on page 105.

home address space. In a 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 that is 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.

hot-spot update. A series of repeated updates made to the same rows over a short period of time.

HTML. See "Hypertext Markup Language".

Hypertext Markup Language (HTML). A markup language that uses tags to specify the format of a document on the Web. These tags define the page layout, graphics, and hypertext links within the document and to other documents on the Internet.

I

ICAPI. Internet Connection API.

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

IDCAMS. In a 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 a 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 for z/OS and OS/390) issues through the MVS subsystem interface to inform DB2 for z/OS and OS/390 of its existence and to initiate the process of becoming connected to DB2.

identity column. A column that provides a way for DB2 to automatically generate a numeric value for each row that is inserted into the table. Identity columns are defined with the AS IDENTITY clause. A table can have no more than one identity column.

idle agent. A database agent that currently does not have a database connection or an application attachment.

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

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

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

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

ILU. See “independent logical unit” on page 50.

image copy. An exact reproduction of all or part of a table space. DB2 for z/OS and 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. (1) 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. (2) A privilege that is granted to a user who has the privilege to execute a package on data objects used within the package that do not require granted explicit privileges. See also “privilege” on page 75 and “explicit privilege” on page 39.

import. (1) To copy data from an external file, using formats such as PC/IXF, DEL, WSF or ASC, into database tables. See also “export” on page 39. (2) In the Information Catalog Center, to read the contents of a tag language file to initially populate the information catalog, change the information catalog contents, or copy the contents of another information catalog.

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

import utility. A transactional utility that inserts user-supplied record data into a table. See also “load utility” on page 57 and “export utility” on page 39.

IMS. Information Management System.

IMS attachment facility. A DB2 for z/OS and OS/390 subcomponent that lets users access DB2 for z/OS and OS/390 from IMS. The IMS attachment facility receives and interprets requests for access to DB2 databases by using exits that are provided by IMS subsystems.

IMS DataPropagator. A product that provides replication between IMS and DB2 databases. See also “replication” on page 82.

IMS DB. Information Management System Database Manager.

IMS TM. Information Management System Transaction Manager.

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

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

incremental backup. A copy of all database data that has changed since the most recent successful full backup operation. This is also known as a cumulative backup image, because a series of incremental backups taken over time will each have the contents of the previous incremental backup image. The predecessor of an incremental backup image is always the most recent successful full backup of the same object.

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 “bind” on page 8.

independent. In DB2 for z/OS and 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. See also "dependent logical unit" on page 32 and "system services control point" on page 96.

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. When you request an index, the database manager builds the structure and maintains it automatically. The index is used by the database manager to improve performance and ensure uniqueness.

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 that is 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 database partition. An index defined on a table is implemented by multiple index partitions, one per table partition.

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

index specification. In a federated system, a set of metadata about a data source object index .The query optimizer uses this information to expedite the processing of distributed requests.When a nickname is created for a data source object, the federated server gathers index information about that object and stores the information in the global catalog. If the object has no index, you can create an index specification for that object to tell the optimizer which column or columns in the object to search on to find data quickly. Use the CREATE INDEX statement to supply index specification information to the global catalog.

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

indicator variable. A variable that represents 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.

individual privilege. A privilege that is granted on a single data object. See also "privilege" on page 75.

indoubt. The status of a unit of recovery. If the database manager 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 the database manager lacks the information that it needs to make this decision, the status of the unit of recovery is *indoubt* until the database manager obtains this information from the coordinator. More than one unit of recovery can be indoubt at restart.

indoubt resolution. 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.

infix operator. An operator that is used in comparison expressions. See also “comparison operator” on page 16.

inflight. A status of a unit of recovery. If DB2 for z/OS and 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*.

informational configuration parameter. A type of configuration parameter that holds information that cannot be modified. See also “configurable configuration parameters” on page 17 and “configurable online configuration parameters” on page 17.

information catalog. A collection of metadata, managed by the Information Catalog Center, that contains descriptive data (business metadata) that helps users identify and locate data and information that is available to them in the organization. An information catalog also contains some technical metadata.

Information Catalog Center. A DB2 graphical interface for organizing, maintaining, finding, and using business information. The Information Catalog Center is a part of the Information Catalog Manager.

Information Catalog Manager application program interface (API). A set of Java classes that can be used to write programs that read, create, and update the metadata that is stored in the information catalog.

Information Catalog Manager. A set of tools designed to help organize, maintain, find, and use business information. The Information Catalog Manager is comprised of the Information Catalog Center, the Manage Information Catalog wizard, and the Information Catalog Manager sample catalogs. A Web version of the Information Catalog Center is also available.

information source. An item of data or information, such as a table or chart, that is represented by an Information Catalog Center object.

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. See “join” on page 55. See also “outer join” on page 69.

I/O parallelism. See “parallelism” on page 70.

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. See also “invalid package” on page 54.

inoperative trigger. A trigger that depends on an object that has been dropped or made inoperative or on a privilege that has been revoked. See also “trigger” on page 102.

inoperative view. (1) A view that is no longer usable because the SELECT privilege on a table or view that the view is dependent on is revoked from the definer of the view. (2) An object on which the view definition is dependent was dropped (or possibly made inoperative in the case of another view).

input relationship type. In the Information Catalog Center, a relationship type that is used to connect objects that transform to their input data resource. See “transformation relationship category” on page 101. See also “relationship type” on page 80.

insensitive cursor. A cursor that is not sensitive to inserts, updates, or deletes made to the database (rows underlying the result table) after the result table is materialized (a row of the result table is materialized when the values for the row are captured from the database). Consequently, the size of the result table, the order of the rows, and the values for each row do not change after the cursor is opened. The SELECT statement cannot contain a FOR UPDATE clause, and the cursor cannot be used for positioned updates or deletes. A positioned UPDATE or DELETE using an insensitive scrollable cursor results in an error. See also “sensitive cursor” on page 87.

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

insert trigger. A trigger that is defined with the triggering SQL operation INSERT.

isolation level. (1) A security feature that determines how data is locked from other processes while it is being accessed. See also repeatable read, read stability, cursor stability, and uncommitted read. (2) An attribute that defines the degree to which an application process is isolated from other concurrently executing application processes.

installation program. A program that prepares a software package to run on the computer. During installation, a component of the setup program is commonly copied to and the disk drive to allow the user to customize the program’s default settings.

installation verification scenario. A sequence of operations that exercises the main DB2 Universal Database functions and tests whether DB2 was correctly installed.

instance. (1) See also “database manager instance” on page 24. (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.

instance-owning partition. The first database-partition server that is installed in a partitioned database environment.

instantiable structured type. A structured type that can be used for creating database objects. A structured type that is not instantiable cannot be used for creating database objects; however, such a type can be used to define subtypes which, in turn, can be instantiable.

instrumentation facility component identifier (IFCID). In DB2 for z/OS and 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 for z/OS and OS/390, to submit DB2 for z/OS and OS/390 commands, and to pass data to DB2 for z/OS and OS/390.

interactive SQL. A set of SQL statements that is provided through an interface such as the Command Center or command line processor. These statements are processed as dynamic SQL statements. For example, an interactive SELECT statement can be processed dynamically using the DECLARE CURSOR, PREPARE, DESCRIBE, OPEN, FETCH, and CLOSE statements.

Interactive System Productivity Facility (ISPF). In a OS/390 environment, an IBM licensed program that provides interactive dialog services. Users can perform most DB2 tasks interactively through ISPF panels.

inter-DB2 R/W interest. In DB2 for z/OS and 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 database server. The target of a request from a local application or a remote application requester that is being forwarded to another database server because the object doesn't exist on the target database server. The remote request is forwarded transparently to another database server if the object referred to by the three-part name does not refer to the local location. See also "database server" on page 24.

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 because it is not a registered replication source. It does not have its own row in the register table; it is identified by the CCD_OWNER and CCD_TABLE columns for the row of the associated registered replication source. See also "external CCD table" on page 40.

internal resource lock manager (IRLM). A DB2 for z/OS and OS/390 component that allows serial access to data. DB2 requests locks from IRLM to ensure data integrity when applications, utilities, and commands attempt to access the same data.

internationalization. In the z/OS or OS/390 environment, the support for an encoding scheme that is able to represent the code points of characters from many different geographies and languages. See also "Unicode" on page 104.

Internet Protocol (IP). A protocol that is used to route data from its source to its destination in an Internet environment. See also "Transmission Control Protocol/Internet Protocol" on page 101.

inter-partition parallelism. A single database operation (for example index creation) that is executed in parallel across the partitions of a partitioned database. See also "intra-partition parallelism".

Inter-Process Communication (IPC). A mechanism of an operating system that allows processes to communicate with each other within the same computer or over a network.

inter-query parallelism. The ability of multiple applications to query a database at the same time. Each query executes independently of the others, but DB2 runs all of them at the same time. See also "intra-query parallelism" on page 54.

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. See also "event timing" on page 38.

intra-partition parallelism. The subdivision of a single database operation (for example, index creation) into multiple parts, which are then executed in parallel within a single database partition. See also "inter-partition parallelism".

intra-query parallelism. The ability to process parts of a single query at the same time using either "intra-partition parallelism" on page 53, "inter-partition parallelism" on page 53, or both.

invalid package. A package that depends on an object that has been dropped. See also "inoperative package" on page 51.

invariant character set. (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. See also "syntatic character set" on page 95.

IP. See "Internet Protocol" on page 53.

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

IRLM. See "internal resource lock manager" on page 53.

ISAPI. Microsoft® Internet Server API.

ISPF. See "Interactive System Productivity Facility" on page 53.

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

J

Java archive. A file format that is used for aggregating many files into one. Commonly known as a *JAR* file.

Java Database Connectivity (JDBC). A set of database APIs for use in the Java programming language. APIs that allows access to database management systems from Java applications using callable SQL, which does not require the use of an SQL preprocessor. The JDBC architecture allows users to add modules, called JDBC database drivers, that link the Java 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.

JCL. See "job control language".

JDBC driver. A program included with database management systems to support the JDBC standard access between the databases and Java applications.

JES. See "Job Entry Subsystem".

JFS. See "Journaled File System (JFS)" on page 55.

job control language (JCL). A command 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 that is 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. See also “right outer join” on page 84, “left outer join” on page 56, “outer join” on page 69, “full outer join” on page 42 and “inner join” on page 51.

joined table. An intermediate result table that is the result of either an inner join or an outer join.

journal. (1) For iSeries systems, a system object that identifies the objects being journaled, the current journal receiver, and all the journal receivers on the system for the journal. The system-recognized identifier for the object type is *JRN. See also “journal receiver”. (2) The destination pages from which you can view all available historical information about task history, database history, PM alerts, messages, and the notification log.

Journaled File System (JFS). The native file system in the AIX operating system.

journal receiver. For iSeries systems, a system object that contains journal entries added when events occur that are journaled, such as changes to a database file, changes to other journaled objects, or security-relevant events. The object type is *JRNRCV. See also “journal”.

K

Kerberos. A network authentication protocol that is designed to provide strong authentication for client/server applications by using secret-key cryptography. See also “Kerberos ticket”.

Kerberos ticket. A transparent application mechanism that transmits the identity of an initiating principal to its target. A simple ticket contains the principal’s identity, a session key, a timestamp, and other information, which is sealed using the target’s secret key.

key. A column or an ordered collection of columns that is identified in the description of a table, index, or referential constraint. The same column can be part of more than one key.

key-sequenced data set (KSDS). In a 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. (3) In the Information Catalog Center, an element of the tag language that identifies the meaning of a data value that is imported into an information catalog.

KSDS. See “key-sequenced data set”.

L

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

large object (LOB). A sequence of bytes with a size ranging from 0 bytes to 2 gigabytes less 1 byte. It can be any of three types: binary large object (binary), character large object (single-byte character or mixed), or double-byte character large object (double-byte character).

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

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

latency. The time required for updates made to a replication source to appear in a replication target.

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

LDS. See "linear data set".

leaf page. A page that contains pairs of keys and record identifiers and that points to actual data. See also "nonleaf page" on page 66.

left outer join. 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" on page 55. See also "right outer join" on page 84 and "full outer join" on page 42.

LEN node. See "low-entry networking node" on page 60.

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

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.

link. The action that the DB2 Data Links Manager takes to control a file that is referenced in a table that contains a DATALINK column. A file can be linked as the result of such database actions as an SQL UPDATE, INSERT, IMPORT, or LOAD.

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.

linked file. In DB2 Data Links Manager, a file that is referenced in a table's DATALINK column that is defined with LINK CONTROL. To guarantee referential integrity, a linked file is maintained under the control of the DLFF component.

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

linked relationship type. In the Information Catalog Center, a relationship type that is used to connect two or more objects in an information catalog. Objects in a linked relationship are peers, rather than a parent-child relationship.

For example, in the sample information catalog that is included with the Information Catalog Center, the object called **CelDial Sales Information** is linked with objects that describe CelDial advertisements for the year. See also "relationship type" on page 80.

list. A type of object, which DB2 utilities can process, that identifies multiple table spaces, multiple index spaces, or both. A list is defined with the LISTDEF utility control statement.

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

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

load authority. An access level that gives LOAD utility or AutoLoader utility privileges to load data into tables.

load copy. A backup image of data that was previously loaded and that can be restored during rollfoward recovery.

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

load utility. A nontransactional utility that performs block updates of table data. See also “import utility” on page 49 and “export utility” on page 39.

LOB. See “large object” on page 55.

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

LOB lock. In DB2 for z/OS and OS/390, a lock on a LOB value.

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

local. Pertaining to any object that the local subsystem maintains. In DB2 for z/OS and OS/390, for example, a local table is a table that is maintained by the local DB2 subsystem. See also “remote” on page 81.

local database. A database that is physically located on the workstation in use. See also “remote database” on page 81.

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 for z/OS and OS/390, the definition of a subset of a user’s environment that combines characters that are defined for a specific language and country or region, and a CCSID. A collection of processing variables that are used to specify how a process executes. Computer locales include the conventions for a specific language and culture, with appropriate date and time formatting, character classification, sorting, and text handling.

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

local subsystem. The unique relational database management system to which the user or application program is directly connected (in the case of DB2 for z/OS and OS/390, by one of the DB2 for z/OS and 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. (1) The name by which DB2 for z/OS and OS/390 refers to a particular DB2 subsystem in a network of subsystems. (2) The unique name of a database server. An application uses the location name to access a DB2 database server.

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" on page 57.

locator variable. A host variable that contains the locator that represents a LOB value on the application server.

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.

lock duration. The interval over which a DB2 for z/OS and OS/390 lock is held. For example, locks on LOBs are taken when they are needed and are usually released at commit.

lock escalation. 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 for z/OS and OS/390 lock is holding.

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

lock parent. For explicit hierarchical locking in DB2 for z/OS and 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 for z/OS and OS/390 lock to a higher level.

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

lock structure. In DB2 for z/OS and 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 that records changes made in a system. (2) A collection of records that describe the events that occur during DB2 for z/OS and OS/390 execution and that indicate their sequence. The recorded information is used for recovery in the event of a failure during DB2 for z/OS and OS/390 execution. (3) See also "database log" on page 23.

log file. (1) A record that is used to monitor a database's activity. Log files are essential to the backup and recovery process. (2) A file that is produced by the Information Catalog Center when it imports a tag language file into the information catalog. This file records the times and dates when the import process started and stopped and any error information for the process.

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

logical agent. An agent that represents the client or application connection.

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

logical data group. A collection of data elements that gather database system monitoring information at a specific scope of database activity. The snapshot monitor and event monitor each have their own sets of logical data groups. See "data element" on page 25.

logical data modeling. The process of documenting the comprehensive business information requirements in an accurate and consistent format. Data modeling is the first step in designing a database.

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

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

logical lock (L-lock). In DB2 for z/OS and OS/390, the lock type that transactions use to control intra-DB2 and inter-DB2 data concurrency between transactions. See also "physical lock" on page 72.

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 for z/OS and 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, even though the actual media (coupling facility or DASD) might not contain any errors. Usually, a connection to the media has been lost.

logical partition. (1) In DB2 for z/OS and OS/390, a set of key or RID pairs in a nonpartitioning index that are associated with a particular partition. (2) In a partitioned database environment, a database partition server on a processor that has more than one database partition server assigned to it

logical recovery pending (LRECP). In DB2 for z/OS and 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 a OS/390 environment, an access point through which an application program accesses the SNA network to communicate with another application program. See also “LU name” on page 61.

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 a 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 for z/OS and OS/390 attempts to locate the current end of the log.

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 header (LRH). In DB2 for z/OS and OS/390, a prefix for a logical record that contains control information. Only the first segment contains the entire LRH; later segments include only the first two fields. When a specific log is needed for recovery, all segments are returned and presented together as if the record were stored continuously.

log record sequence number (LRSN). A number that DB2 for z/OS and OS/390 generates and associates with each log record. The LRSN is also used for page versioning. The LRSNs that a particular DB2 for z/OS and 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 for z/OS and OS/390, a process by which an explicit starting relative byte address is established. This RBA is the point at which the next byte of log data is to be written.

long string. (1) A variable-length string whose maximum length is greater than 254 bytes. (2) In DB2 for z/OS and OS/390, a string whose actual length, or a variable-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. See also “short string” on page 89.

long table space. See “large table space” on page 56.

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 “local page list” on page 59.

LRECP. See “logical recovery pending” on page 59.

LRH. See “log record header”.

LRSN. See “log record sequence number” on page 60.

LU. See “logical unit” on page 59.

LU name. In a OS/390 environment, the name by which VTAM refers to a node in a network. See also “location name” on page 58.

LU 6.2. See “logical unit 6.2” on page 60.

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 and the usage of data stream controls, function management headers, request unit parameters, sense data values, and presentation services protocols such as those associated with function management headers.

LUW. See “logical unit of work” on page 60.

LUWID. See “logical unit of work identifier” on page 60.

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. See also “basic conversation” on page 8.

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.

mass delete. The deletion of all rows of a table.

massively parallel processing (MPP). More than one uniprocessor or symmetric multiprocessor (SMP) computer linked together by a high-speed network.

master table. In update-anywhere replication, the original source table for data in the replica table. If replication conflict detection is enabled, changes made to the master table are retained, whereas changes made to the replica table are rejected. See also “update-anywhere replication” on page 105, “replica table” on page 81, and “conflict detection” on page 17.

materialize. (1) In DB2 for z/OS and OS/390, to put rows from a view or nested table expression into a work file for additional processing by a query. (2) To place a LOB value into contiguous storage. Because LOB values can be very large, DB2 for z/OS and OS/390 avoids materializing LOB data until doing so becomes absolutely necessary.

materialized query table. A table whose definition is based on the result of a query and whose data is in the form of precomputed results that are taken from the table, or tables, that the materialized query table definition is based on.

MBCS. See “multibyte character set” on page 63.

member. See “subscription set member” on page 94.

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

member scope. See “command” on page 15.

member state. In DB2 for z/OS and OS/390, the state of the DB2 member (subsystem) of the data sharing group.

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

merge. To update and insert new content into a table.

message processing program (MPP). An IMS online program that can access DB2 for z/OS and OS/390 databases, full-function databases, data entry databases, and main storage databases.

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. Database catalogs and information catalogs contain metadata.

metadata publication process. A process created by the Data Warehouse Center that contains all the steps to keep 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.

mixed-character string. A string that contains a mixture of single-byte and multibyte characters. Synonym for *mixed-data string*.

mobile client. The node, usually a mobile 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.

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 a OS/390 environment, a VTAM name for the collection of physical and logical characteristics and attributes of a session.

MODEENT. In a 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 does not need to currently exist in the database.

modeling database. In the OS/390 environment, a DB2 database that you create on your workstation which models a DB2 subsystem in the OS/390 environment. A modeling database can be used for indexing and query optimization.

modify lock. In DB2 for z/OS and 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.

Monitor control server. A database that contains the replication Monitor control tables, which store information about alert conditions that the Replication Alert Monitor will monitor.

Monitor qualifier. A case-sensitive character string that identifies an instance of a Replication Alert Monitor process.

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

monotonic decreasing expression. An expression or function used to derive a generated column that has the property that for every possible pair of values x1 and x2, if x2>x1 then fn(x2)<fn(x1).

monotonic increasing expression. An expression or function used to derive a generated column that has the property that for every possible pair of values x1 and x2, if x2>x1 then fn(x2)>fn(x1).

monotonic nondecreasing expression. An expression or function used to derive a generated column that has the property that for every possible pair of values x1 and x2, if x2>x1 then fn(x2)>=fn(x1).

monotonic nonincreasing expression. An expression or function used to derive a generated column that has the property that for every possible pair of values x1 and x2, if x2>x1 then fn(x2)<=fn(x1).

MPP. See "message processing program" on page 62 or "massively parallel processing" on page 61.

MQSeries. An IBM product that facilitates asynchronous message queuing between applications.

MQT. See "materialized query table" on page 61.

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

multibyte character set (MBCS). A set of characters in which each character is represented by 1 or more bytes. Contrast with "double-byte character set" on page 36 and "single-byte character set" on page 89. See also "ASCII" on page 5, "single-byte character set" on page 89, "EBCDIC" on page 37, and "Unicode" on page 104.

multidimensional. In the DB2 OLAP Server, 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. See also "business dimension" on page 9 and "dimension" on page 33.

multidimensional analysis. The process of assessing and evaluating an enterprise on more than one level.

multidimensional database. In the DB2 OLAP Server, a nonrelational database into which you copy relational data for OLAP analysis. See also “relational cube” on page 80.

multidimensional clustering (MDC) table. A table whose data is physically clustered on more than one table column simultaneously.

multiple logical partition configuration. In a partitioned database environment, a configuration in which more than one database partition server is assigned to a computer, and these database partition servers are recorded in the same db2nodes.cfg file.

multisite update. 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.

multi-tier replication. A replication configuration in which changes are replicated from a replication source in one database to a replication target in another database, and changes from this replication target are replicated again to a replication target in another database.

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

MVS. (Multiple Virtual Storage) The primary operating system used on IBM mainframe computers. This operating system manages large amounts of memory and disk space.

MVS/ESA. Multiple Virtual Storage/Enterprise Systems Architecture. Renamed, and more commonly known as z/OS.

N

NAU. See “network addressable unit”.

negotiable lock. In DB2 for z/OS and 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. A fullselect in a FROM clause (surrounded by parentheses).

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 identifier (NID). In a OS/390 environment, the network identifier that is assigned by IMS or CICS, or if the connection type is RRSAF, the OS/390 RRS unit of recovery identifier (URID).

Network Information Service (NIS/NIS+). On AIX, a central record of passwords, nodes, and so on that can be used with the DB2 Administration Server in the administration of user and group names.

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

NETWORK netid. The identifier of the SNA network where the remote LU resides. This network ID is a string of one to eight characters that follows the naming convention for SNA.

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. See also "Advanced Peer-to-Peer Networking" on page 2.

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) In a federated system, identifiers used to reference the object located at the data sources that you want to access. The objects that nicknames identify are referred to as *data source objects*. Examples of data source objects include tables, views, synonyms, table-structured files, and search algorithms. (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-DB2 relational database.

NID. See "network identifier" on page 64.

NIS/NIS+. See "Network Information Service" on page 64.

NN. See "network node".

node. (1) 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. (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 nodes connected by a high-speed network. (3) An obsolete term for database partition. See "database partition" on page 24.

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

nodegroup. An obsolete term for database partition group. See "database partition group" on page 24.

noncomplete CCD table. A CCD table that is initially empty and has rows appended to it as changes are made to the replication source. See also "complete CCD table" on page 16.

noncondensed attribute. A table attribute that indicates 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. A CCD table that can contain more than one row for each key value. These duplicate rows represent the history of changes for the values in rows of a table. Contrast with "condensed CCD table" on page 17. See also "consistent-change-data (CCD) table" on page 18.

noncumulative backup image. See "delta backup" on page 31.

non-DB2 relational database server. An Informix database server or a relational database server from a vendor other than IBM.

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

nonleaf page. 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. See also "leaf page" on page 56.

nonpartitioning index. In DB2 for z/OS and OS/390, any index that is not a partitioning index. For example, if you define a nonpartitioning index and a partitioning index on the same table, you lose some of the benefits of partition-level independence for utility operations, because access to a nonpartitioning index must be sequential. See also "partitioning index" on page 71.

nonscrollable cursor. A cursor that can be moved only in a forward direction. Nonscrollable cursors are sometimes called forward-only cursors or serial cursors. See also "scrollable cursor" on page 86.

normalization. The process of restructuring a data model by reducing its relations to their simplest forms. It is a key step in the task of building a logical relational database design. Normalization helps avoid redundancies and inconsistencies in your data. An entity is normalized if it meets a set of constraints for a particular normal form (first normal form, second normal form, and so on). See also "denormalization" on page 32 and "repeating group" on page 81.

NOT DETERMINISTIC function. In DB2 for z/OS and 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" on page 107. Contrast with "DETERMINISTIC function" on page 32.

not-fenced. Pertaining to a type of user-defined function or stored procedure that is defined to be run in the database manager process. There is no protection for the database manager from changes by this function. See also "fenced" on page 41.

notification log. See "administration notification log" on page 2.

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. Synonym for "DETERMINISTIC function" on page 32. See also "VARIANT function" on page 107.

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

NTFS. One of the native file systems in the Windows NT and later operating environments such as Windows 2000.

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

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

NULL terminator. In C language, the value that indicates the end of a string. For character strings, the NULL terminator is X'00'.

null. A value that indicates the absence of information.

null indicator. A column (by byte position) in a nondelimited ASCII file that contains the "null indicator flag" for the data being loaded into a table column. The null indicator can be any valid positive integer.

null indicator flag. A one-byte character that is contained in a "null indicator" column of a nondelimited ASCII file. When the load process looks at each data row, the null indicator flag indicates whether the data in the column that is defined by the start and end positions is null.

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.

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

O

OASN. See "origin application schedule number" on page 69.

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 that consists of data and operations associated with that data. (3) In the Information Catalog Center, an item that represents a unit or distinct grouping of information. Each Information Catalog Center object identifies and describes information but does not contain the actual information. For example, an object can provide the name of a report, list its creation date, and describe its purpose.

object type. (1) A categorization or grouping of object instances that share similar behaviors and characteristics. (2) In the Information Catalog Center, a classification for objects. An object type is used to reflect a type of business information, such as a table, report, or image.

OBID. In DB2 for z/OS and OS/390, data object identifier.

occasionally connected. A replication configuration that contains target servers that are not always connected to the network. This configuration allows users to connect to a primary data source for a short time to synchronize their local database with the data at the source.

ODBC. See "Open Database Connectivity" on page 68.

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. During an offline backup, the backup database utility acquires exclusive use of the database until the backup is complete. See also "online backup" on page 68.

offline restore. A restoration of a copy of a database or table space from a backup. The restore database utility has exclusive use of the database until the restore is completed. See also “online restore”.

OLAP. See “online analytical processing”.

old structure. See “primary group buffer pool” on page 74.

online analytical processing (OLAP). In the DB2 OLAP Server, a multidimensional, multi-user, client server computing environment for users who need to analyze consolidated enterprise data in real time. OLAP systems feature zooming, data pivoting, complex calculations, trend analysis, and modeling.

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. See also “offline backup” on page 67.

online index create. Creating a new index while allowing the underlying table and any previously existing indexes to be read and updated by concurrent transactions.

online index reorganization. Reorganizing indexes on a table while allowing the table and existing indexes to be read and updated by concurrent transactions.

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. See also “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.

operator. An action that must be performed on data, or the output from a table or an index, when the access plan for an SQL statement is executed.

operational data. Data that is used to run the day-to-day operations of an organization.

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.

option. In the Information Catalog Center tag language, a parameter of the ACTION tag that defines the action to be performed on objects or object types in the information catalog when the tag language file is imported.

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 for z/OS and OS/390, an *uppercase* letter followed by zero or more characters, each of which is an *uppercase* letter, a number, a digit, or the underscore character. An ordinary identifier must not be a reserved word. See also “delimited identifier” on page 31.

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

origin application schedule number (OASN). In a 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.

originating task. In DB2 for z/OS and 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. (2) 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 "join" on page 55. See also "inner join" on page 51, "full outer join" on page 42, "left outer join" on page 56, and "right outer join" on page 84.

outline. In the DB2 OLAP Server, the structure that defines all elements of a database within the DB2 OLAP Server. 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.

output relationship type. In the Information Catalog Center, a relationship type that is used to connect objects that transform to their output data resource. See "transformation relationship category" on page 101. See also "relationship type" on page 80.

overflow record. (1) 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. (2) 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. (3) In the event monitor, a record that is 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.

ownership privilege. Control privilege that allows all privileges for the owned data object. See also "privilege" on page 75.

P

package. (1) A control structure produced during program preparation that is used to execute SQL statements. (2) In Java programming, a program statement that defines the location of a Java class within the directory structure, or library, of a Java application.

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

package name. The name of an object that is created by BIND, PRECOMPILE, or REBIND 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) A unit of storage within a table space (4 KB, 8 KB, 16 KB, or 32 KB) or index space (4 KB). In a table space, a page contains one or more rows of a table. In a LOB table space, a LOB value can span more than one page, but no more than one LOB value is stored on a page. (3) In a notebook in the graphical interface, a predefined display image that typically provides fields and controls that help users accomplish tasks.

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

page set recovery pending (PSRCP). In DB2 for z/OS and OS/390, a restrictive state of an index space in which the entire page set must be recovered.

panel. In DB2 for z/OS and 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 a OS/390 environment, a set of consecutive operations that execute in parallel and have the same number of parallel tasks.

parallel I/O processing. (1) A form of I/O processing in which DB2 for z/OS and OS/390 initiates multiple concurrent requests for a single user query and performs I/O processing concurrently (in parallel) on multiple data partitions. (2) The process of reading from or writing to two or more I/O devices at the same time to reduce response time.

parallelism. The ability to perform multiple database operations at the same time in parallel. See also “inter-partition parallelism” on page 53, “intra-partition parallelism” on page 53, and “I/O parallelism” on page 51.

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” on page 88.

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

parallel task. In a OS/390 environment, the execution unit that is dynamically created to process a query in parallel.

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.

parameter-name. A long identifier that names a parameter that can be referenced in a procedure or user-defined function.

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 for z/OS and OS/390, a table space that contains a parent table. See also "dependent table space" on page 32.

partial declustering. In a partitioned database environment, the storage of table data on a named subset of database partitions (database partition group), rather than on all database partitions for the database.

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

partition. In a 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 gigabytes, 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. See also "join" on page 55 and "collocated join" on page 14.

partitioned database. A database with two or more database partitions. Each database partition stores a subset of table data for each table that resides on it. See "database partition" on page 24.

partitioned data set (PDS). In a 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 a OS/390 environment, a partitioned table space or an 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 a 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.

partitioning agent. Within AutoLoader, the process used to create the partition files for loading. This is done by splitting or partitioning the input file.

partitioning index. An index that determines how rows are physically ordered in a partitioned table space. See also "clustered index" on page 14.

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 database partition group.

partitioning map index. A number that is 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 for z/OS and OS/390 subsystem by way of a VTAM conversation.

pass-through. In a federated system, pertaining to a special DB2 session used to submit SQL statements directly to DBMSs using the SQL dialect associated with that data source. Use a pass-through session when you want to perform an operation that is not possible with DB2 SQL/API, or to perform actions not supported by SQL.

path. (1) In an operating system, the route through a file system to a specific file. (2) In a network environment, the route between any two nodes. See also "SQL path" on page 91.

PCT. In CICS, program control table.

PDS. See "partitioned data set" on page 71.

peer table. A replication source or target table defined as part of a peer-to-peer replication configuration.

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

peer to peer relationship category. In the Information Catalog Center, a category of relationship types that are used to connect objects that have a peer relationship.

peer-to-peer replication. A replication configuration in which all peer tables are both registered sources and read-write targets, and there is no primary source table for full refresh. In this configuration, there is no replication hierarchy among the peer tables. Contrast with "update-anywhere replication" on page 105. See also "multi-tier replication" on page 64.

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

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 that is derived from performance data that is obtained from the database manager. The expression for this variable can be user-defined by the user.

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 for z/OS and OS/390, a claim on an entire nonpartitioning index.

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

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

physical lock (P-lock). A lock type that DB2 for z/OS and OS/390 acquires to provide consistency of data that is cached in different DB2 for z/OS and OS/390 subsystems. Physical locks are used only in data sharing environments. See also "logical lock (L-lock)" on page 59.

physical lock contention. In DB2 for z/OS and OS/390, conflicting states of the requesters for a physical lock. See also "negotiable lock" on page 64.

physically complete. In DB2 for z/OS and 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" on page 19.

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

plan. See "application plan" on page 4.

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

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

plan segmentation. In DB2 for z/OS and 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" on page 72.

PLT. In CICS, program list table.

point-in-time table. A type of replication target table whose content matches all or part of a source table, with an added 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 that a program accesses is consistent. The point of consistency occurs when updates, insertions, and deletions are either committed to the physical database or rolled back. See also "rollback" on page 84 and "commit point" on page 15.

policy. See "CFRM policy" on page 12.

populate. To add object types, objects, or metadata to the Information Catalog Center.

possibly uncommitted. A state assigned by the index manager to an index key where the completion of the COMMIT of the insertion or deletion of that key cannot be determined.

postponed abort UR. In DB2 for z/OS and 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.

power user. (1) A person who has special privileges to perform object management tasks, such as creating and updating objects. (2) In the Information Catalog Center, a person who has access to the information available in an information catalog but who is not an administrator. This person can also perform some object management tasks, such as defining objects and updating or deleting objects that this person has already defined. See also "administrator" on page 2 and "user" on page 105.

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

precision. In numeric data types, the total number of binary or decimal digits, excluding the sign. The sign is considered positive if the value of a number is zero.

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.

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

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

prefetch processing. In DB2 for z/OS and OS/390, an operation in which data is read by one of the following mechanisms: sequential prefetch or list sequential prefetch (also referred to as list prefetch).

prefix. In a DB2 Data Links Manager environment, an absolute path in a DLFS under which linked files are stored.

prepare. (1) To convert an SQL statement from text form to an executable form, by submitting it to the SQL compiler. (2) 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 is processed by the PREPARE statement.

primary authorization ID. The authorization identifier that is used to identify the application process to DB2 for z/OS and OS/390.

primary group buffer pool. For a duplexed group buffer pool, the DB2 for z/OS and 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. See also “secondary group buffer pool” on page 86.

primary index. In DB2 for z/OS and 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. It is a column or combination of columns that uniquely identifies a row in a table.

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. See also “secondary log” on page 86.

principal. An entity that can communicate securely with another entity. In Kerberos, principals are represented as entries in the Kerberos registry database and include users, servers, computers, and others.

private connection. A communications connection that is specific to DB2 for z/OS and OS/390. For example, when the application server is a DB2 subsystem, DB2 private connections are allocated as necessary to support references to objects at other DB2 subsystems. Like SQL connections, DB2 private connections are initially in the held state and can be placed in the release pending status.

private protocol access. A method of accessing distributed data by which you can direct a query to another DB2 system. See also “DRDA access” on page 36.

private protocol connection. A DB2 private connection of the application process. For example, if the first phase of your application program uses DB2 private protocol access and the second phase uses

DRDA access, then open DB2 private protocol connections from the first phase might cause a CONNECT operation to fail in the second phase. See also “private connection” on page 74.

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. For example, privileges include rights such as creating, deleting, and selecting data from tables. (2) In the Information Catalog Center, 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 creating, updating, and deleting objects from the information catalog. (3) In DB2 for z/OS and OS/390, the capability of performing a specific function, sometimes on a specific object. See also “explicit privilege” on page 39, “implicit privilege” on page 49, and “authority” on page 6.

privilege set. In the installation SYSADM ID, the set of all possible privileges. For any other authorization identifier, the set of all privileges that are recorded for that identifier in the DB2 for z/OS and OS/390 catalog.

procedure. See “stored procedure” on page 93.

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) The unit to which the database manager allocates resources and locks. A process involves the execution of one or more programs. The execution of an SQL statement is always associated with a process. The means of initiating and terminating a process are dependent on the environment. Synonym for “application process” on page 4.

programs object type. An object type that identifies and describes applications capable of processing the actual information that is described by Information Catalog Center objects.

The programs object type is included with the Information Catalog Center. Administrators specify which programs can be used to access certain object types.

projected coordinate system. In DB2 Spatial Extender, a reference system that defines the locations of points on a planar surface.

promote. To copy replication definitions for subscription sets or registered sources from one database to another database, without having to register the sources again or recreate the subscription sets.

propagation. A process in which groups of configuration parameters are updated and take effect at different rates.

property. (1) 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. (2) In the Information Catalog Center, a characteristic or attribute that describes a unit of information. Each object type has a set of associated properties.

property display name. A 254 character name that is used by the Information Catalog Center to display the name of a property in the Properties window.

property name. The 254-byte descriptive name of a property that is displayed in the Information Catalog Center user interface.

protected conversation. In a 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.

pruning. The task of removing obsolete data from replication control tables, CD tables, CCD tables, or Capture or Apply log files.

PSRCP. See "page set recovery pending" on page 70.

pseudo deleted. A key that is marked as deleted but has not yet been physically removed from the index page.

pseudo delete index. An index in which deletion of keys is not done until sometime after the deletion is committed. See also "type 2 indexes" on page 103.

pseudo empty. An index page is pseudo empty if all the keys on the page are marked as pseudo deleted. See also "pseudo deleted".

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

PU. See "physical unit" on page 73.

pull configuration. A replication configuration in which the Apply program runs at the target server; the Apply program pulls updates from the source server to apply them to the target. Contrast with "push configuration".

push configuration. A replication configuration in which the Apply program runs at the source server or a replication server other than the target server; the Apply program pushes updates from the source server to apply them to the target. Contrast with "pull configuration".

push-down processing. In a federated system, the processing of segments of a query at a data source instead of at the federated server.

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

Q

QBIC. See "Query by Image Content".

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 for z/OS and OS/390, a component of certain SQL statements that specifies a result table.

query block. In DB2 for z/OS and 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 for z/OS and 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 for z/OS and OS/390, parallel execution of a single query, which is accomplished by using multiple tasks. See also "Sysplex query parallelism" on page 96.

query I/O parallelism. In DB2 for z/OS and OS/390, parallel access to data, which is accomplished by triggering multiple I/O requests within a single query.

query optimization class. A set of query rewrite rules and optimization techniques for compiling queries.

query 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.

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.

quiesce point. A point at which data is consistent as a result of running the DB2 QUIESCE utility.

quiesced member state. In DB2 for z/OS and 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 DB2 for z/OS and OS/390 system fails before the command takes effect, the member state is failed.

quoted name. See "delimited identifier" on page 31.

QSAM. See "queued sequential access method ".

R

RACF. See "Resource Access Control Facility" on page 82.

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

RBA. See "relative byte address" on page 81.

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

RDB. See "relational database" on page 80.

RDBMS. See "relational database management system" on page 80.

RDBMS catalog. In the Information Catalog Center, a collection of tables that contains descriptions of SQL objects, such as tables, views, and indexes, maintained by an RDBMS.

RDBNAM. See "relational database name" on page 80.

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

readahead prefetching. A method of prefetching pages by looking ahead in a scan, which results in asynchronous retrieval of pages even though those pages are not located sequentially on disk. See also "sequential prefetch" on page 87 and "list prefetch" on page 57.

read only. A file or project can be read, but not updated or deleted.

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 concurrency than cursor stability. See also "cursor stability" on page 22, "repeatable read " on page 81, and "uncommitted read (UR)" on page 103.

read token. The authorization key embedded in a READ PERMISSION DB DATALINK column value, returned as a simple column value or by using the scalar functions DLURLCOMPLETE or DLURLPATH. A read token is required for reading a file referenced in a READ PERMISSION DB DATALINK column.

real-time replication. See "synchronous replication" on page 95.

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. See also "automatic rebind" on page 7 and "bind" on page 8.

recapture. In update-anywhere replication, to capture changes at a replica table and forward these changes to the master table or to other replica tables.

RECONCILE. A DB2 utility that is used to validate and repair references to files in the DATALINK columns of a table.

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

record identifier (RID). A 3-byte page number followed by a 1-byte slot 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. See also "row identifier" on page 85.

record identifier (RID) pool. In DB2 for z/OS and 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.

recording activity monitor. An object created by the Activity Monitor to record monitor data for database activities. The recorded data can be viewed at a later time.

record length. The sum of a length of all the columns in a table, which is the length of the data as it is physically stored in the database. Records can be fixed or variable in length, depending on how the columns are defined. If all columns are fixed-length columns, the record is a fixed-length record. If one or more columns are variable-length columns, the record is a variable-length column.

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. See also "circular log" on page 13.

Recoverable Resource Manager Services (RRSAF). Recoverable Resource Manager Services attachment facility, which is a DB2 for z/OS and OS/390 subcomponent that uses OS/390 Transaction Management

and Recoverable Resource Manager Services to coordinate resource commitment between DB2 for z/OS and OS/390 and all other resource managers that also use OS/390 RRS in an OS/390 system.

recovery. The process of rebuilding a database or table space that has become unusable because of hardware or software failure, or both. The process includes by restoring a backup image and may include rolling database logs forward in time. See also "forward recovery" on page 42 and "roll-foward recovery" on page 84.

recovery log. See "database log" on page 23.

recovery pending. A state of the database or table space 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 for z/OS and OS/390, an identifier for an element that is used in recovery (for example, *NID* or *URID*).

RECP. In DB2 for z/OS and OS/390, recovery pending.

REORP. See "REORG pending" on page 81.

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 for z/OS and 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 constraints. 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 cycle. A set of referential constraints in which each table is a descendent of itself.

referential integrity. The state of a database in which all values of all foreign keys are valid. Maintaining referential integrity requires the enforcement of "referential constraints" on all operations that change the data in a table upon which the referential constraints are defined.

referential structure. In DB2 for z/OS and 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.

registration. (1) The process of registering a DB2 table, view, or nickname as a replication source. Contrast with subscription. (2) See "replication source" on page 82.

registration process. In DB2 replication, the process of defining a replication source.

registry database. In a OS/390 environment, a database of security information about principals, groups, organizations, accounts, and security policies.

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

rejected transaction. A transaction that contains one or more updates from replica tables that are in conflict with the master table.

related view. A view that uses or is dependent on another object, such as the parent view or a table.

Relational Connect. A separately orderable DB2 product that provides access to other DBMSs such as Oracle, Sybase, and Microsoft SQL Server.

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” on page 64.

relational database. A database that can be perceived as a set of tables and manipulated in accordance with the relational model of data. Each database includes a set of system catalog tables that describe the logical and physical structure of the data, a configuration file containing the parameter values allocated for the database, and a recovery log with ongoing transactions and archivable transactions.

relational database management system (RDBMS). A collection of hardware and software that organizes and provides access to a relational database.

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

relationship. In DB2 for z/OS and 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.

relationship category. In the Information Catalog Center, a basis to define the relationship type. There are four relationship categories:

- Support
- Hierarchical
- Transformational
- Peer to Peer

Each of these relationship categories has roles associated with it that define how an object can relate to other objects. For example, the support relationship category has ″object″ and ″support object″ roles available.

relationship type. In the Information Catalog Center, a definition that defines the roles an object type can play in a relationship. The default relationship types are:

- Attachment
- Contact
- Contains
- Dictionary
- Input
- Output
- Linked
- Supported

Each default relationship has a specific set of roles that object types can play. For example, the contains relationship type allows parent and child roles. If you added a contains relationship between two objects, one object takes on the ″parent″ role and the other object takes on the ″child″ role.

relative byte address (RBA). In a 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 Universal Database following a fallback to a previous release. This procedure constitutes another migration process.

remote. In DB2 for z/OS and 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. See also “local” on page 57.

remote attach request. In DB2 for z/OS and 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. See also “local database” on page 57.

remote subsystem. In DB2 for z/OS and 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 DB2 for z/OS and OS/390 system.

remote unit of work (RUOW). A unit of work that lets a user or application program read or update data at one location per unit of work. Remote unit of work supports access to one database within a unit of work. An application program can update several remote databases, but it can only access one database within a unit of work. See also “unit of work” on page 104.

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

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

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. See also “read stability” on page 78, “uncommitted read (UR)” on page 103, and “cursor stability” on page 22.

repeating group. A situation in which an entity includes multiple attributes that are inherently the same. The presence of a repeating group violates the requirement of first normal form. In an entity that satisfies the requirement of the first normal form, each attribute is independent and unique in its meaning and its name. See also “normalization” on page 66.

replacement file. In DB2 Data Links Manager, a file whose contents are intended to take the place of an existing file.

replica table. In update-anywhere replication, a type of target table that can be updated locally and also receives updates from the master table through a subscription-set definition. If replication conflict

detection is enabled, changes made to the replica table are rejected, whereas changes made to the master table are retained. See also “update-anywhere replication” on page 105, “master table” on page 61, and “conflict detection” on page 17.

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 registering replication sources and creating subscription sets. This user can also run the Capture and Apply programs.

Replication Alert Monitor. A set of programs that can monitor the activity of the Capture and Apply programs, and depending on user-defined alert conditions, can send notifications to specific users.

Replication Analyzer. A program that can analyze a replication environment for setup problems, configuration errors, and performance issues.

Replication Center. A graphical user interface for DB2 Replication that shows Capture and Apply control servers, registered sources, subscription sets, and Monitor control servers. From the Replication Center, a replication administrator can also perform operational tasks for the Capture and Apply programs.

replication control table. A table in which replication definitions or control information is stored.

replication source. A database table, view, or nickname that is registered as a source for replication. Changes made to this type of table are captured and copied to a target table defined in a subscription-set member. See also “subscription set” on page 94 and “subscription set member” on page 94.

replication subscription. See “subscription set” on page 94.

request commit. In DB2 for z/OS and 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. (1) The source of a request to access data at a remote server. Also, the system that requests the data. For DB2 for z/OS and OS/390, the requester function is provided by distributed data facility to access a remote RDBMS. Depending on the level of DRDA protocol used, a requester can be described as an application requester or a database server. (2) The target of a request from a remote requester.

reserved word. (1) A word that is 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.

residual recovery entry (RRE). A unit of recovery about which DB2 could be in doubt. For example, in a OS/390 environment with IMS, IMS builds a list of residual recovery entries.

resource. In DB2 for z/OS and 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 Access Control Facility (RACF). The Resource Access Control Facility protects the system by giving access to those individuals who have authority to use the resource. RACF is a component of the SecureWay Security Server for z/OS and OS/390.

resource allocation. In DB2 for z/OS and OS/390, the part of plan allocation that deals specifically with database resources.

resource control table (RCT). In DB2 for z/OS and 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 a 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 for z/OS and OS/390 code that prevents dynamic manipulative SQL statements from exceeding specified time limits. Also known as the *governor*.

resource limit specification table. In DB2 for z/OS and OS/390, a site-defined table that specifies the limits to be enforced by the resource limit facility.

response file. An ASCII file that can be customized with the setup and configuration data that will automate an installation. The setup and configuration data must be entered during a interactive install, but with a response file, the installation can proceed without any intervention.

response file generator. A utility that creates a response file from an existing installed and configured DB2 product. You can use the generated response file to recreate the same setup on other computers.

restart pending (RESTP). In DB2 for z/OS and 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 for z/OS and OS/390 invokes after restart if the system parameter LBACKOUT=AUTO.

RESTP. See “restart pending”.

RESTORE. A DB2 utility that is used to rebuild a damaged or corrupted database or tablespace from a backup image that is produced with the BACKUP utility.

restore. To rebuild a damaged or corrupted database or table space from a backup image produced with the BACKUP utility.

restore set. A backup copy of a database or table space plus zero or more log files, which, 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 for z/OS and 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. See also “temporary table” on page 99.

retained lock. A MODIFY lock that a DB2 for z/OS and 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 for z/OS and OS/390 subsystem.

retention-limit pruning. The pruning of CD and UOW tables by the Capture program that are older than a user-specified limit.

retrieve_query utility. In DB2 Data Links Manager, a standalone DLFM utility that displays the status of all files that are managed by a particular Data Links server. The retrieve_query utility can also be invoked with the **dlfm retrieve** command.

revoke. To remove a privilege or authority from an authorization identifier.

rework. (1) To convert an insert into a replication target table to an update if the insert fails because the row already exists in the target table. (2) To convert an update to a replication target table to an insert if the update fails because the row does not exist in the target table.

RID. See "record identifier" on page 78.

RID pool. See "record identifier pool" on page 78.

right outer join. The result of a join operation that includes the matched rows of both tables that are being joined and preserves the unmatched rows of the second join operand. See "join" on page 55. See also "left outer join" on page 56 and "full outer join" on page 42.

RLF. See "resource limit facility" on page 83.

role. In the Information Catalog Center, a descriptor that is associated with the relationship category. The relationship category that you choose determines what roles are available for each object type.

roll forward. To update the data in a restored database or table space by applying changes recorded in the database log files. See also "forward recovery" on page 42.

rollforward recovery. A process invoked through the rollforward utility that is used to recover a database by applying transactions that were recorded in the database recovery log file. See also "forward recovery" on page 42.

roll back. To restore data that is changed by an SQL statements to the state at its last commit point. See "backout" on page 7. See also "point of consistency" on page 73.

roll out. The efficient deletion of a large portion of a multidimensional clustering (MDC) table, and is possible when a DELETE statement is processed that has certain types of predicates (equality, range, BETWEEN, IN) on one or more dimension columns. Most logging can be avoided and, in certain cases, all per-row processing can also be avoided.

root page. In DB2 for z/OS and 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. A user-defined method, user-defined function, or stored procedure.

row. The horizontal component of a table that consists of a sequence of values, one for each column of the table.

row-capture rules. Rules based on changes to registered columns that define when and whether the Capture program writes a row to a CD table, or when and whether the Capture triggers write a row to a CCD table.

row function. A function that returns one row of values and must be defined as an SQL function.

ROWID. See “row identifier”.

row identifier (ROWID). A value that uniquely identifies a row. This value is stored with the row and does not change.

row lock. A lock on a single row of data. See also “locking” on page 58 and “table lock” on page 97.

row trigger. In DB2 for z/OS and OS/390, a trigger that is defined with the trigger granularity FOR EACH ROW.

row-value expression. In a z/OS or OS/390 environment, a comma-separated list of value expressions enclosed in parentheses.

RR. See “repeatable read ” on page 81.

RRE. See “residual recovery entries” on page 82.

RS. See “read stability” on page 78.

RRSAF. See “Recoverable Resource Manager Services” on page 78.

RUOW. See “remote unit of work” on page 81.

S

satellite. A DB2 server that synchronizes with its group at the DB2 control server.

Satellite Administration Center. A user interface that provides centralized administrative support for satellites.

saved search. In the Information Catalog Center, a set of search criteria that is saved for subsequent use. A saved search is displayed as an object under the **Saved Searches** folder in the tree.

SBCS. See “single-byte character set” on page 89.

SCA. In DB2 for z/OS and 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. See also “column function” on page 15 and “table function” on page 97.

scale. The number of digits in the fractional part of a number.

scattered read. A method of reading contiguous data pages from disk to discontiguous portions of memory. See also “block based I/O” on page 8.

schema. (1) A collection of database objects such as tables, views, indexes, or triggers that define a database. A database schema provides a logical classification of database objects. (2) In DB2 for z/OS

and 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, based on 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.

scrollability. In a z/OS or OS/390 environment, the ability to use a cursor to fetch in either a forward or a backward direction. The FETCH statement supports multiple fetch orientations to indicate the new position of the cursor. See also "fetch orientation" on page 41.

scrollable cursor. A cursor that can be moved in both a forward and a backward direction. See also "nonscrollable cursor" on page 66.

SDK. See "Software Developer's Kit" on page 90.

SDWA. In a OS/390 environment, the system diagnostic work area.

search. In the Information Catalog Center, to request the display of the objects that meet user-specified criteria.

search condition. A criterion for selecting rows from a table. A search condition consists of one or more predicates.

search criteria. In the Information Catalog Center, options and character strings that are used to specify how to perform a search. The search criteria can include object type names, property values, whether the search is for an exact match, and whether the search is case sensitive.

secondary authorization ID. In DB2 for z/OS and OS/390, an authorization identifier 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 for z/OS and 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 z/OS and OS/390 equivalent is *new* structure. See also "primary group buffer pool" on page 74.

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. See also "primary log" on page 74.

section. 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.

segment. A group of pages that hold a row of a single table. See also "segmented table space".

segmented table space. In DB2 for z/OS and 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. 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.

sensitive cursor. A cursor that is sensitive to changes made to the database after the result table has materialized. See also “insensitive cursor” on page 52.

sensitive static cursor. The order of the rows and size of the result table is static. The size of the result table does not grow after the rows are materialized. The order of the rows is established as the result table is materialized. Newly inserted rows are not visible to SENSITIVE STATIC cursors once the rows of the result table have been materialized. Rows in the result table do not move if columns in the ORDER BY clause are updated in rows that have already been materialized.

Static cursors have visibility to changes made by the cursor using UPDATE WHERE CURRENT OF or DELETE WHERE CURRENT OF. Visibility of changes made outside the cursor depends on the type of FETCH that is used with a SENSITIVE STATIC cursor.

sequential data set. A non-DB2 for z/OS and 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 for z/OS and OS/390 database utilities require sequential data sets.

sequential prefetch. 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.

serial cursor. See “nonscrollable cursor” on page 66.

serialization. (1) The consecutive ordering of items. (2) The process of controlling access to a resource to protect the integrity of the resource.

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. See “database server” on page 24. (3) The target of a request from a remote requester. In the DB2 environment, the server function is provided by the distributed data facility, which is used to access DB2 data from remote applications. See also “application server” on page 5.

server definition. In a federated system, the name and information that defines the data sources to the federated database. The server definition is used by the wrapper when SQL statements that use nicknames are submitted to the federated database.

If the data source is a RDBMS, this information includes the type and version of the RDBMS, and the name for the data source on the RDBMS. It also includes metadata that is specific to the RDBMS. For example, a DB2 family data source can have multiple databases, and the definition must specify which database the federated server can connect to.

In contrast, an Oracle data source has one database, and the federated server can connect to the database without needing to know its name. The name is therefore not included in the federated server definition of an Oracle data source.

server option. In a federated system, information within a server definition that either configures the wrapper itself, or affects the way that DB2 uses the wrapper. Server option values are stored in the global catalog.

server profile. A profile that contains information about server instances on a system, and databases within each server instance. See also “client profile” on page 14.

server-side programming. A method for adding DB2 data into dynamic Web pages. Three common types of server-side programs are Common Gateway Interface (CGI), Web server API programs, and Java servlets.

service class. In DB2 for z/OS and 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 definition. In a federated database system, a description of a data source.

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 for z/OS and 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 that correspond to the relational operators union, difference, and intersection. A set operator derives a result table by combining two other result tables.

shadow index. A new index structure created during the index reorganization. It is not visible to users for access until the database manager has fully rebuilt the index.

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 for z/OS and 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. See also “exclusive lock” on page 38.

shift-in character. A special control character (X'0F') that is used in EBCDIC systems to denote that the subsequent bytes represent SBCS characters. See also “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. See also “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 for z/OS and 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, a LOB string is not a short string. See also “long string” on page 60.

signal. A communication mechanism for replication that allows the Capture, Apply, and Monitor programs to communicate with each other asynchronously.

sign-on. A request that is made on behalf of an individual CICS or IMS application process by an attachment facility to enable DB2 for z/OS and OS/390 to verify that it is authorized to use DB2 resources.

simple page set. In DB2 for z/OS and 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 gigabytes, another data set is created, and so on up to a total of 32 data sets. DB2 for z/OS and OS/390 considers the data sets to be a single contiguous linear address space that contains a maximum of 64 gigabytes. Data is stored in the next available location within this address space without regard to any partitioning scheme.

simple table space. In DB2 for z/OS and 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. See also “double-byte character set” on page 36 and “multibyte character set” on page 63.

single-precision floating point number. A 32-bit approximate representation of a real number.

slice. The set of blocks that contain pages with data having a certain value of one of the clustering dimensions. If we consider a slice in each dimension, where a slice contains a particular value for its dimension, a cell is the intersection of these slices.

SMF. See “system management facility” on page 96.

SMS. See “Storage Management Subsystem” on page 93.

SMS table space. See “system-managed space table space ” on page 96.

SNA. See “Systems Network Architecture” on page 96.

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. A record of the current state of the database environment. See also “performance snapshot” on page 72, “explain” on page 38, and “health snapshot” on page 47.

socket. A communications handle used by TCP/IP.

socket interface. 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 if a database restart is required.

Software Developer's Kit (SDK). An obsolete term for Application Development Client. See "Application Development Client" on page 4.

source. In the Data Warehouse Center, a table, view, or file that is input to a step. See also "target" on page 98.

sourced function. 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). See also "external function" on page 40, "user-defined function" on page 106, "built-in function" on page 9, and "SQL function" on page 91.

source program. A set of host language statements and SQL statements that is processed by an SQL precompiler.

source server. A database that contains registered replication sources.

source table. A table that contains data that is to be replicated to a target table. See also "target table" on page 98.

source type. An existing type that is used to internally represent a distinct type.

spatial column. A table column or view column that has a spatial data type. This data type allows the column to contain coordinates that define locations within a particular region of the earth.

spatial reference system. In DB2 Spatial Extender, a set of parameter values that includes:

- Coordinates that define the maximum possible extent of space referenced by a given range of coordinates.
- An identifier of the coordinate system from which the coordinates are derived.
- Numbers that, when applied in certain mathematical operations, convert coordinates received as input into values that can be processed with maximum efficiency.

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 for z/OS and OS/390, a particular user-defined function that is known to the database manager by its specific name. When a user-defined function is defined to the database, every function is assigned a specific name that is unique within its schema. The specific name is important for functions that have the same name but have either a different number of parameters or different data types associated with those parameters. The user can either provide this name or use the default.

spill file. A temporary file created by the Apply program that is used to hold data for updating multiple target tables.

SPUFI. See “SQL Processor Using File Input” on page 92.

SQL. See “Structured Query Language” on page 93.

SQL Assistant. A wizard available in several DB2 tools and centers that generates SQL statements graphically.

SQL authorization ID (SQL ID). In DB2 for z/OS and OS/390, the authorization identifier 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. 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. The symbol that is used to enclose an SQL delimited identifier. The escape character is the quotation mark, except in COBOL applications, where the user assigns the symbol to be either a quotation mark or an apostrophe.

SQL function. A user-defined function in which the CREATE FUNCTION statement contains the source code. The source code is a single SQL expression that evaluates to a single value. The SQL user-defined function can return only one parameter.

SQL ID. See “SQL authorization ID”.

SQLJ. A three-part standard for supporting embedded SQL in Java programs (Part 0), defining and calling Java stored procedures and user-defined functions (Part 1), and using database structured types in Java (Part 2).

SQL path. In DB2 for z/OS and OS/390, an ordered list of schema names that is 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 procedure. An application program written in SQL that can be invoked with the SQL CALL statement. See also “external procedure” on page 40.

SQL processing conversation. Any conversation that requires access to DB2 for z/OS and OS/390 data, either through an application or by dynamic query requests.

SQL Processor Using File Input (SPUFI). In DB2 for z/OS and OS/390, a facility of the TSO attachment subcomponent that lets the DB2I user run SQL statements without embedding them in an application program.

SQL return code. Either SQLCODE or SQLSTATE.

SQL routine. In DB2 for z/OS and OS/390, a user-defined function or stored procedure that is based on code that is written in SQL.

SQL statement coprocessor. In a z/OS or OS/390 environment, an alternative to the DB2 precompiler that lets the user process SQL statements at compile time. The user invokes an SQL statement coprocessor by specifying a compiler option.

SQL string delimiter. In DB2 for z/OS and 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 to be either an apostrophe or a quotation mark (").

SSCP. See "system services control point" on page 96.

SSI. In a OS/390 environment, subsystem interface.

SSM. In DB2 for z/OS and OS/390, subsystem member.

stack. An area in memory that stores temporary register information, parameters, and return addresses of subroutines.

staging table. A CCD table that is used to save data before that data is replicated to the target database. A CCD table used for staging data can function as an intermediate source for updating data to one or more target tables. See also "consistent-change-data table" on page 18.

standalone. An attribute of a program that allows the program to run separately from DB2 for z/OS and OS/390, without using DB2 for z/OS and OS/390 services.

schema. The type of relational database schema used by the DB2 OLAP Server, often created in the Data Warehouse Center.

statement. An instruction in a program or procedure.

statement handle. In the CLI, a handle that refers to the data object that contains information about an SQL statement. Such information includes 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 for z/OS and OS/390 environment, the character string form of the statement.

statement trigger. In DB2 for z/OS and OS/390, a trigger that is defined with the trigger granularity FOR EACH STATEMENT. See also "trigger" on page 102.

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 "bind" on page 8. See also "dynamic bind" on page 36.

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. See also "embedded SQL" on page 37 and "dynamic SQL" on page 36.

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 disks on which DB2 for z/OS and OS/390 data can be stored.

Storage Management Subsystem (SMS). In OS/390, software that automates as much as possible the management of physical storage by centralizing control, automating tasks, and providing interactive controls for system administrators. SMS can reduce users' concerns about physical details of performance, space, and device management.

stored procedure. (1) An application program, possibly containing SQL statements, that can be invoked with the SQL CALL statement. (2) A user-written application program that can be started through the use of the SQL CALL statement.

Stored Procedure Builder. Renamed and enhanced in DB2 Universal Database Version 8. See "Development Center" on page 33.

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. (1) In programming languages, the form of data that is used for storing and manipulating text. (2) A sequence of bytes that may represent characters.

strong typing. 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.

structure. A name that refers collectively to different types of DB2 objects, such as tables, databases, views, indexes, and table spaces.

Structured Query Language (SQL). A standardized language for defining and manipulating data in a relational database.

structured type. A way to gather a collection of object attributes under a single type, which allows for the control of the semantics of an object.

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 that

work on behalf of the same application. All subagents working for the application are initiated by the initiating agent for that application. See also “coordinating agent” on page 20.

subcomponent. A group of closely related DB2 for z/OS and OS/390 modules that work together to provide a general function.

subject area. (1) 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. (2) In the Information Catalog Center, an object type that identifies and groups the processes that relate to a logical area of the business. For example, if you are creating an information catalog of marketing and sales data, you define object types Sales and Marketing and select to make them subject areas. Then any objects of type Sales or Marketing are grouped under the corresponding subject.

subject search. See “browse” on page 9.

subject table. The table for which a trigger is created. When the defined triggering event occurs on this table, the trigger is activated.

subordinate agent. See “subagent” on page 93.

subpage. In DB2 for z/OS and 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. (1) The process of creating subscription sets and subscription-set members. Contrast with registration. (2) See “subscription set”.

subscription cycle. The 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 its status and current progress.

subscription set. A replication definition that controls the replication of changed data during a subscription cycle. A subscription set can contain zero or more subscription-set members.

subscription-set member. A replication definition that maps a registered replication source with a replication target. Each member defines the structure of the target table and which rows and columns that will be replicated from the source table.

subselect. The form of a query that does not include an ORDER BY clause, an UPDATE clause, or UNION operators.

subset. To replicate data from part of a source table, rather than from the entire table, to a target table. You can subset by rows or by columns.

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 for z/OS and OS/390, a distinct instance of a relational database management system (RDBMS).

success code set. One or more expressions that specify the return codes of a successful task. For example, specify > -1 to consider all return codes of zero or more a success.

summary table. A specialized type of materialized query table whose fullselect contains a GROUP BY clause which summarizes data from the tables referenced in the fullselect. See also "materialized query table" on page 61.

superuser. A user that has various system control authorities above and beyond that of the ordinary user. In UNIX environments, the standard superuser is root.

support relationship category. In the Information Catalog Center, a category for relationship types that connects supporting objects to another object. For example, you can connect a News object to a Spreadsheet object.

support relationship type. In the Information Catalog Center, a category of relationship types that provides additional information about your information catalog or enterprise. For example, the "Information Catalog Center News" object type in the sample information catalog. See also "relationship type" on page 80.

surrogate pair. In a z/OS or OS/390 environment, a coded representation for a single character that consists of a sequence of two Unicode values, where the first value of the pair is a high-surrogate in the range U+D800 through U+DBFF, and the second value is a low-surrogate in the range U+DC00 through U+DFFF. Surrogate pairs provide an extension mechanism for encoding 917 476 characters without requiring the use of 32-bit characters.

symbolic destination name. The name of a remote partner. The name corresponds to an entry in the CPI-C 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.

synchpoint. A replication control table value for the DB2 log or journal record sequence number of the last change applied during the most recent Apply cycle. This value is also used to coordinate the pruning of CD tables.

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 on the occurrences of specific events, such as a common timing signal. See also "asynchronous" on page 5.

synchronous replication. Also known as real-time replication, a type of replication that delivers updates continuously and within the scope of source transactions.

sync point. See "point of consistency" on page 73.

synonym. In DB2 for z/OS and OS/390, an alternative name, in SQL, for a table or view.

syntactic character set. A set of 81 graphic characters that are registered in the IBM registry as character set 00640. This set is used for syntactic purposes maximizing portability and interchangeability across systems and country boundaries. It is contained in most of the primary registered character sets, with a few exceptions. See also "invariant character set" on page 54.

Sysplex. See "Parallel Sysplex" on page 70.

Sysplex query parallelism. Parallel execution of a single query that is accomplished by using multiple tasks on more than one DB2 for z/OS and OS/390 subsystem. See also “query CP parallelism” on page 77.

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 for z/OS and OS/390 creates internally, such as prefetch processing, deferred writes, and service tasks. See also “agent” on page 2.

system authority. SYSCTRL and SYSMAINT authority levels with full privileges for managing the system but without the ability to access the data.

system catalog. See “catalog” on page 11.

system conversation. The conversation that two DB2 for z/OS and 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. The directory is created when the first database is created or cataloged on the system. See also “local database directory” on page 57.

system diagnostic work area (SDWA). In a OS/390 environment, the data that is recorded in a SYS1.LOGREC entry that describes a program or hardware error.

system-directed connection. A connection that an RDBMS manages by processing SQL statements with three-part names (or nicknames), providing a level of location transparency. See also “application-directed connections” on page 4.

system management facility (SMF). In DB2 for z/OS and OS/390, a standard feature that collects and records a variety of system and job-related information. For example, statistics, accounting information, and performance data.

system-managed space (SMS) table space. A table space whose space is managed by the operating system. This storage model is based on files that are created under subdirectories and managed by the file system. See also “database-managed space table space” on page 23.

system monitor. See “database system monitor” on page 25.

system services control point (SSCP). The control point in an SNA network that provides network services for dependent nodes.

Systems Network Architecture (SNA). An architecture that describes the logical structure, formats, protocols, and operational sequences for transmitting information units through networks, and also operational sequences for controlling the configuration and operation of networks.

SYS1.DUMPxx data set. In a OS/390 environment, a data set that contains a system dump.

SYS1.LOGREC. In a 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” on page 7, “declared temporary table” on page 30, and “temporary table” on page 99.

table check constraint. A user-defined constraint that specifies the values that specific columns of a base table can contain.

table collocation. In a partitioned database environment, a state that occurs when two tables are stored in the same database partition group and that have the same number of compatible partitioning keys. When this happens, DB2 can choose to perform the join or subquery processing at the database partition where the data is stored.

table designator. A column name qualifier that designates a specific object table.

table expression. An expression that creates a temporary result table from a simple query. For example, a table expression might be a query that selects all the managers from several departments and further specifies that they have over 15 years of working experience and are located at the main branch. See also “common table expression” on page 16.

table function. 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. See also “column function” on page 15 and “scalar function” on page 85.

table locator. In DB2 for z/OS and 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 lock. A lock on a table of data. See also “row lock” on page 85 and “row identifier” on page 85.

table-mode processing. A type of replication subscription-set processing in which the Apply program retrieves all the data from the source CD table, then applies the data (one member at a time) to each target table, and finally commits its work. Contrast with “transaction-mode processing” on page 101.

table queue. A mechanism for transferring rows between database partitions. 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 a single-partition 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. 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 for z/OS and OS/390, a page set that is used to store the records in one or more tables.

table space container. 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 for z/OS and OS/390, a set of table spaces and partitions that should be recovered together if each contains a table that is a parent or descendent of a table in one of the others, or the set contains a base table and associated auxiliary tables. A table space set can contain both types of relationships.

tag. An element of the tag language. Tags indicate actions to be taken when the tag language file is imported to the information catalog.

tag language. A format for defining object types and objects, and actions to be taken on those object types and objects, in the Data Warehouse Center or the information catalog.

tag language file. A file that contains tag language which describes objects and object types to be added, updated or deleted in the Data Warehouse Center or in the information catalog, when the file is imported.

In the Information Catalog Center, a tag language file is produced when you:

- Transfer a delete history log.
- Extract descriptive data from another database system using an extract program.

target. In the Data Warehouse Center, a table, view, or file that is produced or populated by a step; the output of a step. See also “source” on page 90.

target server. A database that contains replication target tables.

target table. A table that is the target for replicated changes from a registered replication source. It can be a user copy table, a point-in-time table, a base aggregate table, a change aggregate table, a CCD table, or a replica table.

task. In the Task Center, a unit of work and its associated schedule and task actions. Tasks can be set to run on schedules and can perform various actions based on the success or failure of the task. DB2 scripts, operating scripts, and warehouse steps are all examples of tasks. See also “task action” and “step” on page 93.

task action. In the Task Center, an action that is performed based on the completion status of a particular task. For example, ″If Task A completes successfully, run Task B,″ and ″If Task Z fails, disable the schedule of Task Y.″ See also “task” and “step” on page 93.

task category. A string that is associated with any number of tasks in the Task Center for easier administration of related tasks. For example, you can create a task category named ″Payroll″ then group all payroll-related tasks in the Payroll category.

Task Center. The DB2 graphical interface for organizing task flow, scheduling tasks, and distributing notifications about the status of completed tasks.

task control block (TCB). A control block that is used to communicate information about tasks within an address space that are connected to DB2 for z/OS and 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” on page 101.

TCP/IP port. A 2-byte value that identifies 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. See also “business metadata” on page 9.

template. In a z/OS or OS/390 environment, a DB2 utilities output data set descriptor that is used for dynamic allocation. A template is defined by the TEMPLATE utility control statement.

temporary table. A table that holds temporary data. For example, temporary tables are useful for holding or sorting intermediate results from queries that contain a large number of rows. The two kinds of temporary tables, which are created by different SQL statements, are the created temporary table and the declared temporary table. See also "result table" on page 83, "created temporary table" on page 21, and "declared temporary table" on page 30.

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.

territory code. Specifies the country or region of the database. The territory code is used by DB2 to preset the default collation order for a SBCS database.

thread. (1) The database manager structure that describes an application's connection, traces its progress, processes resource functions, and delimits its accessibility to the database manager resources and services. Most DB2 for z/OS and OS/390 functions execute under a thread structure. See also "allied thread" on page 3 and "database access thread" on page 23. (2) In some operating systems, the smallest unit of operation to be performed in a process.

three-part name. The full name of a table, view, or alias that consists of a location name, authorization identifier, 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.

throttled utilities. Utilities that have a limit placed on the resources that would otherwise be consumed. The degree to which the resources are limited is based on the current workload of the system. Supported utilities include backup, restore, and table space reorganization.

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 that is 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. An abnormal termination of either the DB2 for z/OS and 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 for z/OS and OS/390 waits for IRLM services after starting, and the amount of time IRLM waits if a resource that an application requests is unavailable. If either of these time specifications is exceeded, a timeout is declared.

Time-Sharing Option (TSO). In a z/OS and OS/390 environment, software that provides interactive communications, allowing a user or programmer to start an application from a terminal and work with the application. TSO is required for binding application plans and packages and for executing several online functions that are provided with DB2 for z/OS and OS/390.

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.

Tivoli Space Manager. A feature of the Tivoli Storage Manager product that moves files in and out of a secondary storage medium based upon actual file accesses in the primary native file system. This feature can be used with DB2 Data Links Manager to enable DATALINK files to be stored in a virtually infinitely sized file system.

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 backups.

TM Database. See "Transaction Manager Database" on page 101.

to-do. A state of a unit of recovery that indicates that the changes by the unit of recovery to recoverable DB2 for z/OS and 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" on page 101.

trace. (1) A DB2 for z/OS and OS/390 facility that provides the ability to monitor and collect monitoring, auditing, performance, accounting, statistics, and serviceability (global) data. (2) For DB2 replication, a facility that provides the ability to collect monitoring, auditing, and performance data for the Capture program, Apply program, or Replication Alert Monitor.

transaction. 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" on page 104.

transaction-based replication. A type of replication processing in which every transaction is replicated to the target table as it is committed in the source table. Contrast with "transaction-consistent replication" on page 101.

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-consistent replication. A type of replication processing in which the net result of all transaction updates is replicated to the target table. Contrast with"transaction-based replication" on page 100.

transaction lock. In DB2 for z/OS and 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-mode processing. A type of replication subscription-set processing in which the Apply program retrieves data from the source CD table, the applies the data to the target table in the same commit sequence used at the source. The Apply program processes transactions for all subscription-set members together, rather than sequentially. Contrast with "table-mode processing" on page 97.

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.

transformation relationship category. In the Information Catalog Center, a category for relationship types that connects transformation objects to data resources. For example, you can connect a Transformation object to a File object. Objects that are connected with this category of relationship are displayed in the Information Catalog Center Show Lineage Tree window.

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. Transformers have corresponding step types for the types of data manipulation that the steps perform; for example, a clean step uses the Clean Data transformer.

transition table. A temporary table that contains all the affected rows of the subject table in their state before or after the triggering event occurs. Triggered SQL statements in the trigger definition can reference the table of changed rows in the old state or the new state.

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 provides peer-to-peer connectivity functions for both local and wide area networks.

tree view. A view that provides a hierarchical view of an object and the objects that it contains.

trigger. (1) 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 is stored in a DB2 database and executed when a certain event occurs in a DB2 table.

trigger activation. 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. An indication in a trigger definition of whether the trigger should be activated before or after the triggered event.

trigger body. The set of SQL statements that is executed when a trigger is activated and its triggered action condition evaluates to true.

trigger cascading. 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) 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) 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.

trigger event. In a trigger definition, the operation, either INSERT, DELETE, or UPDATE, that will cause the trigger to be activated.

trigger granularity. A characteristic of a trigger, which determines whether the trigger is activated either once for the triggering SQL statement, or once for each row that the SQL statement modifies.

trigger package. A package that is created when a CREATE TRIGGER statement is executed. The package is executed when the trigger is activated.

triggered SQL statements. 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*.

triggering event. The event that causes a trigger to be activated. In general, a triggering event is the insertion, deletion, or update of rows in a specific table.

triggering SQL operation. The SQL operation that causes a trigger to be activated when performed on the subject table.

triggering table. 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. See “Time-Sharing Option” on page 100.

TSO attachment facility. A DB2 for z/OS and OS/390 facility that consists 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. See also “tuple”.

tuple. A synonym for a row in a table. See also “tuning parameters table”.

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. See also “distributed transaction” on page 34.

typed parameter marker. A parameter marker that is specified along with its target data type. It has the general form: *CAST (? AS data-type)*.

typed table. A table in which the data type of each column is defined separately or the types for the columns are based on the attributes of a user-defined structured type.

typed view. A view in which the data type of each column is derived from the result table or the tupes for the columns are based on the attributes of a user-defined structure type.

type 1 index. An index that is not a type 2 index. As of DB2 for z/OS and OS/390 Version 8, type 1 indexes are no longer supported. See also “type 2 indexes”

type 2 index. A pseudo delete index that supports variable length key parts with length greater than 255. See also “type 1 indexes”.

U

UCS-2. Universal Character Set, coded in 2 octets, which means that characters are represented in 16-bits per character.

UDF. See “user-defined function” on page 106.

UDT. See “user-defined type” on page 106.

UFS. See “UNIX File System (UFS)” on page 104.

unambiguous cursor. A cursor that allows a DBMs to determine whether blocking can be used with the answer set. A cursor that is defined FOR FETCH ONLY or FOR READ ONLY can be used with blocking, whereas a cursor that is defined FOR UPDATE cannot. See also “ambiguous cursor” on page 3.

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 that it is reading, unless the other application attempts to drop or alter the table. See also “repeatable read ” on page 81, “cursor stability” on page 22, and “read stability” on page 78.

underlying view. In DB2 for z/OS and OS/390, the view on which another view is directly or indirectly defined.

undo. (1) To recover the last edit that has taken place. (2) A state of a unit of recovery that indicates that the changes that the unit of recovery made to recoverable DB2 for z/OS and 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. See also "ASCII" on page 5 and "EBCDIC" on page 37.

uniform resource locator (URL). A Web address, which offers a way of naming and locating specific items on the Web.

union. An SQL operation that combines the results of two select statements. Unions are often used to merge lists of values that are obtained from several tables.

unique constraint. 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 identifier (UI). In the Information Catalog Center, a key for an object. The key is comprised of up to 16 properties, which, when concatenated in a designated order, uniquely identify the object during the import function.

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 for z/OS and OS/390. See also "unit of work".

unit of work. (1) 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 for z/OS and OS/390 *multisite update* operation, a single unit of work can include several *units of recovery*. Synonym for "transaction" on page 100. See also "unit of recovery" and "multisite update" on page 64. (2) In the Information Catalog Center, 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.

unit-of-work (UOW) table. A replication control table stored in the Capture control 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 CD table to produce transaction-consistent change data.

UNIX File System (UFS). The native file system in the Solaris Operating Environment.

unlink. The action that DB2 Data Links Manager takes to give up control of a file that is no longer referenced in a table that contains a DATALINK column. A file can be unlinked as the result of such database actions as an SQL UPDATE, DELETE, or DROP TABLE.

unlinked file. In a DB2 Data Links Manager environment, a file that is controlled by the native file system on an operating system. By contrast, a linked file is controlled by the DLFF component.

unlock. To release an object or system resource that was previously locked and return it to general availability within DB2 for z/OS and 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.

updatability. The ability of a cursor to perform positioned updates and deletes. The updatability of a cursor can be influenced by the SELECT statement and the cursor sensitivity option that is specified on the DECLARE CURSOR statement.

update-anywhere replication. A replication configuration in which all tables are both registered sources and read-write targets. One table is the primary source table for full refresh of all the others. In this configuration, there is an implicit replication hierarchy among the source and target tables. Contrast with "peer-to-peer replication" on page 72. See also "multi-tier replication" on page 64, "master table" on page 61, and "replica table" on page 81.

update hole. A row for a SELECT statement of a cursor that no longer has a corresponding row in the base table because the row was updated. An update hole is created when a row in the base table is updated such that the row no longer qualifies to be in the result set while a cursor is open whose SELECT statement result contains the row that is updated. Such a row is no longer accessible though the cursor. See also "delete hole" on page 31.

update-in-place. In a DB2 Data Links Manager environment, the process of making changes to a linked file while a DATALINK column value in a database is pointing to that file. Any changes to linked files during an update-in-place operation become visible to database users when the DB2 host is notified that the update is complete.

update-in-progress state. The logical state of a file under the control of a DB2 Data Links Manager that is in the process of being updated. A linked file enters this state after it is opened using a write token, and is no longer in this state when the DB2 host is notified that the update is complete.

update rule. A condition enforced by the database manager that must be met before a column can be updated.

update trigger. In DB2 for z/OS and OS/390, a trigger that is defined with the triggering SQL operation UPDATE.

updating. In a DB2 Data Links Manager environment, the act of modifying a linked file.

upstream. In DB2 for z/OS and OS/390, the node in the syncpoint tree that is responsible, in addition to other recovery or resource managers, for initiating the execution of a two-phase commit.

UR. See "uncommitted read (UR)" on page 103.

URE. In DB2 for z/OS and OS/390, unit of recovery element.

URID (unit of recovery ID). In DB2 for z/OS and 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.

URL. See "uniform resource locator" on page 104.

user. In the Information Catalog Center, a person who accesses the information available in the information catalog but who is not an administrator. Some users can also perform object management

tasks normally performed by administrators, such as creating and updating objects. See also "administrator" on page 2 and "power user" on page 73.

user copy table. A replication target table whose content matches all or part of a registered source table and contains only user data columns.

user-defined data type. See "distinct type" on page 34.

user-defined distinct type. See "distinct type" on page 34.

user-defined function (UDF). A function that is defined to DB2 by using the CREATE FUNCTION statement and that can be referenced thereafter in SQL statements. A user-defined function can be an external function, or an SQL function. See also "built-in function" on page 9.

user-defined performance variable. A performance variable that is 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 structure type. See "structure type" on page 93.

user-defined type (UDT). A data type that is not native to the database manager and was created by a user. In DB2 Universal Database, the term distinct type is used instead of user-defined type.

user exit. A program that is used to interact with storage devices that are not directly supported by the operating system. When a program user exit program is invoked, the database manager passes control to the executable file. Only one user exit program can be invoked within a database manager instance.

user mapping. In a federated system, the association between the authorization ID at the federated server and the authorization ID at the data source. User mappings are needed so that distributed requests can be sent to the data source. User mappings are created when a user's authorization ID to access the federated database differs from the user's authorization ID to access a data source. The CREATE USER MAPPING statement is used to define the association. The ALTER USER MAPPING statement is used to modify a user mapping that you have already created.

user options. In a federated system, parameters of the CREATE USER MAPPING and ALTER USER MAPPING statements to which values related to authorization are assigned. For example, suppose that a user has the same ID with different passwords for the federated database and a data source. For the user to access the data source, it is necessary to map the passwords to one another. This is accomplished with the user option REMOTE_PASSWORD. See "user mapping".

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.

user view. In logical data modeling, a model or representation of critical information that the business requires.

UT. In DB2 for z/OS and OS/390, utility-only access.

UTC. See "Coordinated Universal Time" on page 20.

UTF-8. Unicode Transformation Format, 8-bit encoding form, which is designed for ease of use with existing ASCII-based systems. The CCSID value for data in UTF-8 format is 1208. DB2 for z/OS and OS/390 supports UTF-8 in mixed data fields.

UTF-16. Unicode Transformation Format, 16-bit encoding form, which is designed to provide code values for over a million characters and is a superset of UCS-2. The CCSID value for data in UTF-16 format is 1200. DB2 for z/OS and OS/390 supports UTF-16 in graphic data fields.

V

value. (1) The alpha or numeric content of a field or variable. (2) The smallest unit of data manipulated in SQL. (3) 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. See also "constant" on page 18.

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. See also "not-deterministic function" on page 66.

variable-length string. A character, graphic, or binary string whose length is not fixed but can range within set limits. Also referred to as a *varying length string.*

vectored I/O. See "scattered read" on page 85.

version. In DB2 for z/OS and OS/390, a member of a set of similar programs, DBRMs, packages, or LOBs. Some examples are:

- 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.

version recovery. The restoration of a previous version of the database, using an image that was created during a backup operation. See also "crash recovery" on page 21 and "forward recovery" on page 42.

view. (1) A logical table that consists of data that is generated by a query. A view is based on an underlying set of base tables, and the data in a view is determined by a SELECT statement that is run on the base tables. Contrast with base table. (2) A way of looking at the information about, or contained in objects. Each view might reveal different information about its objects. See also "base table" on page 7.

view check option. In DB2 for z/OS and 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 a OS/390 environment, an IBM licensed program that controls communication and the flow of data in an SNA network.

Visual Explain. A tool that provides a graphical interface for database administrators and application programmers 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" on page 107.

VTAM. See "Virtual Telecommunications Access Method".

W

warehouse. See "data warehouse" on page 27.

warehouse agent. In the Data Warehouse Center, a run-time process, capable of running on various operating systems, that performs data extraction, transformation, movement, and loading (ETML) and that can also start user programs. See also "warehouse server".

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 server. In the Data Warehouse Center, the Windows or AIX component that manages and schedules the data extraction, transformation, movement and loading (ETML) tasks run by the warehouse agents. See also "warehouse agent".

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. (2) In DB2 replication, the process of starting the Capture program by using existing data from the Capture control tables. Contrast with "cold start" on page 14.

well-known address. An address that is 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 for z/OS and 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 system, the mechanism that the federated server uses to communicate with and retrieve data from the data sources. To implement a wrapper, the federated server uses routines stored in a library called a wrapper module. These routines allow the federated server to perform operations such as connecting to a data source and retrieving data from it iteratively. The DB2 federated instance owner uses the CREATE WRAPPER statement to register a wrapper for each data source that is to be included in the federated system.

write token. The authorization key that is required for updating a file that is referenced in a WRITE PERMISSION ADMIN DATALINK column.

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 might need to be corrected.

WTO. See “write to operator”.

WTOR. A write to operator (WTO) with reply.

X

XBSA. An industry-standard API set for backup and restore utilities. XBSA is one of the archive options that is available for use in maintaining backup copies of linked files in the DB2 Data Links Manager environment. The XBSA option is specified with the DLFM_BACKUP_TARGET registry variable.

XCF. See “cross-system coupling facility” on page 21.

XES. See “cross-system extended services” on page 21.

XID. Exchange station ID.

XML. See “extensible markup language” on page 39.

XML collection. A collection of relational tables from which XML documents are decomposed or that form content of XML documents to be decomposed.

XML column. A column that has a type of XML Extender user-defined type and the contents of the column are whole XML documents.

XML element. Logical structures in XML documents delimited by a start and an end tag. An element can be specified in the DTD by an element type declaration

XML Shredder. A function that parses an XML document, extracting rows of data from an XML table.

XRF. See “extended recovery facility” on page 39.

Z

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

Appendix A. DB2 Universal Database technical information

Overview of DB2 Universal Database technical information

DB2 Universal Database technical information can be obtained in the following formats:

- Books (PDF and hard-copy formats)
- A topic tree (HTML format)
- Help for DB2 tools (HTML format)
- Sample programs (HTML format)
- Command line help
- Tutorials

This section is an overview of the technical information that is provided and how you can access it.

FixPaks for DB2 documentation

IBM may periodically make documentation FixPaks available. Documentation FixPaks allow you to update the information that you installed from the *DB2 HTML Documentation CD* as new information becomes available.

Note: If you do install documentation FixPaks, your HTML documentation will contain more recent information than either the DB2 printed or online PDF manuals.

Categories of DB2 technical information

The DB2 technical information is categorized by the following headings:

- Core DB2 information
- Administration information
- Application development information
- Business intelligence information
- DB2 Connect information
- Getting started information
- Tutorial information
- Optional component information
- Release notes

The following tables describe, for each book in the DB2 library, the information needed to order the hard copy, print or view the PDF, or locate the HTML directory for that book. A full description of each of the books in

the DB2 library is available from the IBM Publications Center at www.ibm.com/shop/publications/order

The installation directory for the HTML documentation CD differs for each category of information:

htmlcdpath/doc/htmlcd/*%L*/*category*

where:

- *htmlcdpath* is the directory where the HTML CD is installed.
- *%L* is the language identifier. For example, en_US.
- *category* is the category identifier. For example, core for the core DB2 information.

In the PDF file name column in the following tables, the character in the sixth position of the file name indicates the language version of a book. For example, the file name db2d1e80 identifies the English version of the *Administration Guide: Planning* and the file name db2d1g80 identifies the German 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
Arabic	w
Brazilian Portuguese	b
Bulgarian	u
Croatian	9
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
Romanian	8
Russian	r
Simp. Chinese	c
Slovakian	7
Slovenian	l

Spanish	z
Swedish	s
Trad. Chinese	t
Turkish	m

No form number indicates that the book is only available online and does not have a printed version.

Core DB2 information

The information in this category covers DB2 topics that are fundamental to all DB2 users. You will find the information in this category useful whether you are a programmer, a database administrator, or you work with DB2 Connect, DB2 Warehouse Manager, or other DB2 products.

The installation directory for this category is doc/htmlcd/*%L*/core.

Table 1. Core DB2 information

Name	Form Number	PDF File Name
IBM DB2 Universal Database Command Reference	SC09-4828	db2n0x80
IBM DB2 Universal Database Glossary	No form number	db2t0x80
IBM DB2 Universal Database Master Index	SC09-4839	db2w0x80
IBM DB2 Universal Database Message Reference, Volume 1	GC09-4840	db2m1x80
IBM DB2 Universal Database Message Reference, Volume 2	GC09-4841	db2m2x80
IBM DB2 Universal Database What's New	SC09-4848	db2q0x80

Administration information

The information in this category covers those topics required to effectively design, implement, and maintain DB2 databases, data warehouses, and federated systems.

The installation directory for this category is doc/htmlcd/*%L*/admin.

Table 2. Administration information

Name	Form number	PDF file name
IBM DB2 Universal Database Administration Guide: Planning	SC09-4822	db2d1x80
IBM DB2 Universal Database Administration Guide: Implementation	SC09-4820	db2d2x80
IBM DB2 Universal Database Administration Guide: Performance	SC09-4821	db2d3x80
IBM DB2 Universal Database Administrative API Reference	SC09-4824	db2b0x80
IBM DB2 Universal Database Data Movement Utilities Guide and Reference	SC09-4830	db2dmx80
IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference	SC09-4831	db2hax80
IBM DB2 Universal Database Data Warehouse Center Administration Guide	SC27-1123	db2ddx80
IBM DB2 Universal Database Federated Systems Guide	GC27-1224	db2fpx80
IBM DB2 Universal Database Guide to GUI Tools for Administration and Development	SC09-4851	db2atx80
IBM DB2 Universal Database Replication Guide and Reference	SC27-1121	db2e0x80
IBM DB2 Installing and Administering a Satellite Environment	GC09-4823	db2dsx80
IBM DB2 Universal Database SQL Reference, Volume 1	SC09-4844	db2s1x80
IBM DB2 Universal Database SQL Reference, Volume 2	SC09-4845	db2s2x80
IBM DB2 Universal Database System Monitor Guide and Reference	SC09-4847	db2f0x80

Application development information

The information in this category is of special interest to application developers or programmers working with DB2. You will find information about supported languages and compilers, as well as the documentation required to access DB2 using the various supported programming interfaces, such as embedded SQL, ODBC, JDBC, SQLj, and CLI. If you view this information online in HTML you can also access a set of DB2 sample programs in HTML.

The installation directory for this category is doc/htmlcd/*%L*/ad.

Table 3. Application development information

Name	Form number	PDF file name
IBM DB2 Universal Database Application Development Guide: Building and Running Applications	SC09-4825	db2axx80
IBM DB2 Universal Database Application Development Guide: Programming Client Applications	SC09-4826	db2a1x80
IBM DB2 Universal Database Application Development Guide: Programming Server Applications	SC09-4827	db2a2x80
IBM DB2 Universal Database Call Level Interface Guide and Reference, Volume 1	SC09-4849	db2l1x80
IBM DB2 Universal Database Call Level Interface Guide and Reference, Volume 2	SC09-4850	db2l2x80
IBM DB2 Universal Database Data Warehouse Center Application Integration Guide	SC27-1124	db2adx80
IBM DB2 XML Extender Administration and Programming	SC27-1234	db2sxx80

Business intelligence information

The information in this category describes how to use components that enhance the data warehousing and analytical capabilities of DB2 Universal Database.

The installation directory for this category is doc/htmlcd/*%L*/wareh.

Table 4. Business intelligence information

Name	Form number	PDF file name
IBM DB2 Warehouse Manager Information Catalog Center Administration Guide	SC27-1125	db2dix80
IBM DB2 Warehouse Manager Installation Guide	GC27-1122	db2idx80

DB2 Connect information

The information in this category describes how to access host or iSeries data using DB2 Connect Enterprise Edition or DB2 Connect Personal Edition.

The installation directory for this category is doc/htmlcd/*%L*/conn.

Table 5. DB2 Connect information

Name	Form number	PDF file name
APPC, CPI-C, and SNA Sense Codes	No form number	db2apx80
IBM Connectivity Supplement	No form number	db2h1x80
IBM DB2 Connect Quick Beginnings for DB2 Connect Enterprise Edition	GC09-4833	db2c6x80
IBM DB2 Connect Quick Beginnings for DB2 Connect Personal Edition	GC09-4834	db2c1x80
IBM DB2 Connect User's Guide	SC09-4835	db2c0x80

Getting started information

The information in this category is useful when you are installing and configuring servers, clients, and other DB2 products.

The installation directory for this category is doc/htmlcd/*%L*/start.

Table 6. Getting started information

Name	Form number	PDF file name
IBM DB2 Universal Database Quick Beginnings for DB2 Clients	GC09-4832	db2itx80

Table 6. Getting started information (continued)

Name	Form number	PDF file name
IBM DB2 Universal Database Quick Beginnings for DB2 Servers	GC09-4836	db2isx80
IBM DB2 Universal Database Quick Beginnings for DB2 Personal Edition	GC09-4838	db2i1x80
IBM DB2 Universal Database Installation and Configuration Supplement	GC09-4837	db2iyx80
IBM DB2 Universal Database Quick Beginnings for DB2 Data Links Manager	GC09-4829	db2z6x80

Tutorial information

Tutorial information introduces DB2 features and teaches how to perform various tasks.

The installation directory for this category is doc/htmlcd/*%L*/tutr.

Table 7. Tutorial information

Name	Form number	PDF file name
Business Intelligence Tutorial: Introduction to the Data Warehouse	No form number	db2tux80
Business Intelligence Tutorial: Extended Lessons in Data Warehousing	No form number	db2tax80
Development Center Tutorial for Video Online using Microsoft Visual Basic	No form number	db2tdx80
Information Catalog Center Tutorial	No form number	db2aix80
Video Central for e-business Tutorial	No form number	db2twx80
Visual Explain Tutorial	No form number	db2tvx80

Optional component information

The information in this category describes how to work with optional DB2 components.

The installation directory for this category is doc/htmlcd/*%L*/opt.

Table 8. Optional component information

Name	Form number	PDF file name
IBM DB2 Life Sciences Data Connect Planning, Installation, and Configuration Guide	GC27-1235	db2lsx80
IBM DB2 Spatial Extender User's Guide and Reference	SC27-1226	db2sbx80
IBM DB2 Universal Database Data Links Manager Administration Guide and Reference	SC27-1221	db2z0x80
IBM DB2 Universal Database Net Search Extender Administration and Programming Guide **Note:** HTML for this document is not installed from the HTML documentation CD.	SH12-6740	N/A

Release notes

The release notes provide additional information specific to your product's release and FixPak level. They also provides summaries of the documentation updates incorporated in each release and FixPak.

Table 9. Release notes

Name	Form number	PDF file name
DB2 Release Notes	See note.	See note.
DB2 Installation Notes	Available on product CD-ROM only.	Available on product CD-ROM only.

Note: The HTML version of the release notes 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/opt/db2_08_01 on AIX
- /opt/IBM/db2/V8.1 on all other UNIX operating systems

Printing DB2 books from PDF files

You can print DB2 books from the PDF files on the *DB2 PDF Documentation* CD. Using Adobe Acrobat Reader, you can print either the entire book or a specific range of pages.

Prerequisites:

Ensure that you have Adobe Acrobat Reader. It is available from the Adobe Web site at www.adobe.com

Procedure:

To print a DB2 book from a PDF file:

1. Insert the *DB2 PDF Documentation* CD. On UNIX operating systems, mount the DB2 PDF Documentation CD. Refer to your *Quick Beginnings* book for details on how to mount a CD on UNIX operating systems.
2. Start Adobe Acrobat Reader.
3. Open the PDF file from one of the following locations:
   - On Windows operating systems:

     x:\doc*language* directory, where *x* represents the CD-ROM drive letter and *language* represents the two-character territory code that represents your language (for example, EN for English).
   - On UNIX operating systems:

     /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.

Ordering printed DB2 books

Procedure:

To order printed books:

- Contact your IBM authorized dealer or marketing representative. To find a local IBM representative, check the IBM Worldwide Directory of Contacts at www.ibm.com/planetwide
- Phone 1-800-879-2755 in the United States or 1-800-IBM-4YOU in Canada.
- Visit the IBM Publications Center at www.ibm.com/shop/publications/order

You can also obtain printed DB2 manuals by ordering Doc Packs for your DB2 product from your IBM Reseller. The Doc Packs are subsets of the manuals in the DB2 library selected to help you to get started using the DB2 product that

you purchased. The manuals in the Doc Packs are the same as those that are available in PDF format on the *DB2 PDF Documentation CD* and contain the same content as the documentation that is available on the *DB2 HTML Documentation CD*.

Accessing online help

The online help that comes with all DB2 components is available in three types:

- Window and notebook help
- Command line help
- SQL statement help

Window and notebook help explain the tasks that you can perform in a window or notebook and describe the controls. This help has two types:

- Help accessible from the **Help** button
- Infopops

The **Help** button gives you access to overview and prerequisite information. The infopops describe the controls in the window or notebook. Window and notebook help are available from DB2 centers and components that have user interfaces.

Command line help includes Command help and Message help. Command help explains the syntax of commands in the command line processor. Message help describes the cause of an error message and describes any action you should take in response to the error.

SQL statement help includes SQL help and SQLSTATE help. DB2 returns an SQLSTATE value for conditions that could be the result of an SQL statement. SQLSTATE help explains the syntax of SQL statements (SQL states and class codes).

Note: SQL help is not available for UNIX operating systems.

Procedure:

To access online help:

- For window and notebook help, click **Help** or click that control, then click **F1**. If the **Automatically display infopops** check box on the **General** page of the **Tool Settings** notebook is selected, you can also see the infopop for a particular control by holding the mouse cursor over the control.
- For command line help, open the command line processor and enter:
  - For Command help:

```
? 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.

- For Message help:

  ```
  ? XXXnnnnn
  ```

  where *XXXnnnnn* represents a valid message identifier.

  For example, `? SQL30081` displays help about the SQL30081 message.

- For SQL statement help, open the command line processor and 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.

Finding topics by accessing the DB2 Information Center from a browser

The DB2 Information Center accessed from a browser enables you to access the information you need to take full advantage of DB2 Universal Database and DB2 Connect. The DB2 Information Center also documents major DB2 features and components including replication, data warehousing, metadata, and DB2 extenders.

The DB2 Information Center accessed from a browser is composed of the following major elements:

Navigation tree
The navigation tree is located in the left frame of the browser window. The tree expands and collapses to show and hide topics, the glossary, and the master index in the DB2 Information Center.

Navigation toolbar
The navigation toolbar is located in the top right frame of the browser window. The navigation toolbar contains buttons that enable you to search the DB2 Information Center, hide the navigation tree, and find the currently displayed topic in the navigation tree.

Content frame
The content frame is located in the bottom right frame of the browser window. The content frame displays topics from the DB2 Information

Center when you click on a link in the navigation tree, click on a search result, or follow a link from another topic or from the master index.

Prerequisites:

To access the DB2 Information Center from a browser, you must use one of the following browsers:

- Microsoft Explorer, version 5 or later
- Netscape Navigator, version 6.1 or later

Restrictions:

The DB2 Information Center contains only those sets of topics that you chose to install from the *DB2 HTML Documentation CD*. If your Web browser returns a `File not found` error when you try to follow a link to a topic, you must install one or more additional sets of topics from the *DB2 HTML Documentation CD*.

Procedure:

To find a topic by searching with keywords:

1. In the navigation toolbar, click **Search**.
2. In the top text entry field of the Search window, enter one or more terms related to your area of interest and click **Search**. A list of topics ranked by accuracy displays in the **Results** field. The numerical ranking beside the hit provides an indication of the strength of the match (bigger numbers indicate stronger matches).

   Entering more terms increases the precision of your query while reducing the number of topics returned from your query.
3. In the **Results** field, click the title of the topic you want to read. The topic displays in the content frame.

To find a topic in the navigation tree:

1. In the navigation tree, click the book icon of the category of topics related to your area of interest. A list of subcategories displays underneath the icon.
2. Continue to click the book icons until you find the category containing the topics in which you are interested. Categories that link to topics display the category title as an underscored link when you move the cursor over the category title. The navigation tree identifies topics with a page icon.
3. Click the topic link. The topic displays in the content frame.

To find a topic or term in the master index:

1. In the navigation tree, click the “Index” category. The category expands to display a list of links arranged in alphabetical order in the navigation tree.
2. In the navigation tree, click the link corresponding to the first character of the term relating to the topic in which you are interested. A list of terms with that initial character displays in the content frame. Terms that have multiple index entries are identified by a book icon.
3. Click the book icon corresponding to the term in which you are interested. A list of subterms and topics displays below the term you clicked. Topics are identified by page icons with an underscored title.
4. Click on the title of the topic that meets your needs. The topic displays in the content frame.

Finding product information by accessing the DB2 Information Center from the administration tools

The DB2 Information Center provides quick access to DB2 product information and is available on all operating systems for which the DB2 administration tools are available.

The DB2 Information Center accessed from the tools provides six types of information.

Tasks Key tasks you can perform using DB2.

Concepts
Key concepts for DB2.

Reference
DB2 reference information, such as keywords, commands, and APIs.

Troubleshooting
Error messages and information to help you with common DB2 problems.

Samples
Links to HTML listings of the sample programs provided with DB2.

Tutorials
Instructional aid designed to help you learn a DB2 feature.

Prerequisites:

Some links in the DB2 Information Center point to Web sites on the Internet. To display the content for these links, you will first have to connect to the Internet.

Procedure:

To find product information by accessing the DB2 Information Center from the tools:

1. Start the DB2 Information Center in one of the following ways:
   - From the graphical administration tools, click on the **Information Center** icon in the toolbar. You can also select it from the **Help** menu.
   - At the command line, enter **db2ic**.
2. Click the tab of the information type related to the information you are attempting to find.
3. Navigate through the tree and click on the topic in which you are interested. The Information Center will then launch a Web browser to display the information.
4. To find information without browsing the lists, click the **Search** icon to the right of the list.

   Once the Information Center has launched a browser to display the information, you can perform a full-text search by clicking the **Search** icon in the navigation toolbar.

Viewing technical documentation online directly from the DB2 HTML Documentation CD

All of the HTML topics that you can install from the *DB2 HTML Documentation CD* can also be read directly from the CD. Therefore, you can view the documentation without having to install it.

Restrictions:

As the Tools help is installed from the DB2 product CD and not from the *DB2 HTML Documentation CD*, you must install the DB2 product to view the help.

Procedure:

1. Insert the *DB2 HTML Documentation* CD. On UNIX operating systems, mount the *DB2 HTML Documentation CD*. Refer to your *Quick Beginnings* book for details on how to mount a CD on UNIX operating systems.
2. Start your HTML browser and open the appropriate file:
   - For Windows operating systems:

     ```
     e:\program files\IBM\SQLLIB\doc\htmlcd\%L\index.htm
     ```

     where *e* represents the CD-ROM drive, and %L is the locale of the documentation that you wish to use, for example, **en_US** for English.
   - For UNIX operating systems:

     ```
     /cdrom/program files/IBM/SQLLIB/doc/htmlcd/%L/index.htm
     ```

where */cdrom/* represents where the CD is mounted, and %L is the locale of the documentation that you wish to use, for example, **en_US** for English.

Updating the HTML documentation installed on your machine

It is now possible to update the HTML installed from the *DB2 HTML Documentation CD* when updates are made available from IBM. This can be done in one of two ways:

- Using the Information Center (if you have the DB2 administration GUI tools installed).
- By downloading and applying a DB2 HTML documentation FixPak .

Note: This will NOT update the DB2 code; it will only update the HTML documentation installed from the *DB2 HTML Documentation CD*.

Procedure:

To use the Information Center to update your local documentation:

1. Start the DB2 Information Center in one of the following ways:
   - From the graphical administration tools, click on the **Information Center** icon in the toolbar. You can also select it from the **Help** menu.
   - At the command line, enter **db2ic**.
2. Ensure your machine has access to the external Internet; the updater will download the latest documentation FixPak from the IBM server if required.
3. Select **Information Center** —> **Update Local Documentation** from the menu to start the update.
4. Supply your proxy information (if required) to connect to the external Internet.

The Information Center will download and apply the latest documentation FixPak, if one is available.

To manually download and apply the documentation FixPak :

1. Ensure your machine is connected to the Internet.
2. Open the DB2 support page in your Web browser at: www.ibm.com/software/data/db2/udb/winos2unix/support.
3. Follow the link for Version 8 and look for the ″Documentation FixPaks″ link.

4. Determine if the version of your local documentation is out of date by comparing the documentation FixPak level to the documentation level you have installed. This current documentation on your machine is at the following level: **DB2 v8.1 GA**.
5. If there is a more recent version of the documentation available then download the FixPak applicable to your operating system. There is one FixPak for all Windows platforms, and one FixPak for all UNIX platforms.
6. Apply the FixPak:
   - For Windows operating systems: The documentation FixPak is a self extracting zip file. Place the downloaded documentation FixPak in an empty directory, and run it. It will create a **setup** command which you can run to install the documentation FixPak.
   - For UNIX operating systems: The documentation FixPak is a compressed tar.Z file. Uncompress and untar the file. It will create a directory named `delta_install` with a script called **installdocfix**. Run this script to install the documentation FixPak.

Copying files from the DB2 HTML Documentation CD to a Web server

The entire DB2 information library is delivered to you on the *DB2 HTML Documentation CD* and may be installed on a Web server for easier access. Simply copy to your Web server the documentation for the languages that you want.

Note: You might encounter slow performance if you access the HTML documentation from a Web server through a low-speed connection.

Procedure:

To copy files from the *DB2 HTML Documentation CD* to a Web server, use the appropriate source path:

- For Windows operating systems:

  E:\program files\IBM\SQLLIB\doc\htmlcd*%L**.*

  where *E* represents the CD-ROM drive and *%L* represents the language identifier.
- For UNIX operating systems:

  /cdrom/program files/IBM/SQLLIB/doc/htmlcd/*%L*/*.*

  where *cdrom* represents the mount point for the CD-ROM drive and *%L* represents the language identifier.

Troubleshooting DB2 documentation search with Netscape 4.x

Most search problems are related to the Java support provided by web browsers. This task describes possible workarounds.

Procedure:

A common problem with Netscape 4.x involves a missing or misplaced security class. Try the following workaround, especially if you see the following line in the browser Java console:

```
Cannot find class  java/security/InvalidParameterException
```

- On Windows operating systems:

  From the *DB2 HTML Documentation CD*, copy the supplied *x*:`program files\IBM\SQLLIB\doc\htmlcd\`*locale*`\InvalidParameterException.class` file to the `java\classes\java\security\` directory relative to your Netscape browser installation, where *x* represents the CD-ROM drive letter and *locale* represents the name of the desired locale.

  Note: You may have to create the `java\security\` subdirectory structure.

- On UNIX operating systems:

  From the *DB2 HTML Documentation CD*, copy the supplied */cdrom*`/program files/IBM/SQLLIB/doc/htmlcd/`*locale*`/InvalidParameterException.class` file to the `java/classes/java/security/` directory relative to your Netscape browser installation, where *cdrom* represents the mount point of the CD-ROM and *locale* represents the name of the desired locale.

  Note: You may have to create the `java/security/` subdirectory structure.

If your Netscape browser still fails to display the search input window, try the following:

- Stop all instances of Netscape browsers to ensure that there is no Netscape code running on the machine. Then open a new instance of the Netscape browser and try to start the search again.
- Purge the browser's cache.
- Try a different version of Netscape, or a different browser.

Searching the DB2 documentation

You can search the library of DB2 documentation to locate information that you need. A pop-up search window opens when you click the search icon in the navigation toolbar of the DB2 Information Center (accessed from a browser). The search can take a minute to load, depending on the speed of your computer and network.

Prerequisites:

You need Netscape 6.1 or higher, or Microsoft's Internet Explorer 5 or higher. Ensure that your browser's Java support is enabled.

Restrictions:

The following restrictions apply when you use the documentation search:

- Search is not case sensitive.
- Boolean searches are not supported.
- Wildcard and partial searches are not supported. A search on *java** (or *java*) will only look for the literal string *java** (or *java*) and would not, for example, find *javadoc*.

Procedure:

To search the DB2 documentation:

1. In the navigation toolbar, click the **Search** icon.
2. In the top text entry field of the Search window, enter one or more terms (separated by a space) related to your area of interest and click **Search**. A list of topics ranked by accuracy displays in the **Results** field. The numerical ranking beside the hit provides an indication of the strength of the match (bigger numbers indicate stronger matches).

   Entering more terms increases the precision of your query while reducing the number of topics returned from your query.
3. In the **Results** list, click the title of the topic you want to read. The topic displays in the content frame of the DB2 Information Center.

Note: When you perform a search, the first (highest-ranking) result is automatically loaded into your browser frame. To view the contents of other search results, click on the result in the results list.

Online DB2 troubleshooting information

With the release of DB2® UDB Version 8, there will no longer be a *Troubleshooting Guide*. The troubleshooting information once contained in this guide has been integrated into the DB2 publications. By doing this, we are able to deliver the most up-to-date information possible. To find information on the troubleshooting utilities and functions of DB2, access the DB2 Information Center from any of the tools.

Refer to the DB2 Online Support site if you are experiencing problems and want help finding possible causes and solutions. The support site contains a large, constantly updated database of DB2 publications, TechNotes, APAR

(product problem) records, FixPaks, and other resources. You can use the support site to search through this knowledge base and find possible solutions to your problems.

Access the Online Support site at www.ibm.com/software/data/db2/udb/winos2unix/support, or by clicking the **Online Support** button in the DB2 Information Center. Frequently changing information, such as the listing of internal DB2 error codes, is now also available from this site.

Accessibility

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. These are the major accessibility features in DB2® Universal Database Version 8:

- DB2 allows you to operate all features using the keyboard instead of the mouse. See "Keyboard Input and Navigation".
- DB2 enables you customize the size and color of your fonts. See "Accessible Display".
- DB2 allows you to receive either visual or audio alert cues. See "Alternative Alert Cues" on page 130.
- DB2 supports accessibility applications that use the Java™ Accessibility API. See "Compatibility with Assistive Technologies" on page 130.
- DB2 comes with documentation that is provided in an accessible format. See "Accessible Documentation" on page 130.

Keyboard Input and Navigation

Keyboard Input

You can operate the DB2 Tools using only the keyboard. You can use keys or key combinations to perform most operations that can also be done using a mouse.

Keyboard Focus

In UNIX-based systems, the position of the keyboard focus is highlighted, indicating which area of the window is active and where your keystrokes will have an effect.

Accessible Display

The DB2 Tools have features that enhance the user interface and improve accessibility for users with low vision. These accessibility enhancements include support for customizable font properties.

Font Settings

The DB2 Tools allow you to select the color, size, and font for the text in menus and dialog windows, using the Tools Settings notebook.

Non-dependence on Color

You do not need to distinguish between colors in order to use any of the functions in this product.

Alternative Alert Cues

You can specify whether you want to receive alerts through audio or visual cues, using the Tools Settings notebook.

Compatibility with Assistive Technologies

The DB2 Tools interface supports the Java Accessibility API enabling use by screen readers and other assistive technologies used by people with disabilities.

Accessible Documentation

Documentation for the DB2 family of products is available in HTML format. This allows you to view documentation according to the display preferences set in your browser. It also allows you to use screen readers and other assistive technologies.

DB2 tutorials

The DB2® tutorials help you learn about various aspects of DB2 Universal Database. The tutorials provide lessons with step-by-step instructions in the areas of developing applications, tuning SQL query performance, working with data warehouses, managing metadata, and developing Web services using DB2.

Before you begin:

Before you can access these tutorials using the links below, you must install the tutorials from the *DB2 HTML Documentation* CD.

If you do not want to install the tutorials, you can view the HTML versions of the tutorials directly from the *DB2 HTML Documentation CD*. PDF versions of these tutorials are also available on the *DB2 PDF Documentation CD*.

Some tutorial lessons use sample data or code. See each individual tutorial for a description of any prerequisites for its specific tasks.

DB2 Universal Database tutorials:

If you installed the tutorials from the *DB2 HTML Documentation* CD, you can click on a tutorial title in the following list to view that tutorial.

Business Intelligence Tutorial: Introduction to the Data Warehouse Center
Perform introductory data warehousing tasks using the Data Warehouse Center.

Business Intelligence Tutorial: Extended Lessons in Data Warehousing
Perform advanced data warehousing tasks using the Data Warehouse Center.

Development Center Tutorial for Video Online using Microsoft® Visual Basic
Build various components of an application using the Development Center Add-in for Microsoft Visual Basic.

Information Catalog Center Tutorial
Create and manage an information catalog to locate and use metadata using the Information Catalog Center.

Video Central for e-business Tutorial
Develop and deploy an advanced DB2 Web Services application using WebSphere® products.

Visual Explain Tutorial
Analyze, optimize, and tune SQL statements for better performance using Visual Explain.

DB2 Information Center accessed from a browser

The DB2® Information Center gives you access to all of the information you need to take full advantage of DB2 Universal Database™ and DB2 Connect™ in your business. The DB2 Information Center also documents major DB2 features and components including replication, data warehousing, the Information Catalog Center, Life Sciences Data Connect, and DB2 extenders.

The DB2 Information Center accessed from a browser has the following features if you view it in Netscape Navigator 6.1 or later or Microsoft Internet Explorer 5 or later. Some features require you to enable support for Java or JavaScript:

Regularly updated documentation
Keep your topics up-to-date by downloading updated HTML.

Search
Search all of the topics installed on your workstation by clicking **Search** in the navigation toolbar.

Integrated navigation tree
Locate any topic in the DB2 library from a single navigation tree. The navigation tree is organized by information type as follows:

- Tasks provide step-by-step instructions on how to complete a goal.
- Concepts provide an overview of a subject.
- Reference topics provide detailed information about a subject, including statement and command syntax, message help, requirements.

Master index

Access the information installed from the *DB2 HTML Documentation CD* from the master index. The index is organized in alphabetical order by index term.

Master glossary

The master glossary defines terms used in the DB2 Information Center. The glossary is organized in alphabetical order by glossary term.

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/region 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/region 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 that has been exchanged, should contact:

IBM Canada Limited
Office of the Lab Director
8200 Warden Avenue
Markham, Ontario
L6G 1C7
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 document 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 illustrate 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 are trademarks of International Business Machines Corporation in the United States, other countries, or both, and have been used in at least one of the documents in the DB2 UDB documentation library.

ACF/VTAM
AISPO
AIX
AIXwindows
AnyNet
APPN
AS/400
BookManager
C Set++
C/370
CICS
Database 2
DataHub
DataJoiner
DataPropagator
DataRefresher
DB2
DB2 Connect
DB2 Extenders
DB2 OLAP Server
DB2 Universal Database
Distributed Relational Database Architecture
DRDA
eServer
Extended Services
FFST
First Failure Support Technology
IBM
IMS
IMS/ESA
iSeries
LAN Distance
MVS
MVS/ESA
MVS/XA
Net.Data
NetView
OS/390
OS/400
PowerPC
pSeries
QBIC
QMF
RACF
RISC System/6000
RS/6000
S/370
SP
SQL/400
SQL/DS
System/370
System/390
SystemView
Tivoli
VisualAge
VM/ESA
VSE/ESA
VTAM
WebExplorer
WebSphere
WIN-OS/2
z/OS
zSeries

The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the documents in the DB2 UDB documentation library:

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Contacting IBM

In the United States, call one of the following numbers to contact IBM:

- 1-800-237-5511 for customer service
- 1-888-426-4343 to learn about available service options
- 1-800-IBM-4YOU (426-4968) for DB2 marketing and sales

In Canada, call one of the following numbers to contact IBM:

- 1-800-IBM-SERV (1-800-426-7378) for customer service
- 1-800-465-9600 to learn about available service options
- 1-800-IBM-4YOU (1-800-426-4968) for DB2 marketing and sales

To locate an IBM office in your country or region, check IBM's Directory of Worldwide Contacts on the web at www.ibm.com/planetwide

Product information

Information regarding DB2 Universal Database products is available by telephone or by the World Wide Web at www.ibm.com/software/data/db2/udb

This site contains the latest information on the technical library, ordering books, client downloads, newsgroups, FixPaks, news, and links to web resources.

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) to order products or to obtain general information.
- 1-800-879-2755 to order publications.

For information on how to contact IBM outside of the United States, go to the IBM Worldwide page at www.ibm.com/planetwide

Printed in U.S.A.

Spine information:

IBM® DB2 Universal Database™ DB2 Glossary Version 8